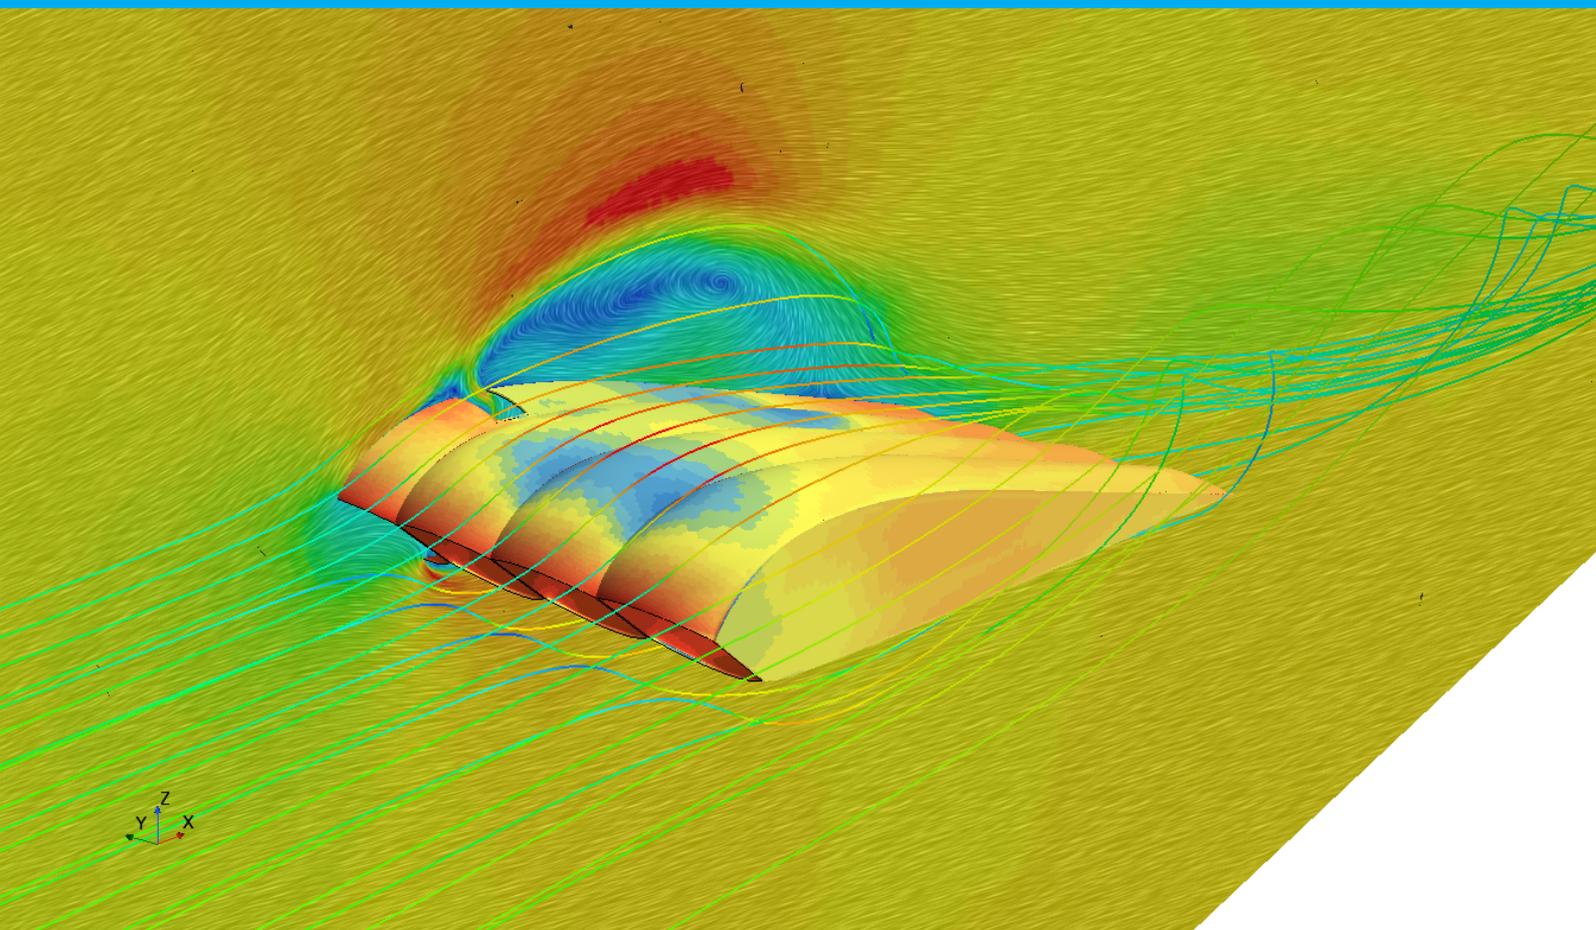Master of Science Thesis

Bleed Air Spoiler Optimisation for AWE Purposes

An aerodynamic optimisation of upper surface bleed air control to minimise reel-in power loss for soft-kite-based Airborne Wind Energy systems

G. J. Fimpel

July 31, 2021

TUDelft   DTU   EWEM   **SkySails POWER**

Bleed Air Spoiler Optimisation for AWE Purposes

An aerodynamic optimisation of upper surface bleed air control to minimise reel-in power loss for soft-kite-based Airborne Wind Energy systems

by

G. J. Fimpel

to obtain the degree of Master of Science at the Delft University of Technology and the Danish University of Technology. To be defended publicly on August 12, 2021, at 10:00 am.

This thesis is confidential and cannot be made public until Aug 1, 2023.

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

Abstract

A key challenge in the Airborne Wind Energy industry is the minimisation of the time and energy required for the transfer and return phases inherent to the power cycles of ground-generator-based kite power systems. This thesis contributes to solving this problem by exploring the use of bleed air spoilers during the return phase in terms of their potential to improve the average cycle performance. To this end, a mathematical model is first introduced to evaluate the individual phases of the characteristic power pumping cycles in a general sense. This model is further used to demonstrate that a given kite configuration with the capacity to selectively lower its glide slope and/or resultant aerodynamic loading could improve the net power output of the system by around 30%. A CFD-based simulation environment for investigations of the effects induced by bleed air spoilers is developed using Simcenter STAR-CCM+. To establish the reliability of the setup, a first study revolves around its validation with experimental data from independent wind tunnel tests on bleed air spoilers effects on ram-air parachutes. The second CFD study is based on a standard kite geometry employed by SkySails Power GmbH. Two parameter studies testing the influence of spoiler size and chordwise position are carried out to investigate how bleed air spoilers are best used to reach the desired aerodynamic effects. The results show that bleed-air spoilers should ideally be applied in the vicinity of the suction peak and that relatively small spoiler geometries can already entail a major effect on the resultant flow field. It is predicted that the evaluated configurations would improve the system's overall performance by 9 to 18.3%, mostly due to substantial reductions in the glide ratio.

Acknowledgements

Contents

List of Tables

List of Figures

Nomenclature

Abbreviations

AWE	Airborne wind energy
BAS	Bleed air spoiler
CFD	Computational fluid dynamics
CV	Control volume
ERCOFTAC	European Research Community on Flow, Turbulence and Combustion
GCI	Grid convergence index
NEG	Net energy gain
RANS	Reynolds-averaged Navier–Stokes equations
SST	Shear stress transport
TKE	Turbulence kinetic energy

Greek letters

α	Angle of attack
Δ	Difference or change in
δ	Boundary layer thickness
ϵ	rate of dissipation of turbulence kinetic energy
μ	Dynamic viscosity
μ	Reeling factor, $\dot{l} = \mu v_w$
ν	Kinematic viscosity, $\nu = \frac{\mu}{\rho}$
ν_t	turbulence (/eddy) viscosity
ω	Specific rate of dissipation of turbulence kinetic energy
ϕ	Azimuth angle
ϕ	Monitoring variable for grid refinement studies
ψ	Angle of the kite around its yaw axis to the wind
ρ	Air density
θ	Elevation angle
ϑ	Spatial angle between the x-axis (wind direction) and the connecting line between the ground station and the kite

Tensor and vector notations

Δ	Laplace operator
δ_{ij}	Kronecker delta

$\mathbf{v}$	Velocities, $\mathbf{v} = \left(v_x, v_y, v_z\right)^{\mathrm{T}}$
∇	Nabla operator
Φ	Arbitrary tensor-, vector-, or scalar-based property of a substance

Other symbols

$\dot{l}$	Reeling (winch out/in) speed
b	Half span
c	Chord length
$C_{\mathrm{onset,1}}$	Transition onset parameter
C_D	Drag coefficient
C_L	Lift coefficient
C_M	Pitch moment coefficient around control pod
C_P	Pressure coefficient, $C_P = \frac{p - p_\infty}{q_\infty}$
C_R	Resultant aerodynamic loading coefficient, $C_R = \sqrt{C_L^2 + C_D^2}$
Co	Courant number
E	Glide ratio, $E = \frac{L}{D}$
h	Representative cell size
I	Turbulence intensity
k	Turbulence kinetic energy, $k = \frac{1}{2}\overline{v_i' v_i'}$
L	Characteristic length of a flow
l	Tether length
Ma	Mach number
N	Number of cells in volume mesh
n	Number of prism layers in wall normal direction
p	Spatial discretisation order
p	Static pressure
q	Dynamic pressure
Re	Reynolds number, $Re = \frac{vd}{\nu}$
S	Projected surface area of kite
s	Cell growth rate
T	Total height of prism-layers
t	Time
V	Cell volume
v_∞	Ambient wind speed
v_a	Apparent wind as measured by an on-board anemometer at the control pod
v_k	Path velocity of the kite
y^+	Non-dimensional wall distance, $y^+ = \frac{y u_\tau}{\nu}$

1

Introduction

The energy supply of the world's steadily growing population is arguably one of the central challenges faced by our generation. Although advanced fusion technologies could eventually offer a reliable, abundant source of clean energy, they are not expected to play a significant role in the energy market until 2050 [9, p. 19]. Until then, it is clear that the pursuit of the United Nations' Sustainable Development Goals for affordable and clean energy, sustainability, climate action, and life below water and on land [31] is inherently linked to a balanced energy mix that is based on as many complementing types of renewable resources as possible. Although several established technologies already exist, such as photovoltaics, conventional wind turbines and biomass plants, these still only scratch the surface of the vast powers of nature which surround us. Engineering pioneers hence continue to tap into new sources of renewables in an ever-growing number of ways.

This thesis will focus on a current challenge in one of these new fields: the harnessing of high-altitude winds by means of ground-generator-based tethered kite power systems. Pursued by a handful of companies and institutions worldwide, these plants work on the basis of power pumping cycles, also coined the yo-yo-principle, where the kite is periodically reeled out and retracted. Naturally, the average power output over such a cycle is strongly linked to the time and energy lost for the retraction phase, which continues to pose a pivotal challenge to the industry. SkySails Power GmbH, one of the major drivers of this technology, aims to partially solve this issue by employing bleed air spoilers to adjust the aerodynamic performance of the kite during the return phase appropriately. This study shall investigate how these bleed air spoilers are best implemented to achieve this effect from an aerodynamic viewpoint.

The remainder of this chapter will set the framework for the study. First, Airborne Wind Energy is introduced in an industry context and SkySails currently pursued concept, which forms the basis of the treated evaluations, is described in some more detail. Next, an overview of previous research on bleed air spoiler technologies is given, and relevant conclusions, as well as research gaps for the context of the thesis, are pointed out. After establishing the objective and research question which guided this work, the methodology and thesis outline are presented.

1.1. Airborne Wind Energy

Airborne Wind Energy (AWE) encompasses a novel field of technology that addresses the conversion of wind energy into electricity using automated tethered flying devices. Since first prototypes were tested by visionary pioneers in the early 2000s, research groups across the globe have been exploring the potential of various such approaches, from rigid-wing drones with on-board generators to soft-kite systems which leverage on the traction force of the canopy to drive a winch on the ground that connects to a generator.

Currently, about 20 manufacturers are actively developing AWE technologies, and more than 50 institutions worldwide are involved along the supply chain, as highlighted by the world map in Figure 1.1. To name a few, these include the German companies SkySails Power and EnerKite, the Dutch company KitePower, the technical universities in Berlin, Munich and Stuttgart along with TU Delft and DTU, as well as project developers and utility companies such as EnBW and RWE.
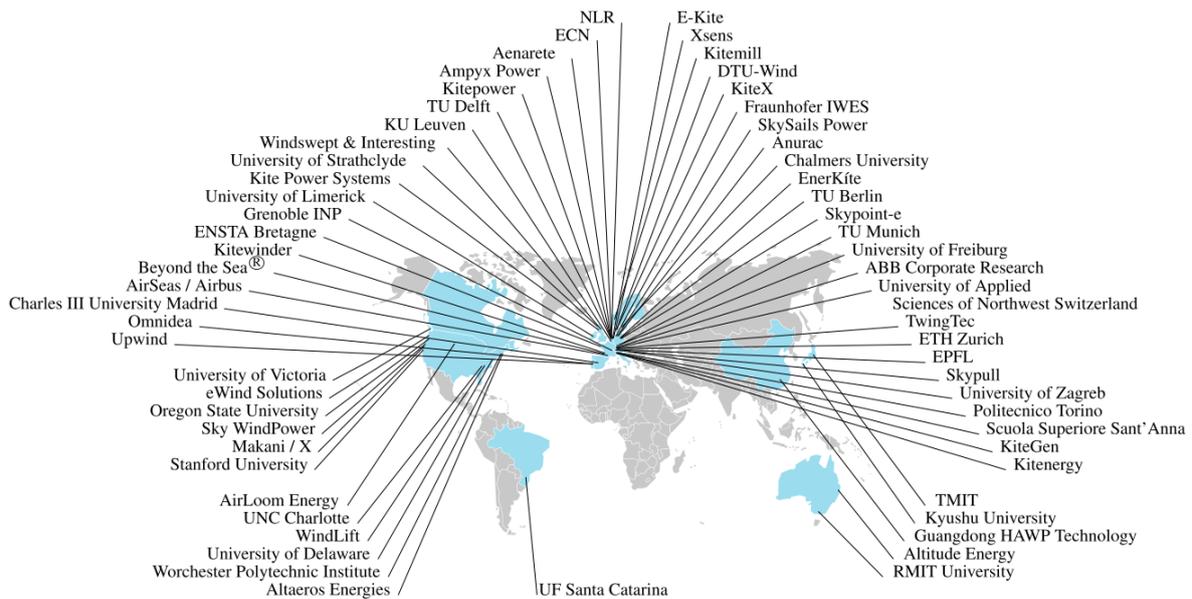
Figure 1.1: Global overview of companies and institutions involved in the AWE sector (state: 2017; from [27, p. 5])

AWE currently appears to evolve into a promising complement to existing renewables in meeting future energy demands thanks to a multitude of key strengths. Whereas state-of-the-art horizontal axis wind turbines (HAWTs) can achieve capacity factors of up to 35% (onshore) or 55% (offshore) [21, p. 18], that of AWE systems is already as high as 69% [27, p. 10] due to stronger and steadier high altitude winds, moving potential electricity production into the base-load range.

HAWTs generate a heap of the mechanical power that drives their generators in the outer region of their rotors, where the blades are slender and relative velocities are largest. Meanwhile, airborne wind energy systems embrace this feature by moving in a circular motion in the air, like the tips of a rotor, only that the massive tower structure and heavy rest of the blades are replaced by the tether. Because the aerodynamic forces acting on the kites are transferred to the ground purely via tensile loads as opposed to the heavy duty structures of HAWTs, kite power systems are very lightweight [27, p. 110], making them much more resource-friendly in comparison.

While continually improving on their respective technologies, companies like SkySails Power have already secured contracts with early adopters, including large European utility companies [26] and remote island nations [29], which put their hopes in AWE to become less reliant on conventional sources of electricity. In these projects, the technology shall mature to a point where it becomes promising enough to scale it further up. This is theoretically possible up to capacities of several MW as demonstrated in [27, p. 296].

Despite the many potential competitive advantages, over 130 years after the erection of the world's first electricity-generating wind turbine in 1887, AWE is entering a market which is already largely dominated by sophisticated, three-bladed HAWTs at a point where the technology is still at a fledgling stage. Whereas the HAWT industry was able to benefit from many years of large-scale subsidies before it became cost effective enough to stand on its own feet, the emerging economical and technological landscape around the AWE sector needs to progress much quicker for it to remain viable. To achieve this, fast paced innovations are of central importance when tackling the remaining challenges.

An overview of the state of the art in Airborne Wind Energy would not be complete without a consideration of these challenges still faced by today's developers. Although the general public has yet to form an opinion on this emerging technology, noise and visual impact will likely be a hot topic in densely populated regions, since the kite's trajectory covers a substantial projected area on the ground. Additionally, AWE systems are still rather susceptible to malfunctions: A wind turbine can easily be stopped in such a case; Airborne Wind Energy systems, on the other hand, run the risk of crashing and being destroyed when an unplanned shutdown occurs. Research efforts therefore also focus on the development of reliable software and control algorithms to ensure dependable autonomous flight capacity in a wide range of weather conditions. Another key hurdle is the persistently high demand for energy during

the return phase of the kite, which narrows the net energy gain of a given power cycle by a significant margin, as further explained in Chapter 3.2. As mentioned before, the contribution to solving this issue is also the main focus of this work.

1.1.1. SkySails' kite power concept

Because the investigations carried out in the context of this thesis are based on the present system design of the company SkySails Power GmbH, their concept is presented in some more detail. For a more elaborate description of the characteristic pumping cycle, by which the system is operated, the reader is again referred to Chapter 3. A labelled illustration of the currently marketed system 'SKS PN-14' is shown in Figure 1.2.

1. Kite

2. Control pod

3. Tether

4. Launch & landing mast

5. Ground station

6. Drive train

7. Tripod with ring mount
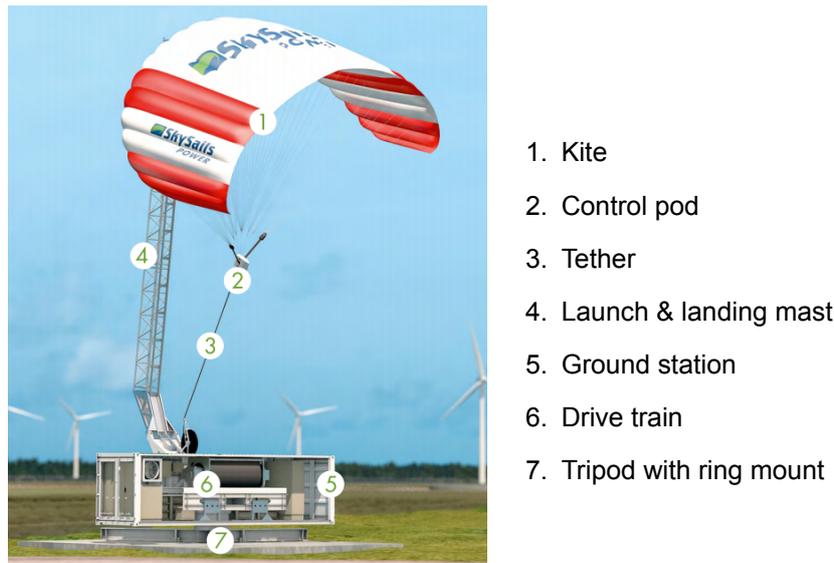
Figure 1.2: Illustration of SkySails' SKS PN-14
(from [30])

As highlighted by the graphic, a majority of the electrical components is housed by a regular 30 ft container, which is installed on a tripod with a ring mount. The latter allows the system to turn itself into the wind, similar to the yaw actuator of a conventional wind turbine. The winch inside the ground station is installed on a movable sled to enable the controlled reeling of up to 800 m of tether length. The tether is led out of an opening on the side of the container and passes over a pulley, the so called towing point, which redirects the forces of the tether appropriately, before it is attached to the control pod.

The control pod is already a remarkable feat of engineering in itself. Since it has no electrical connection to the ground station and yet has to remain airborne for long periods of time, it uses an autonomous energy supply based on an on-board ram-air turbine and a battery pack for buffer storage. Like the human pilot of a ram-air parachute, it can pull on active risers connected to the tips of the canopy to steer it. Since the cogged belt on which it pulls in this process leads through the actuator, only differential forces between the tips have to be overcome which allows for a high degree of control authority and relatively small use of energy. Data from on-board accelerometers and anemometers are continuously transmitted to the ground station, where the autopilot further processes them along with the position of the kite to compute control directives. These directives dictate the imposed reeling speed of the winch and the steering of the kite with the addressed belt.

SkySails Power pursues a soft ram-air kite concept in contrast to other players in the industry, who choose to employ tube kites or rigid-wing drones (e.g. by KitePower and TwingTec, respectively). As such, the kite has two intakes located on the leading edge of the second chamber (counted from its symmetry plane) which fill the canopy with stagnation pressure ($C_p = 1$), whereby it takes its ballooned shape, similar to an open plastic bag that is dragged through the air. This approach has a number of advantages, including a longer durability of the kite, less mass and easier stowability. On the other

hand, it also complicates the implementation of control surfaces and tightens the requirements for the automated guidance system.

With kite sizes between 90 to 180 m^2 and an operating wind range of 3 to 25 m s^{-1}, the current model is designed to reach average cycle powers of 80 to 200 kW [30]. Because all components fit into two regular sized containers, the technology can be flexibly deployed, which enables fast response use and the installation in remote locations which are difficult to reach. Since the kite can easily be stowed away in a bag, the system is also not very susceptible to storms.

1.2. Research Base on Bleed Air Spoilers

As briefly introduced before, SkySails plans to leverage on the potential of bleed air spoilers (BAS) to further improve the average cycle power of its system. In some prototype models, these are already successfully used to keep the kite steady during the final landing approach. The goal is to further understand and improve the feature and to ideally use it as a means to decrease the negative impact that the cyclical return phase, which is required to bring the kite back to its starting position after the maximum tether length has been reached, has on the net power output.

Bleed air spoilers are based on the principle that air is vented from the high pressure region which develops on the inside of a ram-air canopy to the parafoil's upper surface. The vortical structures which this inflicts on the region around and in the wake of the outlet can fulfil a similar purpose as lift spoilers on rigid aircraft wings: the flow over the airfoil is impeded, thus drastically altering the lift and drag characteristics of the wing and lowering its glide slope, as illustrated below.

(a) closed   (b) open
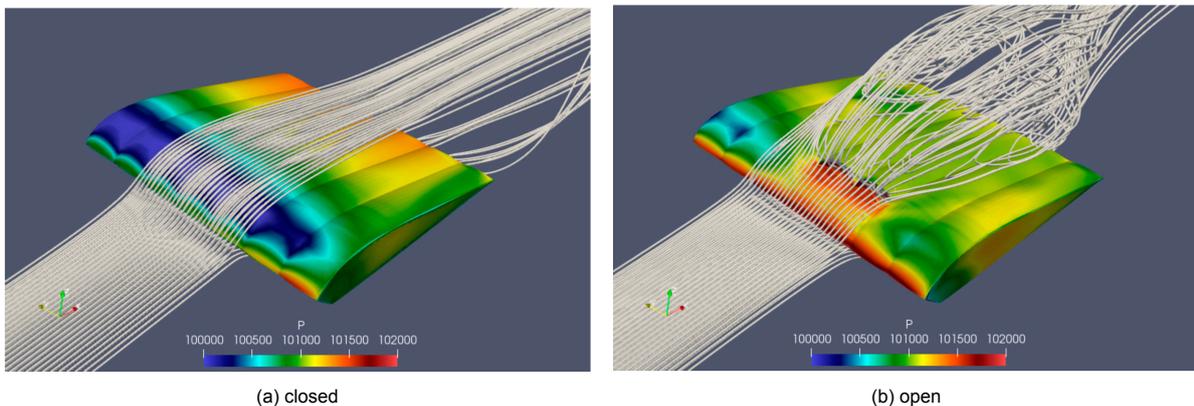
Figure 1.3: CFD simulations of the flow-field around a simplified canopy with and without BAS by Berexa et al. [3]

While, to the extent known, SkySails Power GmbH would be the first to implement bleed air spoilers on its kites for Airborne Wind Energy purposes, their use on parafoils is not new in general. In the mid-1970s, H. Bergeron was the first to implement upper-surface BAS on parafoils for increased control authority of personnel airdrop systems as documented in [4]. Although the idea to employ BAS for an improved lateral control of ram-air gliding parachutes on autonomous systems was already patented by Higgins in 1979 [20], it was not until 2011 that Gavrilovski et al. [16] developed the idea further and systematically investigated their potential by contrasting lower-surface flap spoilers to three upper-surface BAS configurations with identical chordwise (25% at minimum pressure point) and varying spanwise location. They concluded that the latter approach is more expedient and "generates large glide slope changes of 1.5 in a linear manner" with respect to the spoiler deflection. In a following paper, the same group of researchers expanded on their investigations and found that BAS are most effective to control the canopy's glide ratio when the slits are deflected on their leading edge side (air vented in the flight direction) and when they are placed at a chordwise position of around 30% [17].

First computational fluid dynamics (CFD) based investigations on upper-surface BAS were also carried out by Bergeron et al. and are documented in [5]. Simulations were carried out in 2D on rigid airfoil sections with an inlet and an outlet (BAS). Spoilers were placed at $0.2c$ and opened by $0.018c$, where c is the chord length, and 1.7M cells were used to resolve the flow field. The researchers noted that the lift and drag coefficients showed signs of quadratic decay when moving the spoiler further towards the trailing edge, and recorded a mean pressure decrease of roughly 40-60% in the front third

of the canopy. Following experimental trials showed that this decrease is much lower in 3D (roughly 25%), and also highlighted that large spoiler deflections lead to significant fabric deformation in the area around the opening. The authors also noted highly unsteady behaviour when the spoilers are placed at $0.3c$ or further to the rear. In a subsequent, more detailed experimental study by Bergeron et al. of a prototype in flight, the lateral and longitudinal control capacity of BAS located in the front or rear of the canopy were compared. The results showed that an activation of the BAS could also decrease flow reversal at the inlet and thus improve the flow behaviour on the lower surface of the airfoil [4].

In a parallel study by Ghoreyshi et al., researchers established recommendations with respect to grid quality and resolution for two-dimensional cases, noting that reliable time-averaged simulations required grid sizes between 500k-2M cells using wall-orthogonal prism layers in the boundary layer [18]. The most extensive study carried out on BAS up to this point was again by Berexa et al., documented in [3], who performed a number of wind tunnel experiments to analyse the aerodynamic forces and moments induced by spoiler configurations varying in chordwise position and number (further described in Chapter 5).

1.2.1. Conclusions and research gaps

Although the topic is not entirely new, the insights they provide are not easily transferable to the field of airborne wind energy. One aspect which all of the studies described in the above have in common is that they were carried out on the basis of ram-air parachutes for the airdrop of persons and the aerial delivery of equipment. It comes as no surprise that many were also issued by the United States military. The aim of these studies was always to employ BAS to achieve a symmetric or asymmetric reduction of the parafoil's glide slope: a powerful tool to increase the control authority over the canopy and improve the accuracy of automated guidance systems.

Whereas obvious parallels exist between ram-air parachutes and the power kites employed by SkySails, there are a number of important distinctions, too. A fundamental difference, for instance, is the tethered flight of SkySails' kites as opposed to the free flight of parachutes, which has important implications for the trim angle of attack and hence the loads acting on the kite.

Furthermore, the intentions for the use of bleed air spoilers do not exactly coincide. While the chief purpose of BAS on ram-air parachutes is to achieve a local reduction of the glide slope, the intents for AWE purposes are more intricate. In order to improve the average cycle power of the system, the cyclical return phase should both be as short as possible, so that the net share of the power phases in a given time-window is increased, and require as little energy as possible in order to maximise the net energy gain (NEG) of a given cycle. While the first goal indeed aligns with the capacity to lower the kite's glide slope, as will be demonstrated in Section 3.3.2, the latter calls for a means to decrease the resultant loading which acts on the tether when reeling the kite back in to its starting position. A configuration which achieves a lower glide slope, $E = L/D$, through a massive increase in drag could for instance lead to further increased tether forces and thus potentially be counterproductive in lowering the average cycle power.

1.3. Objective and Research Question

The potential use of bleed air spoilers for AWE purposes clearly entails a wide range of questions which have yet to be answered, to which this study strives to contribute. This section outlines the scope of the investigations that were carried out in this context.

Research Objective

In line with SkySails' aim to leverage on BAS as an instrument to further improve the overall performance of their system, the research objective is formulated as follows:

To make recommendations on how bleed air-control can be used to improve the cycle performance of soft-kite-based airborne wind energy systems by investigating and comparing the aerodynamic effects of various bleed air spoiler configurations by means of computational fluid dynamics simulations.

Research Questions

Based on the this objective the underlying research question of the thesis is derived, which the investigations discussed in the following chapters aim to answer:

To what extent can upper-surface bleed air control be used to minimise reel-in power loss for soft-kite-based airborne wind energy systems?

In addition to the main research question, a number of lower level sub-questions are defined which break the overall target up into smaller, more tangible problems. These serve as a practical guideline, and, ideally, in answering the individual sub-questions conclusions with respect to the main research question automatically fall into place.

- How should the aerodynamic performance of the kite ideally change during the return phase in order to increase the net power output of the system?
    - What are the important metrics, and how do they compare in their significance?
    - How much room is there for improvements?
- How are the important aerodynamic phenomena induced by bleed air spoilers best captured in CFD simulations?
    - How can dependable solutions be established?
    - To what degree of complexity should the resultant flow field be resolved by CFD in order to obtain reliable results?
    - What are the limitations of CFD in predicting the effects induced by bleed air spoilers?
- How do bleed air spoilers affect the performance of a ram-air power kite?
    - What impact do various spoiler configuration parameters (e.g. positioning) have on the system?
    - What is the expected impact of the bleed air spoilers on the kite's stability?

1.4. Methodology and Thesis Outline

The remainder of this thesis will be structured as follows. First, some fundamental theoretical background knowledge is introduced in Chapter 2 that is helpful in understanding the methods and conclusions drawn in the subsequent chapters. It is followed by an analysis of the potential impact which an improved reel-in performance could have on the net power output of the system in Chapter 3. In doing so, the key metrics which are later used for the aerodynamic investigations are identified, and a numerical model is introduced to calculate the average mechanical power output of the system as a basis for comparison between the evaluated bleed air spoiler configurations. Chapter 4 covers the setup of the simulation environment with a particular focus on the meshing procedure and implemented physics and solver settings. Chapter 5 is dedicated to the thorough CFD validation which preceded the actual investigations of bleed air spoilers in an AWE context. In this phase, the accuracy as well as the limitations of the developed simulation environment are established through various sensitivity studies with respect to the spatial and temporal discretisation, and not least also through the comparison of polars to an independent experimental study by Berexa et al.. Finally, Chapter 6 deals with the evaluation of various bleed air spoiler configurations on the basis of a common power kite geometry used by SkySails. The thesis closes with a review of the conclusions which were drawn from Chapters 3, 5 and 6 along with some final recommendations for promising designs as well as an outlook on potential future studies.

2

Theoretical Background

In today's world of research and development, it is only rarely the case that engineering problems can be solved with the help of closed analytical formulas alone. Hence, numerical methods have been established for some time in almost all scientific fields as a common means for investigations. By the application of suitable mathematical models, they make highly complex physical processes tangible. Whereas prototypes were initially developed at great expense of time and money through an iterative process that centred on experiments, today's numerical models provide developers with a powerful tool for a cost-effective and risk-free investigation of their designs. In addition, significantly more precise insights into the physical mechanisms behind multilayered processes are often made possible.

In the world of aerodynamics, a cornerstone of such numerical methods is the application of computational fluid dynamics (CFD). This chapter deals with the theoretical foundations of CFD required for a better understanding of the contents of this thesis. In this context, detailed discussions and derivations are intentionally omitted. Rather, emphasis is placed on conveying the most important connections as well as relevant implications for the work. For further information, the reader is referred to more elaborate, subject-specific textbooks and papers through literature references in the corresponding passages.

2.1. Fundamentals of Fluid Mechanics

CFD simulations approximate the laws of fluid mechanics by means of numerical methods. As such, it is instructive to first review the related central theorems, equations, and models before moving on to how they are discretised in time and space.

2.1.1. The Reynolds Transport Theorem

Two distinct approaches exist to track the motion of substances, which differ in the type of reference frame that is used. In the Lagrangian approach, the reference system moves with the observed particle. This is for instance ideal for solid mechanics, where the system under observation is simultaneously our object of interest (e.g. a piston, moving beam, or wing). In fluid mechanics, one is however rarely interested in the exact path taken by a fluid particle and rather wants to record the overall impact of the flow on an object (e.g. aerodynamic loads which move the beam, exert pressure on the piston, or which act on the wing). The latter requires the definition of a control area or volume in the relevant surrounding of our body, which is generally fixed in space. This second type of reference system is central to the Eulerian approach.

The transfer of quantities from the global system to the control volume is summarised in the Reynolds Transport Theorem, also referred to as the integral form of the Lagrangian or material derivative. For this, the flow in a time-varying space, $V(t)$, with the surface $\partial V(t)$ is considered. If Φ represents an arbitrary tensor-, vector-, or scalar-based property of the substance, then the temporal change of the property within the control volume is given by

$$\frac{\mathrm{d}}{\mathrm{d}t} \int_{V(t)} \Phi \, \mathrm{d}V = \int_{V(t)} \frac{\partial \Phi}{\partial \mathbf{t}} \, \mathrm{d}V + \oint_{\partial V(t)} \Phi \mathbf{v_r} \cdot \mathbf{n} \, \mathrm{d}S. \qquad (2.1)$$

The first summand on the right-hand side describes the change in Φ over time in the integration domain, for instance through source terms, while the second takes into account the change due to inflows and outflows of Φ across the integration boundaries. $\mathbf{v_r} = \mathbf{v} - \dot{\mathbf{x}}$ is the relative velocity between the substance under consideration and the surface of the control volume moving at the velocity $\dot{\mathbf{x}}$. The outward unit vector of the surface normal is denoted by $\mathbf{n}$. Together, $\mathbf{v_r} \cdot \mathbf{n}$ yields the flux across a cell interface.

2.1.2. The Navier-Stokes equations

The Navier-Stokes equations form the physical basis of fluid mechanics and describe the flow of a compressible, Newtonian fluid in a Eulerian reference frame. The complete set of equations is formed on the basis of the continuity equation (in the form of Reynolds Transport Theorem applied to the velocity field), the momentum equation (Newton's second law) and the law of conservation of energy (first law of thermodynamics). If no source terms occur, they can be combined to form the Navier-Stokes momentum equation (derived in more detail for instance in [33, p. 242]):

$$\rho \underbrace{\frac{\mathrm{D}\mathbf{v}}{\mathrm{D}t}}_{\text{A}} = \rho \left(\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla) \mathbf{v} \right) = \underbrace{-\nabla p}_{\text{B}} + \underbrace{\mathbf{f}}_{\text{C}} + \underbrace{\mu \nabla^2 \mathbf{v}}_{\text{D}}, \tag{2.2}$$

where the underscored terms relate to the following quantities:

- A: Lagrangian/material derivative
- B: pressure gradient
- C: body forces (e.g. Coriolis or gravitational forces)
- D: viscous and surface (shear) forces

This system of nonlinear partial differential equations forms the mathematical basis employed by computational fluid dynamics (CFD) solvers when predicting flow phenomena in all of its facets, including boundary layers and turbulence. A non-dimensionalisation of equation (2.2) further yields some key classification numbers for Newtonian fluid flows, such as the Reynolds number, which is the ratio of inertial over viscous forces in the flow and is defined as $Re = \frac{vL}{\nu}$, with L denoting the characteristic length of the flow.

2.1.3. Modelling of turbulence

According to one of the most renowned experts for the modelling of turbulence, Stephen B. Pope, turbulent flows are characterised by a "velocity field that varies significantly and irregularly in both position and time" [25, p. 3]. Turbulence always naturally develop when inertial forces in the flow become large enough to overcome viscous and other stabilising forces, causing small perturbations to the flow field to evolve into highly chaotic, unsteady flow structures [10]. Theoretically, turbulent flows could also be captured accurately using the Navier-Stokes equations through so-called direct numerical simulations (DNS). However, this requires enormous computational resources and is thus generally not feasible for engineering applications. Fortunately, several methods exist which model the turbulence without replicating it directly. The most common method is the application of the Reynolds Averaged Navier Stokes (RANS) equations, which is for instance derived in detail in [25, p. 83]. In a process commonly referred to as Reynolds decomposition, the Navier-Stokes equations are modified by splitting a given quantity Φ (such as the velocity field) into its time-averaged mean, $\overline{\Phi}$, and a fluctuating component, Φ', as shown below:

$$\Phi = \overline{\Phi} + \Phi'. \tag{2.3}$$

If we apply Reynolds decomposition to the flow quantities in the Navier-Stokes equations as done in equation (2.4), the fluctuations, Φ', yield additional terms, which are also referred to as Reynolds stresses and are summarised in the Reynold's stress tensor: $\tau_{ij} = \rho \overline{v_i' v_j'}$.

$$\rho \left(\frac{\partial \overline{v_i}}{\partial t} + \overline{v_j} \frac{\partial \overline{v_i}}{\partial x_j} \right) = -\frac{\partial \overline{p}}{\partial x_i} + \nu \frac{\partial^2 \overline{v_i}}{\partial x_j \partial x_j} - \rho \underbrace{\frac{\partial \overline{v_i' v_j'}}{\partial x_j}}_{\text{new term}} \tag{2.4}$$

These stresses all together introduce six new variables, leading to a total of 10 unknowns in equation (2.4). Meanwhile, only four equations are available to solve for them, namely the three Reynolds equations and the Poisson equation or mean continuity equation. This dilemma is generally referred to as the closure problem since the equations thus cannot be solved. When solving for the mean flow characteristics, one hence needs to model this last, newly added term.

In the widely used Boussinesq approach, the turbulent stresses and the transverse derivatives of the averaged velocity gradients are assumed to be proportional to each other. For this purpose, the molecular viscosity (eddy viscosity), ν_t, is introduced as a necessary proportionality factor along with the eddy viscosity approximation for τ_{ij}, shown below:

$$-\overline{v_i' v_j'} = 2\nu_t S_{ij} - \frac{2}{3} k \delta_{ij} \, , \tag{2.5}$$

where S_{ij} is the mean strain rate tensor and k is the mean turbulence kinetic energy, and k and ν_t are unknown. Turbulence modelling attempts to approximate the eddy viscosity, which is generally done using linear eddy viscosity models that are subdivided as follows:

- Algebraic (zero equation) models: linear function for ν_t; no partial differential equations involved

- One-equation models: Prandtl's model or Spalart-Almaras model

- Two-equation models: k-ϵ or k-ω models

The simulations carried out in this work use both the k-ϵ and the k-ω models, combined in the shear stress transport (SST) k-ω model introduced by Menter in [24]. The k-ϵ model provides two partial differential, transport equations for the dissipation rate of turbulence kinetic energy, ϵ, and the turbulence kinetic energy itself, k. While it is well suited for the calculation of flows within the fully turbulent layer away from boundary-layer flows, it is bad at predicting flows close to walls or flows with dissipation. Here, the k-ω model should rather be used, which is in turn quite sensitive to inflow conditions and should thus not be used for free flows. Menter's SST model combines both models through a blending function, thus taking the best from both sides [25, p. 384]. For more insights on the modelling of turbulence and detailed derivations of the two-equation models, the reader is referred to Pope [25].

2.1.4. The Law of the Wall

An important aspect of turbulence modelling is the behaviour of the flow in regions close to the wall. Due to the adhesion (no-slip) conditions that apply here, these are characterised by their particularly high velocity gradients in the wall-normal direction. The SST model can only reproduce these conditions with sufficient accuracy at a very fine mesh resolution (see Section 4.1.2 for a description of such a mesh). Alternatively, so-called wall models can be used, which are based on a combination of empirical and physical relationships and predict the velocity, turbulence kinetic energy and specific dissipation rates at a given distance from the wall. In order to know whether a wall model is required, and if so, which is suitable to make this prediction, the dimensionless wall distance, y^+, is used as a distinguishing characteristic, which describes the stratification of a flow independent of its velocity or the fluid properties. It is composed of the shear velocity, u_τ, the kinematic viscosity, ν, and the dimensional distance normal to the wall, y, and is defined as:

$$y^+ = \frac{u_\tau y}{\nu} \, . \tag{2.6}$$

Using y^+, a boundary layer flow can be divided into the following three regions following the 'law of the wall', which is for instance more closely described in [25, p. 303]:

- $y^+ \leq 5 (\leq 10)$: linear viscous sublayer. The Reynolds shear stresses can be neglected in this region with respect to the viscous shear stresses so that a flow which remains laminar is formed.

- $30 \leq y^+ \leq 500 \ldots 1000$: log-law region. Up to the end of this layer, the universal wall law $u^+ = f(y^+)$ can be used to determine the velocity profile.

- $y^+ \geq 500 \ldots 1000$: Exterior region. Above the near-wall region, flow states can only be determined numerically or experimentally.

In the transition (buffer) region between the laminar sublayer and the near-wall region, a transition function is used to characterise the velocity profile, since both viscous and turbulent shear stresses are relevant here.

State-of-the-art CFD solvers typically have a setting[1] which automatically selects the correct equations to be solved for velocities in the first cell layer depending on its y^+ value. In this sense, the solver automatically disables wall models when the lowest cell layer height falls within $y^+ < 1$, where the velocity is zero due to the no-slip condition and applies the suitable equations when it falls into one of the regions listed in the above.

2.2. Numerical Discretisation for Computational Fluid Dynamics

In order to solve the partial differential equations presented in Section 2.1.3 by numerical methods, they have to become tangible for the mathematical models first. For this purpose, a temporal and spatial discretization has to be carried out.

The first step of spatial discretization consists of the decomposition of the solution area into a numerical computational mesh. Subsequently, the continuous differential equations (2.4) are discretized on this mesh. Besides spectral, boundary element methods, etc., the most important methods are as follows:

- Finite Difference Method (FDM)

- Finite Element Method (FEM)

- Finite Volume Method (FVM)

The finite difference method uses the differential form of the conservation equations. At each mesh node, the spatial and temporal derivatives are approximated by means of difference quotients; for higher derivatives, a Taylor series expansion is usually used.

The finite element method is mainly used for the discretization of structural-mechanical problems, but can also be applied in CFD. It is useful for elliptic and parabolic problems in the non-compressible domain, and not so much for hyperbolic ones. Although care must be taken to ensure a conservative solution, it is said to be more stable than FVM methods. On the downside, it entails larger memory and computational requirements than the finite volume method and does not work with all types of spatial discretisations [28]. Hence, the latter is typically chosen. The FVM is also used by the CFD solver which was used in this thesis and is thus discussed in some more detail below.

2.2.1. The Finite Volume Method

The Finite Volume Method is based on the integral form of the conservation law and the divergence theorem. It works with arbitrary volumetric meshes that have to be carefully tailored to the flow domain in order to accurately capture the flow phenomena one is interested in (see the subsection on computational meshes below for a more detailed explanation). This mesh is used to define control volumes (CVs; see Figure 2.1), which can be done using either a cell-vertex or cell-centred scheme. In the cell-vertex scheme, the flow variables are stored on the vertices of the constructed mesh and CVs are defined by combining all cells that share a common vertex (CVs can hence overlap). In the cell-centred approach (the method used in this thesis), the control volumes are the cells themselves.
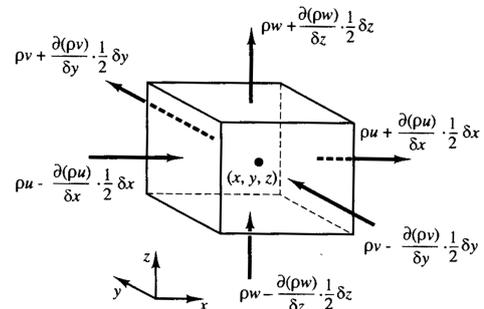
Figure 2.1: Representative illustration of mass fluxes in and out of a control volume (from [32])

---

[1] In Star-CCM+, the commercial CFD suite which was used for this thesis, this is done when the "all y^+" wall-treatment is applied.

To determine the flux, $\Phi v_r \cdot n$, at the interface between
two control volumes, polynomial interpolation is applied on a corresponding stencil of neighbouring nodes in the domain. The size of the stencil used for this interpolation determines the maximum order of the spatial discretisation scheme. By balancing the fluxes calculated on each cell face, the average flow variables within the CV are found, which are then stored for the next iteration. Over the course of this procedure, a consistent definition of the fluxes guarantees global conservation of the transported variables to the order of machine precision. In this sense, the outflow from one cell is always equal to the inflow of the next, and thus mass or energy is not created or lost in the process.

For more detailed descriptions of the Finite Volume Method for CFD applications, the reader is referred to standard literature such as [32] by Versteeg and Malalasekera and [13] by Ferziger and Peric.

2.2.2. Spatial discretisation

The convergence behaviour, as well as the qualitative result of a CFD simulation, depends significantly on the quality of the underlying computational grid. These are commonly divided into structured and unstructured grids/meshes, which can be further combined into a third group of hybrid meshes. Figure 2.2 includes illustrative examples for each of these mesh types.

(a) (block) structured   (b) unstructured   (c) hybrid
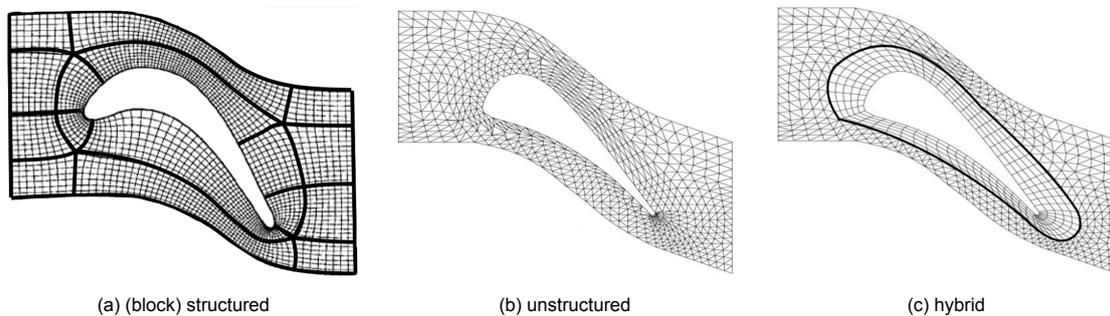
Figure 2.2: Classification of CFD meshes (from [22])

In structured grids, cells take the form of topological quadrilaterals (2D) or hexahedra (3D). These are arranged in an ordered fashion, which allows that neighbouring CVs in the physical space are also neighbouring elements in the matrix storing the flow variables. For two-dimensional grids, an element is therefore uniquely determined by the indices (i,j), for three-dimensional ones by the indices (i,j,k). When meshing complex structures, it is also possible to sub-divide the flow-domain into smaller blocks which are suitable for structured meshing (as was done in Figure 2.2a) even though the blocks are arranged in an unstructured manner amongst each other. Because of the rather simple grid logic, fast implementation and processing on a computer and the possibility to use directional solution methods is achieved, which contributes to a reduction of the overall computational effort. Furthermore, since the cell surfaces can be positioned perpendicularly to large gradients, such as in boundary layer flow, unwanted numerical diffusion effects are greatly reduced, which often allows for the use of larger cells [32]. On the downside, structured meshes tend to require a large amount of user input through time-consuming tailoring of the connectors (edges of structured blocks) and even individual grid points until a satisfactory mesh quality is achieved.

Meanwhile, unstructured grids are characterised by their high flexibility and good capacity for automation. Although they typically consist of triangular (2D) or tetrahedral (3D) elements, principally any polyhedral element can serve as a cell. Star-CCM+, for instance, combines cells from an underlying tetrahedral mesh to polyhedral cells, thus immensely cutting down on the total cell count. With unstructured grids, continuous indexing of the mesh nodes is not possible. As a result, elaborate bookkeeping is required by the CFD solver using a so-called 'connectivity matrix' in order to know which cells share a certain face. This elaborate cell management increases the complexity of the numerical solution algorithms. Other downsides include a need for more resolution in regions with large flow gradients than their structured counterparts because cell faces cannot be tailored to the expected flow as easily.

Hybrid meshes are a common means of combining the strengths of both structured and unstructured grids. The global mesh may contain regions that in themselves have some structured characteristics.

Boundary-layer flows are a typical example where this can be a big benefit. Here, an unstructured surface mesh is then generated which is evenly extruded to form a so-called prism-layer mesh, as was for instance done in Figure 2.2c around the airfoil. The cell faces in the wall-normal direction are thus aligned in a perpendicular manner to the strong gradients inflicted by the no-slip condition at the wall. In the outer (core) region, the grid is again fully unstructured.

Since a CFD solution is very sensitive to spatial discretisation, special attention is required when constructing the mesh. In order to avoid subjective decisions, there are a number of guidelines and indicators that can be used to evaluate the quality of a grid. These should be kept in mind right from the start in order to avoid the need for laborious rectification work. An official guideline [1] by the European Research Community on Flow, Turbulence and Combustion (ERCOFTAC) and [18] detail a number of key directives, the most important of which are summarised in the following:

- Cell Size (for surface meshes on aerodynamic profiles)

    - 1‰ of the local chord in the chordwise direction at the leading and trailing edges
    - 1‰ of the semi span at the root (when applicable) and tip
    - The first cell layer in a boundary layer mesh should fall within a distance of $y+ < 1$ (see equation (2.6)) to the wall.

- Growth rates

    - 1.15 to 1.25 for boundary layers in the wall-normal direction
    - Generally, growth rates larger than 2 should be avoided in all directions across the entire mesh.

- Cell skewness and aspect ratio

    - The skewness of a given cell, defined as $(V_{i,\mathrm{opt}} - V_i)/V_{i,\mathrm{opt}}$, where $V_{i,\mathrm{opt}}$ is the size of an equilateral cell with the same circumradius as our cell of interest, should equal at most 0.875 (0 is a perfect cell).
    - The aspect ratio of a given cell should be as close as possible to one. The boundary layer mesh is an exception to the rule. Due to the low cell layer thicknesses at the wall, much larger values are desired there. Generally, the cell aspect ratio should only change slowly.

2.2.3. Temporal discretisation

In equation (2.4), temporal derivatives appear which also have to be discretized appropriately. In contrast to the spatial coordinates, this however only affects one term. One generally distinguishes between implicit and explicit methods.

Explicit methods are characterised by the fact that they only draw on known values from previous iterations to infer the solution at the next time step, t_{n+1}. For these schemes, a maximum permissible time increment exists for stability reasons, which is dependent on the mesh size of the grid. An important metric to evaluate whether the chosen time-step is adequate is the Courant number, which is defined as follows:

$$Co = \frac{\Delta t v}{\Delta x} \, ,\qquad(2.7)$$

where v and Δx are the velocity and grid width of a cell in a certain direction. It provides information on whether a fluid particle moves by more than one cell width during a time step, which would lead to an instability of the calculation. This is commonly referred to as the Courant-Friedrichs-Lewy condition, which is satisfied when $Co < 1$ [13]. As will become clearer in Section 5.3, an explicit scheme would for instance not be feasible for the simulations treated in this work, because time-steps at the order of 10^{-5} s would be required.

Meanwhile, the Courant number only has a limited impact on the accuracy of an implicit time-stepping scheme. Some, like the first-order accurate implicit Euler scheme shown below, are even unconditionally stable [13].

$$\left.\frac{\partial \Phi}{\partial t}\right|_{n+1} = \frac{\Phi_{n+1} - \Phi_n}{\Delta t} \tag{2.8}$$

As the equation demonstrates implicit schemes work with the solution of the time-step they are trying to predict from the start. Because of the existence of linear convective terms in equation (2.4), this however cannot happen in a single computation and thus iterations are required. In implicit methods, the time-step is not limited by stability reasons but rather by the types of flow phenomena which one aims to capture. This is because implicit time-stepping schemes inherently perform time averages over a given time step. As a consequence, turbulence that occurs at time scales smaller than Δt is not reproduced. Higher-order implicit schemes can also be used. A detailed discussion of the impact which the discretisation order can have on the solution is later included in Section 5.3. For more elaborate descriptions of temporal discretisation schemes and detailed derivations, the reader is referred to Ferziger and Peric [13].

3

Impact Analysis of an Improved Reel-In Performance on the Power Cycle

Before the simulation setup and experimental results are discussed, it is important to establish a proper understanding of how a ground generator based kite power system is generally operated, and in particular how the transfer and return phases should ideally be optimised in order to maximise the overall performance of the system. In doing so, the following sections will also highlight the relevance of the key performance metrics for the bleed air spoilers, namely the glide ratio E_{in} and the resultant aerodynamic loading coefficient $C_{R,in}$, which will be referred to on a regular basis in the subsequent chapters.

In the following, first the resultant tether loading and with it the mechanical power generated/consumed by the ground station for given flight conditions is derived. Then, the characteristic pumping cycle with which kite power systems such as that used by SkySails harvest energy from the wind is explained. After introducing a numerical model to predict the system performance during each of the three cyclical phases, they are jointly evaluated for a range of combinations of E_{in} and $C_{R,in}$ to reveal the potential impact which a given means to change the aerodynamic characteristics of the kite during the return phase may have.

3.1. Basic Analysis of the Resultant Tether Loading

For a better understanding of the loads which the canopy creates in flight, it is instructive to review the definition of the apparent wind, v_a, as measured by an onboard anemometer installed on the kite, and how it relates to the resultant mechanical power out-/input at the ground station. As explained in [12], v_a is given by

$$v_a = v_\infty \cos \vartheta E - \dot{l} E \, , \tag{3.1}$$

where v_∞ is the ambient wind, ϑ is the spatial angle that gives the position of the kite in the wind window (angle between the wind direction and the connecting line between the ground station and the canopy; see Figure 3.1), $\dot{l}$ is the reel in/out speed, and $E = {}^L/_D$ is the kite's glide ratio. Equation (3.1) demonstrates that v_a is low for large values of ϑ, and increases proportionally with E, which propels the kite forward.
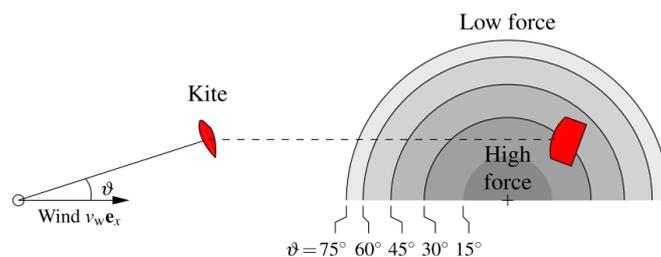
Figure 3.1: Depiction of the wind window, highlighting the importance of ϑ for the aerodynamic forces felt by the kite (from [12])

Neglecting gravitational and inertial forces acting on the kite, there must be a balance between the aerodynamic lift and drag forces acting on the canopy and the reactive tether force for quasi-steady conditions to apply. As a result, the resultant aerodynamic loading coefficient, $C_R = \sqrt{C_L^2 + C_D^2}$, directly relates to the tether loading. It is worthwhile to point out that because there must always be a moment equilibrium around the control pod (otherwise the canopy would pitch up or downwards) the kite always assumes the same trim angle of attack, whereby C_R and E remain constant. Assuming no drag and gravity acting on the tether, this loading also gives the pulling force, F_R, on the drum located in the ground station which drives the generator:

$$F_R = \frac{1}{2}\rho v_a^2 S C_R \, , \tag{3.2}$$

where S is the projected surface area of the canopy and ρ is the air density.

Since the power of a system is calculated by multiplying the acting force by the velocity – in this case the reeling speed of the system – the mechanical power, P_{mech}, is found by

$$P_{\text{mech}} = F_R \dot{l} \, , \tag{3.3}$$

which is positive or negative depending on whether the tether is unwinded or retracted.

3.2. The Pumping Cycle

The cyclical process by which ground generator based kite power systems operate is often referred to as a 'pumping cycle' to reflect the periodic generation of mechanical power as the tether is reeled out, and the expense of energy during the transfer and return phase, which is required to bring the kite back to its starting position once the maximum tether length has been reached. Figure 3.2a depicts a typical power cycle and its individual phases. Figure 3.2b further demonstrates the cyclical gain and expense of power over time for a representative system. When the Net Energy Gain (NEG) per cycle is positive, its operation becomes useful.

(a) typical pumping cycle showing characteristic figures of eight in the power phase (from [12])

(b) Representative time history of (from [11])
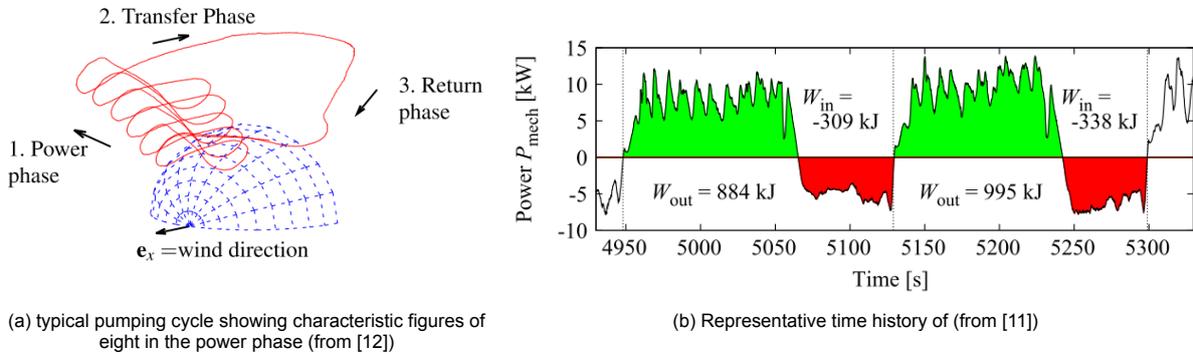
Figure 3.2: Overview of the characteristic kite power pumping cycle

During the power phase, which commences the cycle, the kite is steered above the high-force centre of the 'wind window' shown in Figure 3.1 where it is flown in eight-down patterns, which means that the kite accelerates towards the ground during the straight, rather steady-state segments before it is turned around again. Although winds are generally stronger at larger heights, it is often preferred to fix the altitude to a certain limit in order to dip deeper into the high force zone. Following equations (3.1) and (3.2), the kite thereby reaches very large values of v_a and traction forces are maximised. By periodically flying the kite through the high force zone, according to [11] the closest possible match is achieved to Loyd's limit, established in his pioneering paper [23] from 1980, for a kite operating in perfect crosswind conditions:

$$P_{\text{Loyd}} = \frac{\rho C_R S}{2} \frac{4E^2}{27} v_\infty^3 \, . \tag{3.4}$$

Once the system approaches its maximum tether length, the kite is steered into a transfer trajectory which brings it into the outer zone of the wind window (larger ϑ). Here, $v_\infty \cos(\vartheta_{in})$ is lower and thus the kite may be reeled in again towards the ground station under smaller loads. It is important to point out that approximately steady-state conditions apply during the return phase. As such, the kite is reeled in at a constant $\dot{l}_{in}$ and ϑ_{in}, where the latter solely depends on the glide ratio in the reel-in phase and the reel-in speed (further explained in Section 3.3.2). Once the initial tether length is reached, the kite almost directly enters the power phase again and the cycle is started over.

3.3. Evaluation of Individual Cycle Phases

In this section, models are derived and applied that allow for an individual assessment of the three power cycle phases, and how a means to control the canopy's glide ratio and the resultant aerodynamic loading during the return phase would affect the system's overall performance. To this end, some general assumptions are made to keep the complexity and clarity of the calculations at reasonable levels. In particular, it is assumed that:

- aerodynamic forces are large compared to the gravitational and inertial forces acting on the kite,

- the incoming wind is steady and uniform (no shear),

- the kite experiences no side-slip.

- the aerodynamic characteristics of the flying system can be reduced to the kite's glide ratio and the aerodynamic loading coefficient, C_R. Aerodynamic forces acting on the tether are neglected.

- the kite is always in aerodynamic equilibrium.

- the kite is assumed to be rigid (constant aerodynamic characteristics during each phase).

- the flight trajectory may be simplified as depicted in Figure 3.3. For further explanations and how some of the simplifications are compensated, the reader is referred to the subsection of the respective phase.

For the sample calculations, a representative kite with a projected area $S = 0.78 \cdot 120 \, \text{m}^2$, a baseline glide ratio of $E_{base} = 5$ and corresponding $C_{R,base} = 1$ is assumed. Meanwhile, environmental conditions are set to $\rho = 1.225 \, \text{kg m}^{-3}$ and $V_\infty = 10 \, \text{m s}^{-1}$. With sufficient tether length, the canopy is held at a maximum altitude of 200 m. The minimum tether length (start and end lengths of a cycle) is also set to 200 m.
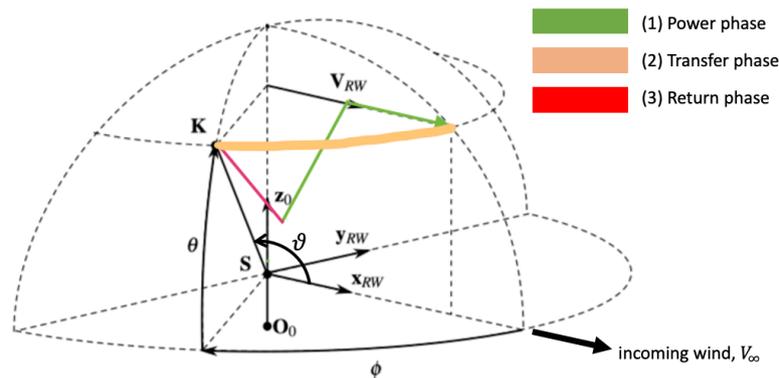
Figure 3.3: Overview of the simplified power cycle as modelled for the analysis, and definition of important angles (adapted from [6])

3.3.1. Power phase

As shown in Figure 3.3, the characteristic power phase consisting of figures of eight is reduced to a flight path made up of two linear segments, both located on the xz-plane. The azimuth ϕ is assumed to remain small during the power phase, and thus the evaluation of ϑ is reduced to that of the elevation

angle, θ (i.e. $\cos\vartheta = \cos\theta\cos\phi \approx \cos\theta$). In the first segment, a constant $\vartheta = 45°$ is maintained. Once the imposed maximum altitude of $h_{max} = 200\,\text{m}$ is reached, the height is held constant and thus ϑ decreases. While the omission of the eight-pattern may seem like a major limitation of the model, it is important to note that the control strategy with which the tether is reeled out strives to achieve a sinusoidal variation of v_a in line with equation (3.1). In the model, it is thus sufficient to directly impose this variation as opposed to modelling the true trajectory. After consultation with the task group responsible for the development of the autopilot at SkySails, approximately ten periods with an amplitude of $2\,\text{m\,s}^{-1}$ around a mean of $25\,\text{m\,s}^{-1}$ are assumed to take place during the power phase at the chosen ambient wind conditions.

The flowchart presented in Figure 3.4 summarises the derived procedure to model the power phase. Its initially unknown duration is discretised into N evenly spaced time-steps, which are denoted by the subscript t. It is assumed that the kite enters the cycle at the minimum velocity ($v_{a,1} = 23\,\text{m\,s}^{-1}$) and leaves it at a peak velocity ($v_{a,N} = 27\,\text{m\,s}^{-1}$). Because the variations of $\dot{l}$ and ϑ are interdependent, the analysis requires an iterative approach based on an initial estimate of ϑ, which is first set to the final elevation angle. With every new iteration, the estimates of ϑ_t are refined until a converged prediction of the phase's NEG is found.
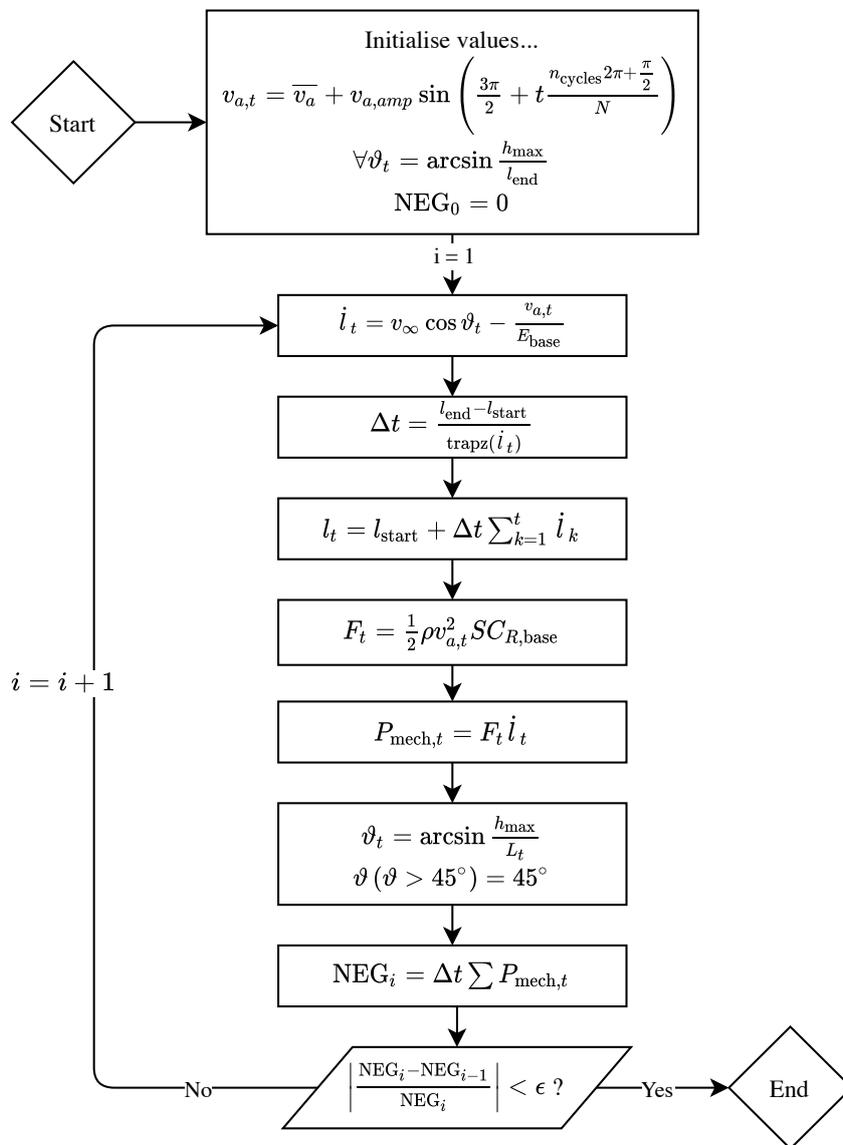
Figure 3.4: Flowchart showing the iterative procedure to model the power phase

Because the solution is independent of the aerodynamic return-phase characteristics, E_{in} and $C_{R,in}$, the calculations for the power phase only need to be carried out once for a given baseline configuration and ambient wind conditions. The results are depicted below.
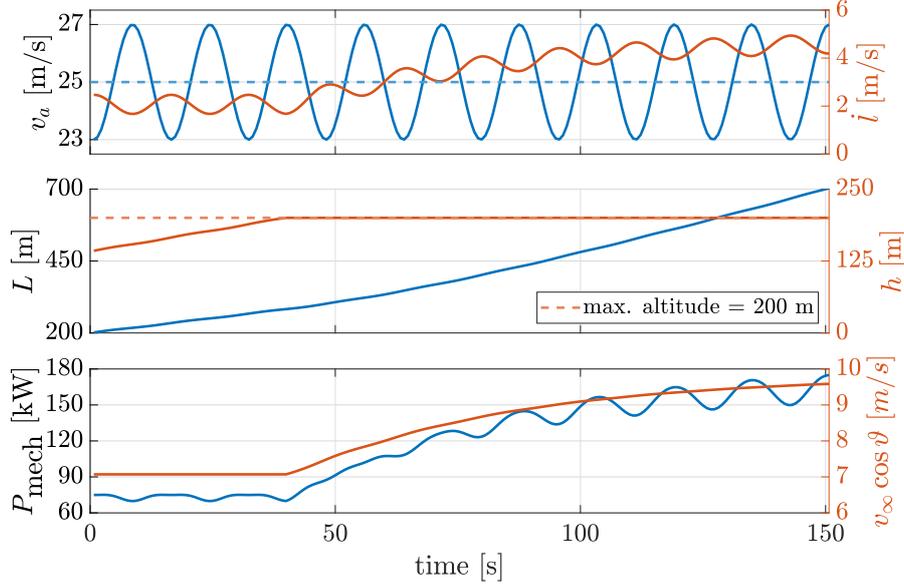
Figure 3.5: Time-history of modelled power phase ($V_\infty = 10\,\text{m/s}$)

The time-history of the modelled power phase clearly shows the imposed sinusoidal variation of v_a, as well as the expected variation of the altitude. Note that as h_{max} is reached, the kite begins to dip deeper into the wind window and thus $v_\infty \cos \vartheta$ increases. As a result, the variation in $\dot{l}$ increases proportionally (see equation (3.2), where the variations in v_a remain constant) and with it P_{mech} following equations (3.2) and (3.3). A summary of the processed results is given in Table 3.1. The predicted average power output amounts to 55.9% of Loyd's limit (see equation (3.4)).

Table 3.1: Summary of system performance during power phase ($V_\infty = 10\,\text{m/s}$)

duration	151 s
avg. P_{mech}	118.7 kW
net energy gain	17.89 MJ

3.3.2. Transfer phase

The transfer phase seamlessly connects to the power phase. As such, it begins at $h = 200\,\text{m}$, an apparent velocity of $v_a = 27\,\text{m s}^{-1}$, and a spatial angle of $\vartheta = \vartheta_{1,end}$ ($\phi = 0°$).

The basic control strategy sets the target $\vartheta_{2,end}$ equal to ϑ_{in}, and imposes a linear interpolation of v_a with respect to ϑ between the start and end points, which can be found analytically. As stated in Section 3.2, the straight reel-in trajectory is inherently dependent on the constant glide ratio, E_{in}, and reel-in speed, $\dot{l}_{in}$, during the return phase. $\dot{l}_{in}$ is a freely selectable design variable, which is defined as a ratio of the ambient wind: $\mu = \dot{l}_{in}/v_\infty$. In order to solve for the static spatial angle at which the kite balances out during the return phase, the respective equation of motion, derived for instance in [11], must be used:

$$\dot{\vartheta} = \frac{v_\infty}{l} \left(E \cos \vartheta \cos \psi - \sin \vartheta \right) - \frac{\dot{l}}{l} E \cos \psi \, , \tag{3.5}$$

where ψ is the kite's angle to the wind that equals 0° following the no-side-slip assumption. By setting the derivative equal to zero, the following implicit equation is found:

$$0 = v_\infty \left(E \cos \vartheta - \sin \vartheta \right) - \dot{l} E \, , \tag{3.6}$$

of which the root for a given $\dot{l}$ can be computed in a predefined realistic interval of 0° to 130°, yielding the static value for ϑ_{in}. Figure 3.6 shows the impact of a lower E_{in} on the static spatial angle which the kite takes during the return phase. After consultation with the autopilot team at SkySails, a maximum value for ϑ of 110° has been imposed in order to prevent excessive overshooting, or in other words to keep the kite from moving too far behind the ground station.
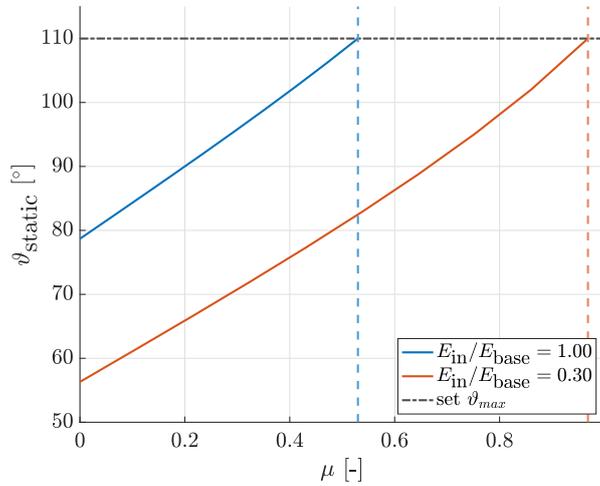
Figure 3.6: $\vartheta_{\mathrm{static}}$ (i.e. ϑ_{in}) over μ for two representative reel-in glide ratios

The plot highlights a key advantage that a lower glide ratio entails, which will also be important for the cycle-evaluation: by lowering E, much larger reel-in speeds are feasible before the kite overshoots, since the higher ratio of drag to lift pushes the kite further into the wind window.

Since ϑ_{in} is now fixed for a given $\dot{l}_{in}$, so is $v_{a,in}$ following equation (3.1). The remaining state variables are derived as follows, where the trajectory is discretised in evenly spaced steps of $\Delta\vartheta$. Note that the return phase should be covered as quickly as possible, and thus a high glide ratio is of advantage. Therefore, the baseline aerodynamic characteristics apply for all states that are unrelated to the return phase.

Knowing v_a for a given spatial angle step i, the respective reeling speed is found by rearranging equation (3.1):

$$\dot{l}_i = v_\infty \cos\vartheta_i - \frac{v_{a,i}}{E_{\mathrm{base}}} \; . \tag{3.7}$$

To relate a given angular step to the respective time, $(\Delta t)_i$, which it requires, the metric distance covered in it, $(\Delta x)_i$, as well as the kite's path velocity, $v_{k,i}$, must be calculated. By definition, v_k is found by subtracting the true wind component from the apparent velocity measured at the kite:

$$v_{k,i} = v_{a,i} - v_\infty \cos\vartheta_i \; . \tag{3.8}$$

To find the corresponding $(\Delta x)_i$, negligible variations in θ and the tether length, l, are assumed[1], so that the kite may be seen to move along the rim of a spherical cap as shown in Figure 3.3. The horizontal distance between this rim and the z-axis if found through $a = l_{end} \cos\vartheta_{1,end}$. With a fixed $\theta = \vartheta_{1,end}$, the azimuthal angle can be calculated for each i through equation (3.9), which in turn yields $(\Delta x)_i$ when applied to equation (3.10).

$$\phi_i = \arccos\frac{\cos\vartheta_i}{\cos\theta} \qquad\qquad (\Delta x)_i = 2\pi a\frac{\phi_i - \phi_{i-1}}{360°} \tag{3.9, 3.10}$$

With the path velocity of the kite for each angular step, $v_{k,i}$, as well as $(\Delta x)_i$, one can finally compute the time-steps $(\Delta t)_i = (\Delta x)_i / v_{k,i}$. With the means to now map v_a, $\dot{l}$ and ϑ over time, the mechanical power and tether length along the trajectory can be computed. Figure 3.7 shows the transfer phase time-histories for sample combinations of $\mu = \{0.6, 0.2\}$ and $E_{in}/E_{base} = \{1, 0.3\}$.

---

[1] a constant θ is deemed to be a fair assumption at large tether length and low altitude
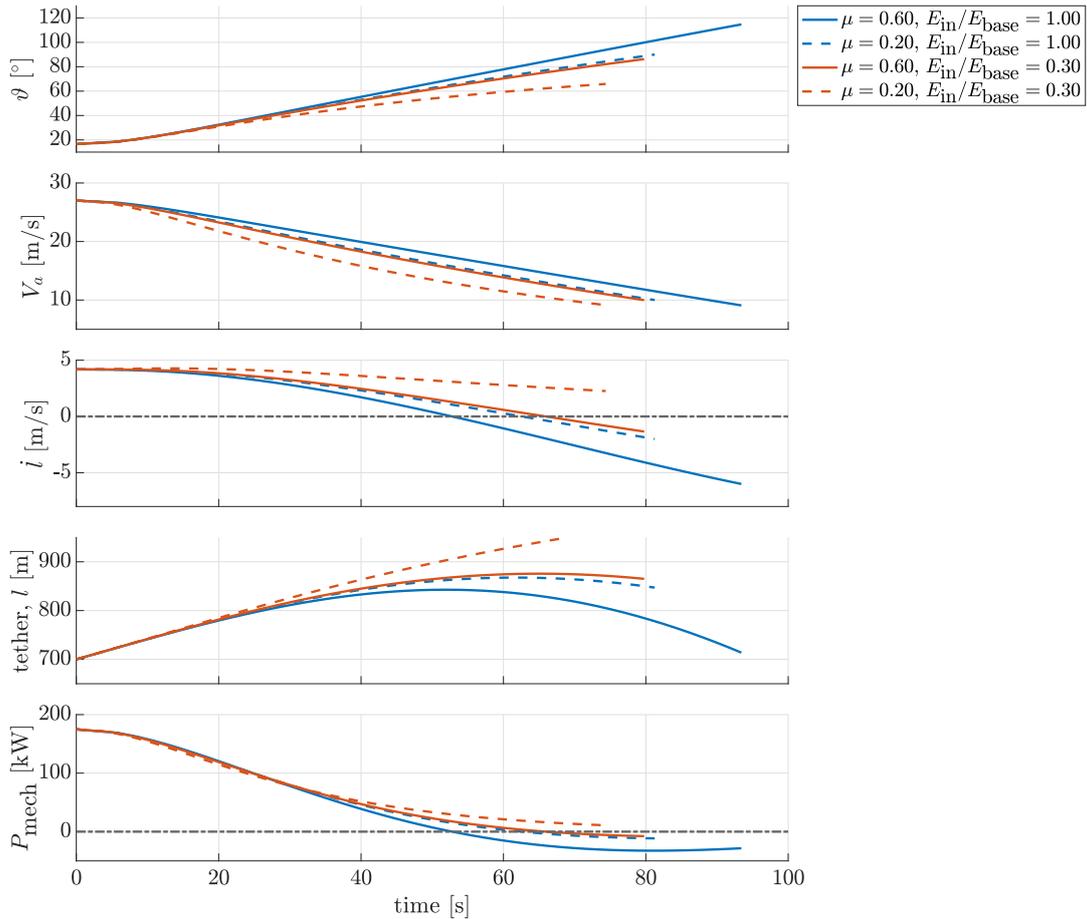
Figure 3.7: Time-histories of modelled transition phase ($V_\infty = 10\,\text{m/s}$)

The time-histories allow for a number of important observations. Firstly, the data again demonstrate that with a lower E_{in} the kite will remain deeper in the wind window. While this means that $v_\infty \cos\vartheta$ will be larger, v_a still remains much lower over the trajectory than the baseline configuration with a smaller E_{in} for both values of μ. This is because of the proportional influence of E_{in} on the return-phase $v_{a,in}$, which is also reduced by 70%. Moving one plot further, a look at the reeling speed $\dot{l}$ shows that the system achieves these low values of v_a by continuing to unwind the tether at larger rates. This comes at the cost that the tether is also longer at the end of the transfer phase for a small E_{in}, which means that more tether will need to be retracted in the return phase.

A comparison of the respective curves for different values of μ shows that given a fixed E, lower reel-in speeds during the return phase mean that the final, static ϑ will be smaller. Meanwhile, the results show that the transfer phase is also shorter for smaller μ since the kite does not have to travel as far. A review of the reeling speeds shows that $\dot{l}$ is larger over the course of the transfer phase for small values of μ. Despite lower apparent winds, the larger $\dot{l}$ causes more energy to be extracted from the wind during this phase, which again comes at the cost of a larger final tether length.

Finally, it is worthwhile to point out that the combination of $\mu = 0.6$ and $E_{in}/E_{base} = 0.3$ clearly yields very similar time-histories to $\mu = 0.2$ and $E_{in}/E_{base} = 0.3$. This shows that similar performance can be achieved in the transfer phase even though the reel-in speed is tripled.

The transfer phases modelled in the above were post-processed to obtain some general performance metrics, which are summarised in Table 3.2.

Table 3.2: Summary of system performance during transition phase for several E_{in} and μ ($V_\infty = 10$ m/s)

E_{in}/E_{base} [-]	μ [-]	duration [s]	L_{end} [m]	avg. mech. P [kW]	NEG [MW]
1.0	0.6	93.4	713.8	39.4	3.68
	0.2	81.2	847.1	59.6	4.84
0.3	0.6	79.8	865.0	62.4	4.98
	0.2	74.4	962.3	72.9	5.42

The data further enforce a corollary that was already suggested based on the plots: although v_a was continuously lower for the smaller glide ratio, the average power output is increased by 58% and 22% for values of μ of 0.6 and 0.2 respectively. As stated before, this is caused by the larger reeling speeds. As a result, the transfer phase does not necessarily have to be seen as 'wasted time' in the power cycle as much, but maybe rather as a logical extension of the power phase that still achieves 34% (and potentially more) of Loyd's power limit (see equation (3.4)) prior to switching straight to the return mode.

3.3.3. Reel-in phase

In comparison to the power and transfer phase, an assessment of the return phase is very straightforward because of the steady-state conditions which apply. The values of $v_{a,in}$, ϑ_{in}, μ and the tether length at the start of the final phase, $l_{in,start}$, are all known from the previous analysis of the transfer phase. It is assumed that the mechanism which allows control of E_{in} also decreases C_R to some extent. Given $C_{R,in}$, the constant tether force, $F_{R,in}$, is easily found with equation (3.2). Meanwhile, the Energy consumed throughout the third phase results from a multiplication of $F_{R,in}$ by the total tether length $\Delta l_{in} = l_{in,start} - l_{in,end}$ that must be covered, where $l_{in,end}$ should again equal 200 m. While the duration follows from the reeling speed, $\Delta t_{in} = \Delta l_{in}/(\mu v_\infty)$, the power consumption is simply equal to the NEG divided by Δt_{in}.

3.4. Cycle Evaluation

With the three models derived in Section 3.3, all necessary means are obtained to evaluate the system performance over the complete cycle when running the computations for a given configuration (i.e. given $C_{R,in}$ and E_{in}) back to back. Figure 3.8 depicts the complete time-histories for cycles with an arbitrary kite design that allows for $E_{in}/E_{base} = 0.3$ and $C_{R,in}/C_{R,base} = 0.9$ for various values of μ. Since the power phase is independent of the return phase, it is the same for all μ.
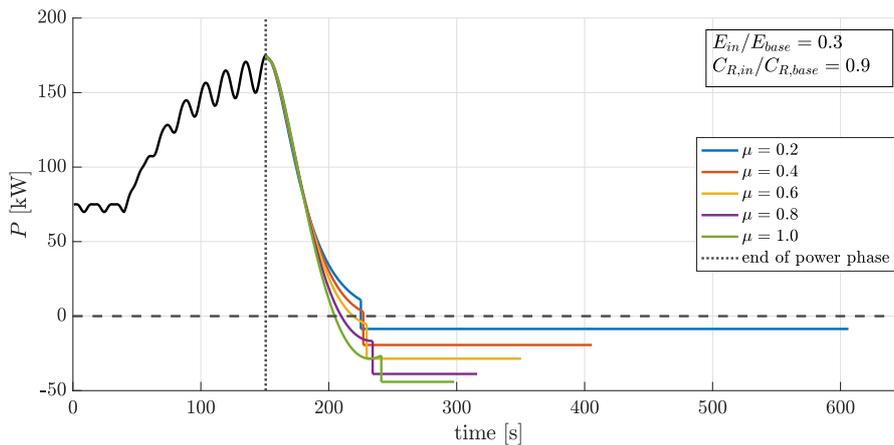
Figure 3.8: Power cycle time histories for different reel-in speeds ($V_\infty = 10$ m/s)

The results show that the smaller μ is, the shorter the transfer phase becomes and the larger the average mechanical power output is throughout it. Meanwhile, whereas the change in the return phase's duration is approximately inversely proportional to μ for low reel-in speeds (e.g. between μ of 0.2 and

0.4), a diminishing return is observed as μ increases further. This is caused by the additional tether length that has to be retracted by the end of the transfer phase when large reel-in speeds are used, as was previously discussed in Section 3.3.2. Note that for very large values of μ, the mechanical power input starts to decrease again towards the end of the transfer phase, and thus a trough is visible for $\mu = 1$. This reflects the fact that as the kite overshoots (i.e. $\vartheta > 90°$), $v_\infty \cos \vartheta$ becomes negative and thus the wind assists the electric motor during the reel-in to some extent.

Figure 3.9a shows the average mechanical power output of the system for kite configurations with various combinations of E_{in} and $C_{R,in}$. Figure 3.9b further depicts the predicted maximum achievable mechanical power outputs for configurations with varying return-phase E and C_R.

(a) $P_{avg,mech}$ depending on the chosen reel-in speed

(b) max. achievable P_{avg} yields by return phase characteristics
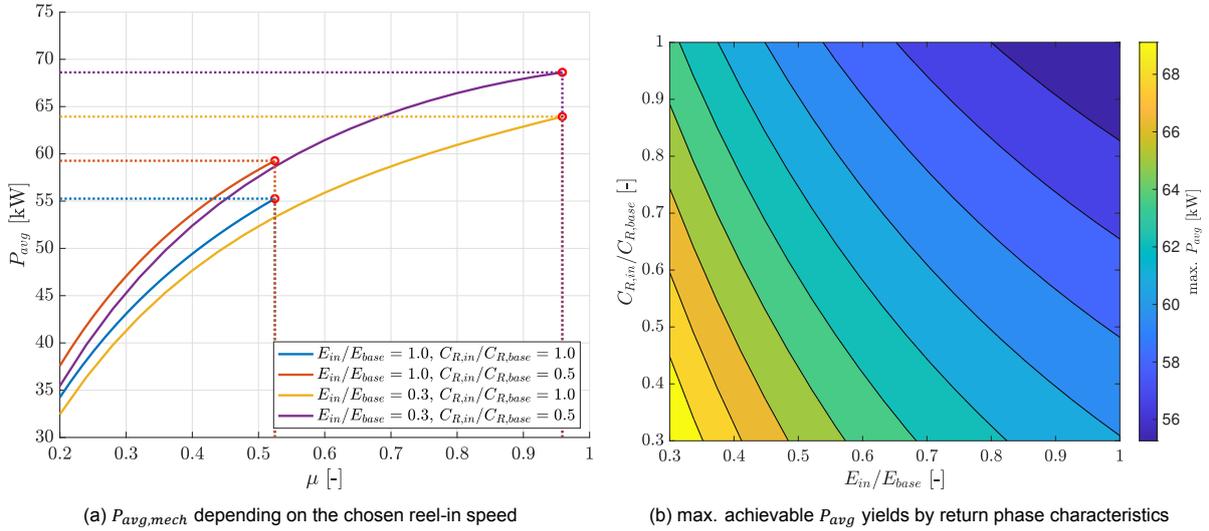
Figure 3.9: P_{avg} yields for various combinations of E_{in} and $C_{R,in}$ ($V_\infty = 10\,\text{m/s}$)

Figure 3.9a again demonstrates the advantage of low E_{in} that they enable larger reel-in speeds. Clearly, this alone leads to much larger average power outputs as the steadily increasing curves are cut off later due to the imposed limit[2] of $\vartheta_{in,max} = 110°$. Meanwhile, P_{avg} also increases when the resultant tether loading is decreased by lowering $C_{R,in}$, which is expected as less energy is then wasted in the return phase while the transfer phase (and thus the boundary conditions of the return phase) remains unaffected by this change.

The same relations between the maximum average power outputs and the kite's configuration can be inferred from the contour plot presented in Figure 3.9b. It also shows that a given percentage decrease in E is generally preferable over an equal decrease in the tether loading during the return phase through a drop in C_R. The most ideal of the investigated representative configurations yields an average mechanical power output that is 28% larger than the baseline case at the given ambient wind.

---

[2]It is worth to acknowledge that the system is in reality also limited by the maximum towing force which it can exert on the tether when reeling the kite back in to its starting position. However, this is regarded as a secondary problem, as it would at most require a redesign of the generator/electric motor in the ground station. For the purpose of the cycle evaluations, it is assumed that the motor is not limited in force and thus no additional constraint is imposed.

3.5. Summary of important findings

In conclusion, although no actual investigations of the aerodynamic effects induced by bleed air spoilers have been carried out yet, this chapter already shed light on the intertwined relationships between the aerodynamic characteristics of the kite and some of the freely selectable variables that make up a power cycle. Some of the most important take-aways for the following chapters are summarised below:

- The system's overall performance is strongly linked to the glide-ratio, E_{in}, and the resultant aerodynamic loading coefficient, $C_{R,in}$, of the kite during the return phase

- For a given kite configuration, an ideal μ can be identified that is dependent on E_{in}. The lower E_{in}, the faster the kite may be retracted without causing excessive overshoot.

- In comparison to each other, a given percentage change in E has a more important effect on the overall performance of the system than an equal change in C_R.

- A low E_{in} can significantly improve the system's performance already in the transfer phase (even though this uses E_{base}).

Having established a dependable metric, $P_{\mathrm{avg,mech}}$, for the evaluation of kite configurations that allow for some degree of control over the glide ratio and resultant tether loading during the return phase, the way is clear to move ahead with some aerodynamic investigations of the actual effects which bleed air spoilers have on a ram-air kite's performance.

4

Pre-Processing and Computational Setup

In this section, the general workflows are described which were used to produce the investigated geometries and to generate the corresponding hybrid meshes. Also, the choices for physics and solver settings are discussed which form the other half of the numerical basis for the treated simulations.

Aside from the geometry cleanup and conversion of the provided, tessellated structure files to smooth `.step` files, all of the pre-processing as well as the actual CFD simulations were carried out using Simcenter STAR-CCM+: a commercial CFD suite which offers ample functionalities that range from modifications to input geometries with an internal CAD module over unstructured mesh generation and advanced CFD solver algorithms to real-time post-processing of results. Matlab was employed for further post-processing steps that encompassed more than individual simulations.

4.1. Meshing

At the beginning of every CFD simulation stands the spatial discretization of the solution area into a computational mesh. This mesh is the heart of the simulation and largely determines the quality of the solution as well as the convergence behaviour of the calculation. In the following, the essentially consistent approach to generate the hybrid meshes which were used for the purpose of this thesis is described. More detailed, simulation specific meshing choices are provided where necessary further down in the document.

4.1.1. General mesh description

All simulations which are discussed in the remainder of this thesis were run on hybrid meshes (see Section 2.2.2) generated with the automated mesher of STAR-CCM+. They have predominantly hexagonal surface grids and may be volumetrically separated into an unstructured 'core' mesh consisting of arbitrary polyhedral cells, and structured, wall-resolved prism layer meshes at no-slip walls where boundary layer flow shall be accurately captured (see Figure 5.3 for illustrations). More specifically, the latter includes all outer surfaces of the canopy whereas the inner sides are treated as no slip-walls, which do not require wall-resolved meshes, in order to reduce the complexity of the mesh and to reduce total cell counts.

4.1.2. Mesh parameterisation

The grid generation is often seen as the most labour intensive part of a CFD study. While the automated mesher of STAR-CCM+ allows reducing the time a human would take to mesh a given geometry from several weeks to less than an hour, this requires that the desired mesh must be fully predefined when handed to the mesher. In order to pave the way for easily iterable design changes to the geometry – as for instance required for the conducted bleed air spoiler parameter studies – as well as to achieve a high degree of reproducibility and order in the mesh, a parameterisation of these meshing directives is virtually indispensable.

Proper mesh convergence studies, more closely described in Section 4.1.3, require a very homogeneous, uniform approach to grid refinement even for unstructured meshes. A useful approach to meet this expectation is to define as many meshing constraints as possible relative to a single reference value, the "base size", which will be frequently referenced in the following and is at the order of magnitude of the respective chord length.

Where no local mesh constraints are imposed, the automatic mesher aims to satisfy the general mesh constraints summarised in Table 4.1. For a closer description of the parameters, the reader is referred to the user manual of STAR-CCM+.

Table 4.1: General mesh constraints

Parameter	Value
Base Size	*specific to case family*
Target Surface Size	$0.05 \cdot \mathrm{base}$
Minimum Surface Size	$y^+ = 1$
Surface Curvature	120 Pts/circle
Surface Growth Rate	1.3
Volume Growth Rate	1.2
Maximum Tet Size	$5 \cdot \mathrm{base}$

Prism layer grid generation

In general, the prism layer mesh which is generated on the outer surface of the canopies must fulfil the following criteria (from [7]):

- the first cell layer height must fall within the viscous sub-layer region, ie. $\Delta y_1^+ \leq y^+ = 1$

- the prism layer mesh should be tall enough to cover the full height of the boundary layer

- at the transition to the core mesh, the cell aspect ratio should be close to 1

- cell growth in the wall-normal direction should not exceed factors of 1.2

The automated mesher requires three inputs for the prism layer generation: (1) the first cell layer thickness, Δy_1, (2) the number of prism layers in the wall-normal direction, n, and (3) the total prism layer thickness, T. Since the first cell layer height, Δy_1, should satisfy $y^+ = \frac{\Delta y u_\tau}{\nu} \leq 1$ (see eq. (4.1)), appropriate values must be found for n and T to guarantee that the aforementioned criteria are satisfied.

$$\Delta y_1 = \frac{1 \cdot \nu}{u_\tau} \tag{4.1}$$

The minimum required height of the prism layer mesh, $T_{\min}$, can be approximated with the boundary layer height that develops on a flat plate with the same chord length at the same Reynolds number, which according to [33, p. 513] may be analytically estimated with

$$T_{\min} = \delta_{\text{flat plate}} \approx \frac{0.37c}{Re^{0.2}} \ . \tag{4.2}$$

From $T_{\min}$, the height of the first cell layer, Δy_1, and the desired thickness of the last layer, Δy_N, which should be equal to the cell width (target surface size) in order to obtain a final aspect ratio of unity, the required growth rate, s, follows from

$$s|_{T_{\min}} = \frac{\Delta y_1 - T_{\min}}{\Delta y_N - T_{\min}} \ . \tag{4.3}$$

When $T_{\min}$ is small or Δy_N is large, values of $s > 1.2$ would be required to reach an aspect ratio of unity at the top of the prism layer. If this is the case, s is set to 1.2 and the corresponding suitable T is instead computed using

$$T|_{s=1.2} = 6\Delta y_N - 5\Delta y_1 \ . \tag{4.4}$$

The number of prism layers, n, is then found with

$$n = \log_s \left(\frac{T}{\Delta y_1}(s-1) + 1 \right) \ . \tag{4.5}$$

Note that whether the first cell layer indeed falls into the viscous sub-layer everywhere must be checked in a given simulation using an adequate y^+-plot to guarantee the soundness of the solution.

Local mesh controls

Local mesh controls may be used to override the default values from Table 4.1 in order to obtain tailored mesh characteristics on specified curves, surfaces or within predefined volumetric regions. Important controls which were imposed on the meshes and their implications are listed below:

- Curve controls

    - Prism layer edges: Wherever a prism layer does not end at a wall, such as at the kite's intake or at the BAS opening, it must gradually shrink in height all the way down to Δy_1 as it transitions to a fully unstructured mesh (see Figure 5.3a for a close-up of such an edge). As cell layers are successively removed from the top inwards, the surface cell size should ideally always match the thickness of the last layer above it in order to maintain a cell aspect ratio of unity at the transition to the core mesh, and thus proper integrity and quality of the grid. This can only be established by forcing a grid spacing of $y^+ = 1$ around all unbounded edges of the prism layer, which inevitably comes at the cost of many extra cells.

    - Canopy folds: In order to better capture flow phenomena at the intersections of the canopy chambers and also to allow for a high-quality transition between their prism layer meshes, canopy folds have a target spacing of 0.5% of the base size.

    - BAS folds: Similar to the canopy folds, folds on the bleed air spoiler have a target spacing of 0.1% of the base size.

- Surface controls

    - Canopy inside: Sections of the investigated canopies are hollowed out in order to allow for a proper assessment of the implemented intakes and bleed air spoilers. Meanwhile, the relevance of the kite-interior flows is largely limited to the pressure levels that develop in the separate chambers, which does not call for proper reproduction of inner boundary layers. Prism layers are thus suppressed on the interior (including the bleed air spoiler flap).

    - Trailing edge: In line with the grid quality guidelines [7] published by ERCOFTAC[1], the resolution of the trailing edge is set to 0.1% of the chord length or base size (whichever is smaller) if no smaller mesh constraints apply.

    - Canopy wake: Wake refinement to target cell sizes of 5% of the base size is applied up to three chord lengths from the kite's surface in alignment with the inflow.

    - Inlet, outlet, symmetry plane: A target cell size of 400% of the base size is applied to the symmetry plane and far-field.

- Volumetric controls

  – BAS region: Where bleed air spoilers are applied, additional volumetric refinement is imposed in its direct surrounding and wake in order to better capture the unsteady phenomena induced by it (see Figure 5.3d for an illustration). The chosen resolution depends on the separate mesh convergence studies described in the respective chapters.

4.1.3. Description of mesh refinement studies

The mesh convergence studies which were carried out in the context of this thesis followed the best-practice guidelines [8] which are encouraged and published by the Journal of Fluids. This section aims to provide a sufficient degree of understanding of the grid convergence index, GCI, and the evaluation steps that lead up to it. For an in-depth description of the process and editorial examples, the reader is referred to the official guideline.

The method proposed in said document is based on the Richardson Extrapolation method. Scrutinised by numerous researchers, it is acknowledged that the mesh resolution may not have a smooth, monotonic influence on the extrapolated values of the predicted variables and that time-dependent computations will entail further sensitivities on the discretizations in space and time. However, the authors assert "it is currently the most reliable method available for the prediction of numerical uncertainty" [8].

Step 1: Definition of a representative cell size, h

First, a representative cell size, h, is established to enable a proper comparison of the analysed meshes. When integral quantities are monitored such as $C_R = \sqrt{C_L^2 + C_D^2}$, which is one of the key metrics in this study, it is advised to use an average global cell size. For three-dimensional grids, h is thus defined as follows:

$$h = \left[\frac{1}{N} \sum_{i=1}^{N} (\Delta V_i) \right]^{1/3} , \tag{4.6}$$

where N is the total number of cells and ΔV_i are their individual volumes, of which the sum is always approximately constant and equal to the fluid volume in the domain.

Step 2: Generation of simulation results with three grids of varying resolution

Three grids of distinctively different resolutions are required for a given mesh convergence study. By varying the overarching base size parameter, as described in Section 4.1.2, which has a proportional effect on the majority of the imposed mesh controls, a rather systematic grid refinement is achieved even though the mesh itself is unstructured. In the course of this, the grid refinement factor, $r = h_\text{coarse}/h_\text{fine}$, should preferably be be larger than 1.3.

When running the simulations, at least one monitoring variable, ϕ, should be chosen and recorded that is central to the study and of which convergence with the degree of spatial discretisation is expected. For the convergence studies in this thesis, ϕ was taken to be $C_R = \sqrt{C_L^2 + C_D^2}$. The simulations were set up on base cases of the respective models without bleed air spoilers where steady solutions could be achieved with a RANS solver, and values of ϕ were recorded after asymptotic convergence of ϕ to $\pm\epsilon = 10^{-4}$ over 50 iterations was reached.

Step 3: Calculation of the order of the method, p

With $h_1 < h_2 < h_3$, $r_{ij} = h_i/h_j$ and $\epsilon_{ij} = \phi_i - \phi_j$, ϕ_k describing the converged solution on the k^th grid, the order of the method, p, is found using fixed-point iteration on the following set of equations after a rough initial guess of p (e.g. $p = 3$):

$$p = \frac{1}{\ln r_{21}} \left| \ln \left| \frac{\epsilon_{32}}{\epsilon_{21}} \right| + q(p) \right| , \tag{4.7a}$$

$$q(p) = \ln\left(\frac{r_{21}^p - s}{r_{32}^p - s}\right),$$ (4.7b)

$$s = 1 \cdot \text{sign}\left(\frac{\epsilon_{32}}{\epsilon_{21}}\right).$$ (4.7c)

Step 4: Calculation of extrapolated values, ϕ_{ext}^{21}

Knowing the observed order of the method, p, one can extrapolate ϕ to the expected value for an infinitely fine mesh.

$$\phi_{ext}^{21} = \frac{r_{21}^p \phi_1 - \phi_2}{r_{21}^p - 1}$$ (4.8)

Step 5: Calculation of error estimates and the fine-grid convergence index, GCI_{fine}

In demonstrating the degree of uncertainty of the tested meshes, one shall finally compute and report the following error estimates as well as the fine-grid convergence index along with the order of the method:

$$\epsilon_a^{21} = \left|\frac{\phi_1 - \phi_2}{\phi_1}\right|, \qquad\qquad \epsilon_{ext}^{21} = \left|\frac{\phi_{ext}^{21} - \phi_1}{\phi_{ext}^{21}}\right|,$$ (4.9a, 4.9b)

$$\text{GCI}_{\text{fine}}^{21} = \frac{1.25\epsilon_a^{21}}{r_{21}^p - 1}.$$ (4.10)

4.2. Physics and Solver Settings

Next to the computational mesh, the physics and solver settings form the second pillar of a CFD simulation. To this end, a vast range of models to make a selection of exists, each entailing even more numerous and often complex extensions, which must be tailored to the study at hand. Fortunately, no single adequate selection of models and settings exist, but it is rather the validation phase that shows if the selected setup performs well.

The choice of settings for the present thesis was made in part on the basis of previous numerical studies on ram-air parafoils for AWE applications by Folkersma et al. [15], Fogell et al. [14], and Ghoreyshi et al. [19], from which, for instance, values for the turbulence intensity and turbulent viscosity ratio were transferred. It is acknowledged that the aerodynamic phenomena in this thesis may vary significantly from those investigated in the aforementioned papers, and that transferability is not inherently guaranteed. Meanwhile, the SST $k-\omega$ turbulence model and the $\gamma - Re_{\theta_t}$ transition model were selected following the advice of the professor for mechanical engineering at DTU, where this is the standard choice. The fluid properties were set in line with the International Standard Atmosphere (ISA).

Table 4.2: Summary of key physics settings imposed within Star-CCM+

Parameter	Symbol/Definition	Value
Gas Model		Constant density
Density	ρ	$1.225\,\text{kg}\,\text{m}^{-3}$
Dynamic viscosity	μ	$1.835\,\text{kg}\,\text{m}^{-1}\,\text{s}^{-1}$
Turbulence model		SST $k-\omega$
Near-wall treatment		"all y^+"
Turbulence intensity	I	2%
Turbulent viscosity ratio	ν_t/ν	10
Transition model		$\gamma - Re_{\theta_t}$
Time discretisation		Steady/Implicit Unsteady (2$^{\text{nd}}$-order)
Solver type		Segregated

The guidelines for proper transition modelling in the user manual of Star-CCM+ contain a dedicated section for large Re-number (i.e. $Re > 10^6$), low Mach-number flows, where the standard $\gamma - Re_{\theta_t}$ model is said to lead to premature transition behaviour. Whereas the validation case does not fall in this category, the flows simulated for the large power kites discussed in Chapter 6 have Reynolds numbers of approximately $5 \cdot 10^6$ despite a low Mach number $Ma = 0.03$ due to the large chord length. To alleviate this problem for the latter case, some advice is given that requires tuning of the transition onset parameter, $C_{\text{onset 1}}$, which controls the initiation of intermittency, to $\min(4.84, \max(2.2, 1.388 \ln(Re \cdot 10^{-6}) + 0.705))$.

Unless otherwise stated in the text, all simulations were initialised with a constant flow field matching the inflow angle and velocity. Automatic boundary layer initialisation was turned on to speed up convergence with a rough approximation of the turbulent kinetic energy and specific dissipation rates in the near-wall regions.

<div style="text-align: right; font-size: 3em;">5</div>

CFD Validation

5.1. Introduction

A pivotal hurdle in any CFD-based investigation is the validation of the employed simulation setup. Ideally, this is done via comparison to previously collected experimental data or to as many other independent numerical studies on common test cases as possible. Otherwise, little trust can be put into the obtained results due to the sheer endless amount of design choices and potential pitfalls entailed by the simulations.

Since the use of bleed air spoilers for AWE purposes is still novel and unresearched to the extent known, finding a proper basis for this comparison is not a trivial task. However, as mapped out in section 1.2, a number of sophisticated studies have been carried out on the aerodynamic effects inflicted by BAS on regular ram-air parachutes both on experimental and numerical bases. One study which stands out in particular for its many parallels to the present investigation is that by Berexa et al. [3], who investigated a number of BAS configurations in their ability to change the aerodynamic forces and moments acting on a rigid ram-air canopy over a range of angles of attack through a series of wind tunnel experiments. Fortunately, Dr. Keith Bergeron – a co-author of the experiment – kindly provided the used canopy geometry in the form of `.step` files, thus allowing for an exact replication of the wind tunnel setup for the purpose of the validation.

It is expected that if a sufficient predictability of the experimental results can be achieved, then the degree of reliability of the configured CFD environment for further simulations on SkySails' kite can be established whilst highlighting potential limitations of the obtained results. A critical purpose of this phase was also the iterative development and fine-tuning of the setup through a gradual increase in complexity of the mesh as well as of the physics and solver settings. However, only the final version is discussed.

5.2. Experimental Setup

In the following, the experimental setup and procedure as employed by Bergeron et al. are briefly described. Then, the adaptation of the setup for the CFD simulations is explained.

5.2.1. Description of comparative wind tunnel study

The underlying wind tunnel tests were carried out by Berexa and her team in the subsonic wind tunnel of the Aeronautics Laboratory at the United States Air Force Academy (USAFA) in Colorado Springs. Some key parameters of the setup are summarised in Table 5.1. The test geometry is a rigid model of a billowed ram-air canopy with an open inlet covering the full span of the leading edge. For the measurements, it is mounted vertically on a support beam located at $0.5c$ which also acts as the pitch axis. To reduce unwanted interactions of the half-model with the boundary layer at the symmetry plane, a splitter plate separates the airfoil from the floor of the wind tunnel by $\sim 5\,\text{cm}$.

Table 5.1: Key parameters of the wind tunnel experiment

Parameter	Definition	Value
Reynolds number	$Re = \frac{U_\infty c}{\nu}$	$5 \cdot 10^5$
Dynamic pressure	$p_0 = \frac{1}{2}\rho U_\infty^2$	$383\,\mathrm{Pa}$
Wind tunnel dimensions	$l \times w \times h$	(4 x 0.9 x 0.9) m
Projected surface area of kite	S	$0.81\,\mathrm{m}$
Chord length	c	$0.3\,\mathrm{m}$
Half-span	b	$0.3\,\mathrm{m}$
range of angles of attack	α	[-6°,24°]

The lift and drag forces as well as roll and yaw moments about the quarter chord point in the symmetry plane were recorded for five spoiler configurations in addition to a baseline model without spoilers. To this end, between 3-6 measurement series were run for each model where α was incremented in small steps of 1° or 2° for the range of α's listed above. Each α was held for 30 s whilst measuring at 300 Hz before moving on to the next. For further details on the study, the reader is referred to the corresponding paper [3].

5.2.2. Adaptation for CFD

For the simulation environment, the original dimensions of the experiment were replicated to the extent possible and reasonable (see Figure 5.1). The only two notable differences between the physical and numerical domains are the definition of the tunnel floor as a slip-wall/symmetry plane (shown in blue), and the omission of the splitter plate. In order to allow for the investigation of arbitrary values of α, the Star-CCM+-internal underlying CAD model allows for a free rotation of the canopy around the center of the supporting beam. By means of the meshing pipeline inherent to Star-CCM+, any meshing constraints imposed on the geometry are readily transferred to new fluid domains, allowing for easy adjustments of α and modifications with BAS. The physics settings are as described in section 4.2.
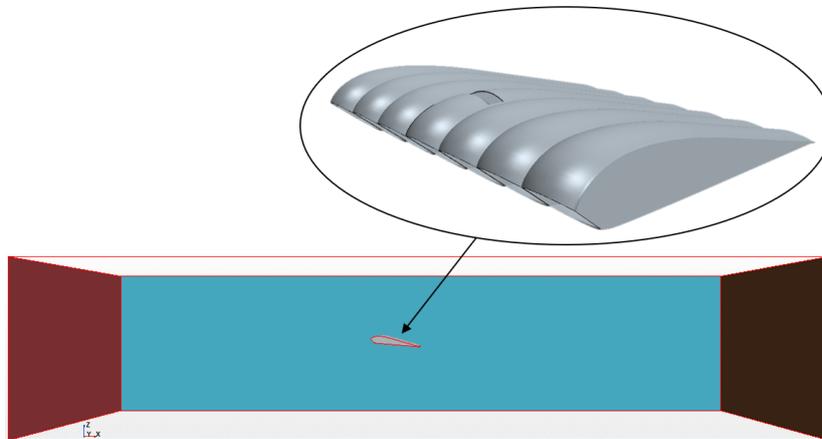
Figure 5.1: Adaptation of the wind tunnel setup for the CFD simulations

Out of the five spoiler configurations which were tested by Berexa et al., three were chosen as candidate models for the CFD validation, namely 'Baseline_R', 'Model_4_25_', and 'Model_4_25_5_25_', which are also shown in Figure 5.2c.

(a) 'Baseline_R' (b) 'Model_4_25_' (c) 'Model_4_25_5_25_'
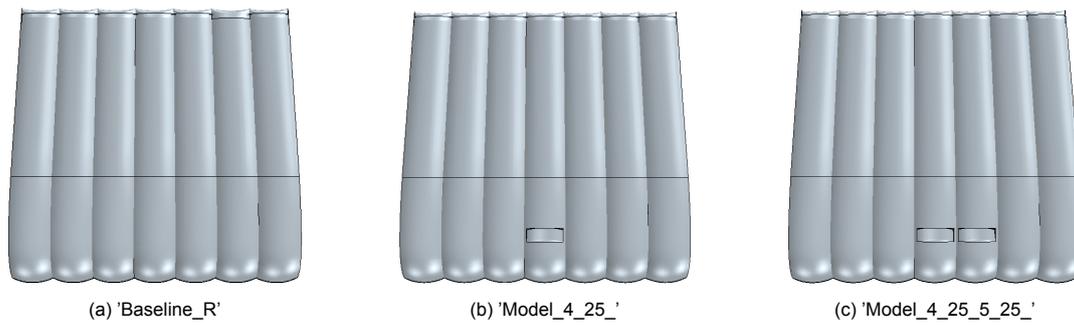
Figure 5.2: The selection of three geometries from Berexa et al. [3] which are used for the CFD validation

Meshing

For a general description of the meshing process and standard mesh constraints, the reader is referred to Section 4.1 of this report. The following will rather highlight some distinctive design choices for the mesh of the validation geometry. A number of impressions of the grid are depicted in Figure 5.3

(a) close-up of prism layer boundary

(b) close-up of trailing edge

(c) close-up of bleed air spoiler

(d) big picture (refined BAS affected region in orange)
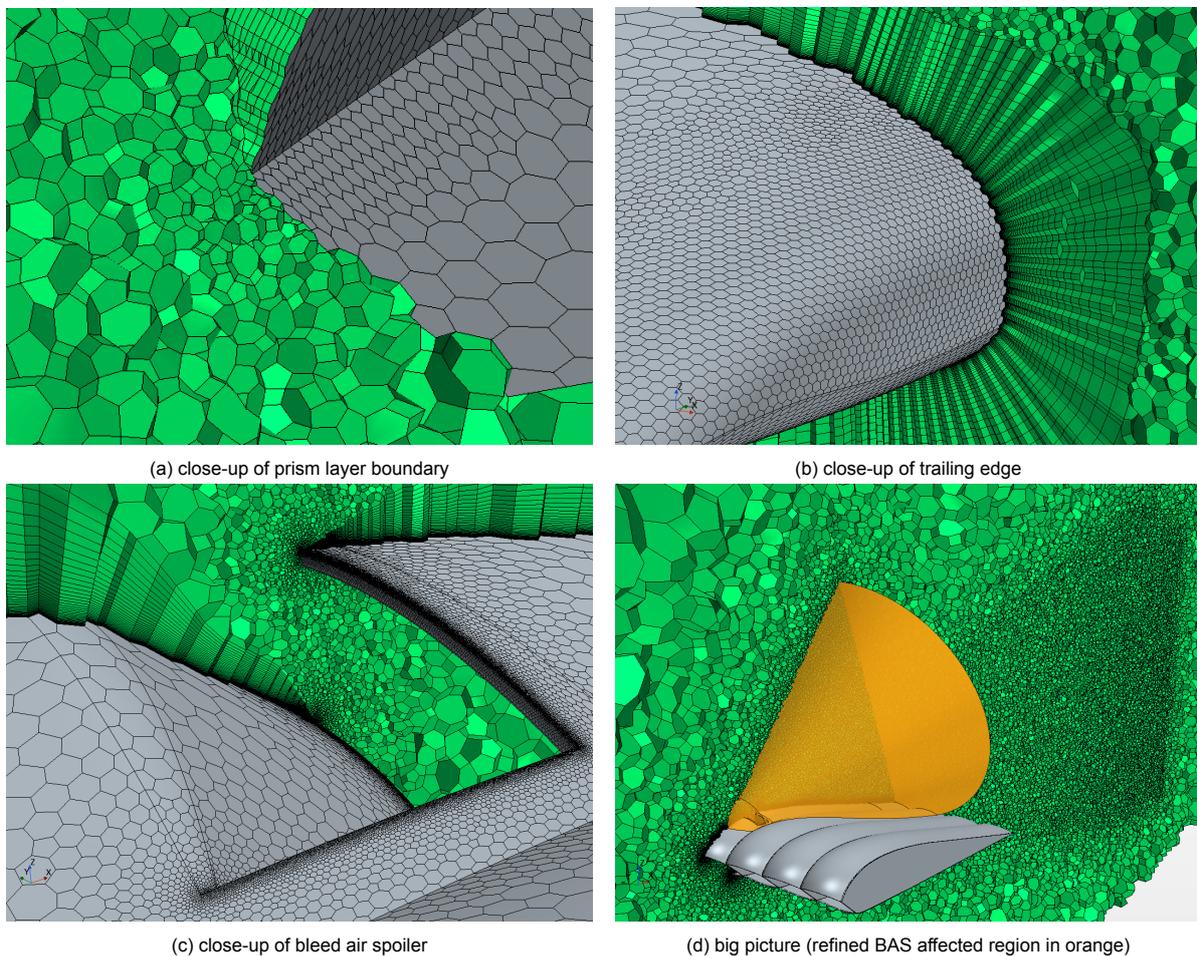
Figure 5.3: Key visuals of a representative validation case mesh

Like the physical models used in the experiment, the CAD representations have thin walls of finite thickness. Parametrised prism layers as defined in Section 4.1.2 cover the complete surface of the canopy except for its blunt edges (i.e. at the leading edge and around the bleed air spoiler). The wake region is more closely refined up to a length of $1\,\mathrm{m}$ ($\sim 3.3c$) beyond the trailing edge of the canopy. In order to reach a sufficient fidelity in the regions subject to the turbulent phenomena induced by the bleed air spoiler(s), such configurations include a further spatial refinement in the region above the

outlet as shown in Figure 5.3d.

5.3. Sensitivity Studies

The first phase of the CFD validation focused on a thorough investigation of the reliability of the created setup through three sensitivity studies to provide insights on the convergence behaviour with the imposed mesh resolution and the employed time-step. First, a global mesh convergence study was carried out as explained in Section 4.1.3 on a baseline case without bleed air spoilers at $\alpha = 5°$, where largely steady flow was expected, using steady RANS calculations in order to identify a reasonable setting for the base size parameter which largely governs the overall resolution of the mesh. Then, a proper time-step was identified and the sensitivity on the time-discretisation settings was investigated after imposing the identified resolution settings on a configuration with a bleed air spoiler. In order to double-check that the grid convergence behaviour was not just limited to small angles of attack with little degree of flow separation, a second, more comprehensive mesh convergence study was carried out over a range of α's at the established time-discretisation settings. Finally, additional volumetric refinement was applied to the directly affected wake region of the bleed air spoiler to determine if the global mesh resolution settings established for the baseline cases are still adequate when additional turbulence is introduced by the BAS.

5.3.1. Initial global mesh refinement study

As described above, the purpose of the initial, precursor mesh refinement study was to make an educated first estimate of the required mesh resolution in preparation for the unsteady simulations which were carried out in the subsequent sensitivity studies. In order to rule out effects of the temporal discretisation on the solution, this set of simulations was run using steady RANS computations, which requires a limited degree of unsteadiness in the flow field in order to converge. This was expected to be met on the baseline case (i.e. without bleed air spoilers) at $\alpha = 5°$, which formed the basis of this evaluation. Following the standard approach for mesh refinement studies explained in 4.1.3, three simulations of uniformly varying resolution and otherwise identical setup were initialised with a constant flow field matching the inflow conditions and run for 2000 iterations. Figure 5.4 shows the results of the simulations.

(a) Solution histories for the tested mesh resolutions   (b) Convergence with respect to the respective char. cell size, h
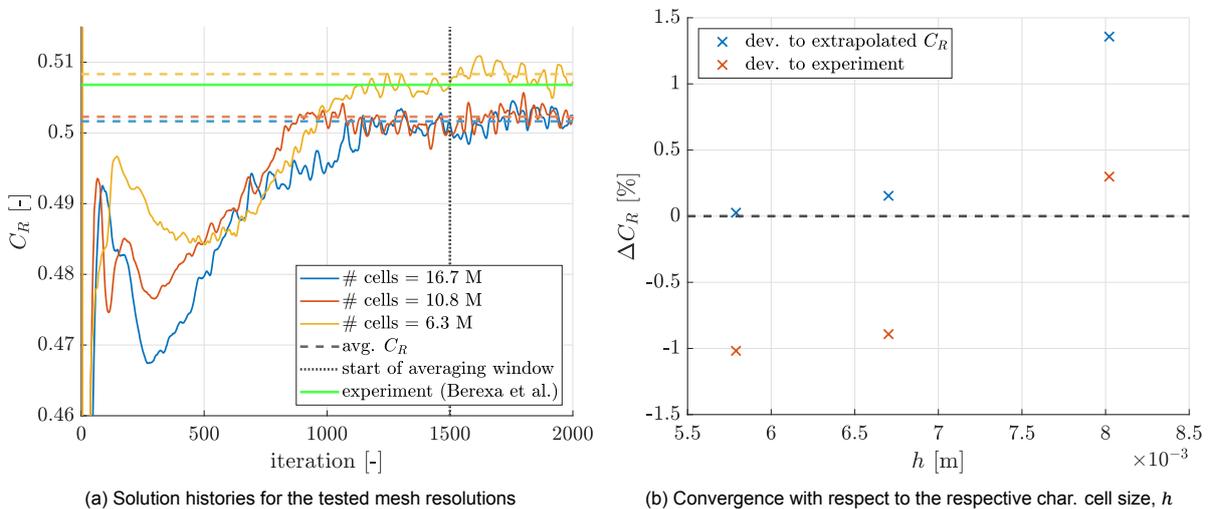
Figure 5.4: Results of initial global mesh refinement study

Plot 5.4a illustrates that all three simulations take approximately 1200 iterations before they begin to plateau, during which the loading on the wing, characterised by $C_R = \sqrt{C_L^2 + C_D^2}$, gradually builds up after some more chaotic first 300 iterations. It is important to point out that the oscillations which persist after the simulations plateau around their respective mean are not physical, as one would observe in an unsteady simulation, but rather demonstrate that there remains some degree of unsteadiness in the flow field which the RANS solver is not able to time-average out, thus hindering it from converging

to a constant solution. Note that an increase in resolution consistently leads to a decrease of the predicted C_R. However, the difference between the solutions of the finest and the coarsest mesh still only amounts to 1.4%, suggesting that all grids have a relatively high degree of resolution which puts them in the asymptotic branch of the convergence characteristic, where further refinement still leads to higher solution accuracy but coarse meshes are also not prone to diverge.

Because no constant solutions are found, results were averaged out over the last 500 iterations, which seem to fluctuate around rather constant mean values, for the purpose of the evaluation. Table 5.2 contains a summary of the resultant findings.

Table 5.2: Results of initial global mesh convergence study on baseline validation case; $\alpha = 5°$

base sizes$_{1,\,2\,\&\,3}$ [m]	0.075, 0.15, 1
$N_{1,\,2\,\&\,3}\ \left(\cdot 10^6\right)$	16.7, 10.8, 6.3
r_{21}, r_{32}	1.16, 1.20
$C_{R,1}, C_{R,2}, C_{R,3}$	0.5016, 0.5023, 0.5083
$\epsilon_a^{21}, \epsilon_a^{32}$	0.0013, 0.0120
$\epsilon_{ext}^{21}, \epsilon_{ext}^{32}$	0.0003, 0.0015
p	12.1
$\mathrm{GCI}_{\mathrm{fine}}^{21}, \mathrm{GCI}_{\mathrm{fine}}^{32}$	0.033%, 0.192%

As expected from visual inspection of Figure 5.4, the approximate and extrapolated relative errors ϵ_a and ϵ_{ext} are very small at 1.2% or lower. This also reflects in the fine grid convergence indices, which are relatively low in comparison to those in the example case discussed in [8] between 0.5% and 2.2%. Also, no oscillatory divergence behaviour is observed which would reflect in a non-monotonic dependence of ΔC_R on the resolution. The spatial discretisation order of ~ 12 which the procedure estimates is far from the actual, 2nd order scheme that is actually employed in the simulations. This is presumably in part due to the high closeness of the results, which makes subtle differences have a very pronounced effect on the observed order. Looking at Figure 5.4b, about 0.2% more on the medium-resolved mesh, for instance, would have led to an almost linear relationship between ΔC_R and h, leading to a significantly lower p. The limitation of the procedure for low values of ϵ is also something that the authors of the guideline acknowledge in [8]. They however state that "the results may be reported as such".

One may argue that coarser meshes should also be tested to establish whether a sufficient degree of accuracy could be reached at even less computational expense. However, when increasing the base size further the cell count quickly reaches an apparent lower limit of about $5 \cdot 10^6$ cells due to the independent mesh constraints which had to be imposed on the mesh in order to maintain certain minimum mesh quality characteristics. Examples of such constraints that are critical to the mesh's integrity but also entail significant implications on the cell count are those imposed on unbounded prism layer edges, the trailing edge and the wall-normal layering of the boundary layer (see Section 4.1.2).

It is important to acknowledge that because of the mesh constraints that are unrelated to the base size, the grid refinement which is applied to the simulation is not entirely uniform. However, because their definitions follow the strict, high standards established in the grid quality guidelines [7] published by the European Research Community on Flow, Turbulence and Combustion (ERCOFTAC), it is assumed that they do not call for additional sensitivity studies.

Based on the findings of this initial convergence study, a base size of 0.15 m (i.e. matching the resolution of the medium-coarse mesh) is picked for the following time-discretisation sensitivity study. While this entails about 70% more cells per simulation than the coarsest tested mesh, which arguably also performed well, there are this way grounds to assume that some safety margin with respect to the resolution exists.

5.3.2. Time-discretisation sensitivity study

The temporal discretisation can have an equally strong impact on the accuracy of the solution as poor spatial discretisation. Special care must therefore be taken when establishing the associated settings for unsteady simulations. This applies particularly to the chosen time-step, Δt, and the temporal discretisation order, both of which will be discussed in the following.

Choice of Δt

As explained in Section 2.2.3, in simulations with first-order temporal discretisation schemes, Δt should be chosen such that the Courant number is less than 1 everywhere in the flow field, or rather everywhere where one wants to accurately reproduce any unsteady aerodynamic behaviour. For the present case setup, a Δt as small as $8 \cdot 10^{-5}$ s is required to satisfy this criterion in the majority of the flow region (see Figure 5.5b). Note that even at this point, the maximum Courant number is 880 due to the extremely fine surface resolutions around open prism layer edges, where cells are as small as $y^+ = 1$ in all directions.

Fortunately, higher order temporal discretisation schemes allow for the choice of significantly larger time-steps that make CFD simulations which are to reproduce time periods that capture a sufficiently high number of flow overs feasible. Turbulence of smaller time scales is then time-averaged over a given time-step. If this is applied properly, averages over a given time-window should yield the same outcome, regardless – within reasonable limits – of the chosen Δt.

A Δt of 10^{-3} s was identified as a promising candidate based on an assessment of several Courant number scenes as the two shown in Figure 5.5. The aim was to maintain Cfl-numbers of around 1 or lower in flow regions where important unsteady effects were expected, such as in the wake region of the bleed air spoiler. Meanwhile, it was tolerated that Courant numbers can reach higher values in places where either not a lot of unsteadiness is expected, such as on the lower side of the canopy, or where the capturing of turbulent effects is assumed to be of little importance to the overall solution, such as at the leading edge or in the immediate vicinity of the bleed air spoiler. The result of this compromise is demonstrated in Figure 5.5a, which shows that in regions with finer resolution, such as at the leading edge, the immediate surrounding of the bleed air spoiler or sections of the prism layer, Courant numbers go significantly higher than 1.

(a) $\Delta t = 1 \cdot 10^{-3}$ s   (b) $\Delta t = 8 \cdot 10^{-5}$ s
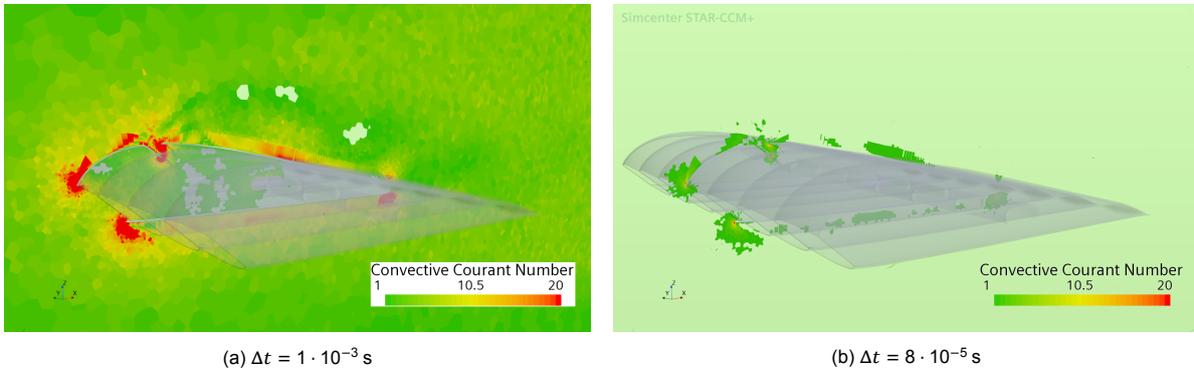
Figure 5.5: Comparison of Courant numbers on a cut-section through the bleed air spoiler at the tested Δt's (light green indicates that $Co < 1$ criterion is satisfied); $\alpha = 5°$; Model_4_24_

In order to test the robustness of the solution with respect to the time-step, a simulation was run with the chosen Δt on the configuration 'Model_4_25_', which has a bleed air spoiler installed on the fourth chamber at 25% of the chord length. After a time interval of 0.5 s, the solver was switched to the smaller time-step of $\Delta t = 8 \cdot 10^{-5}$ s to determine what the effect on the unsteady solution is. Figure 5.6 shows the resulting time-history of the C_R coefficient. Compared to the average C_R for $0.2\,\text{s} < t < 0.5\,\text{s}$ (≈ 25 flow overs) of 0.456, the respective average under the lower Δt comes out at just 0.2% less (i.e. 0.455): a negligible difference bearing in mind the rest of assumptions entailed by the simulation. This confirms that $\Delta t = 0.001$ s is indeed a reasonable time-step to obtain accurate metrics for the evaluation of bleed air spoilers on the validation case.

Figure 5.6: Effect of changing Δt to $8 \cdot 10^{-5}$ s during a running simulation to satisfy $Co < 1$ criterion; $\alpha = 5°$; Model_4_24_
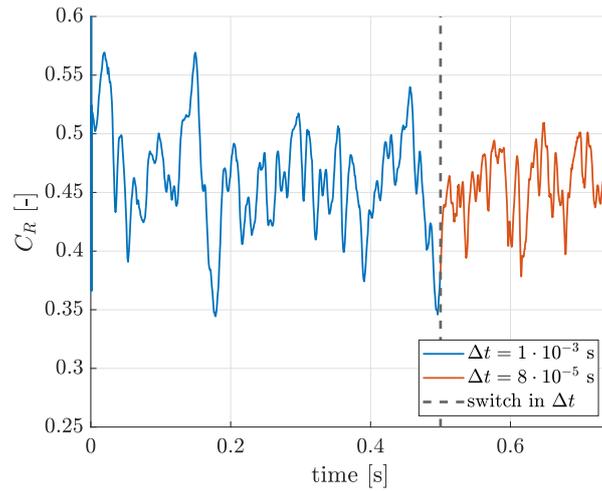
Influence of time-discretisation order

The importance of the correct choice regarding the order of the time-discretisation scheme became clear throughout the development of the simulation environment. Because it proved to be so critical to the accuracy of the solution, it was decided to dedicate a separate discussion to it here.

The results shown in Figure 5.7 depict the simulated time histories of C_L and C_D using two simulation environments which are fully identical aside from the order of the temporal discretisation scheme. Both simulations were initialised with steady RANS calculations covering 500 iterations and were then switched to unsteady, following which the solution was recorded over 0.5 s (≈ 40 flow overs).

(a) 1st-order temp. disc.

(b) 2nd-order temp. disc.
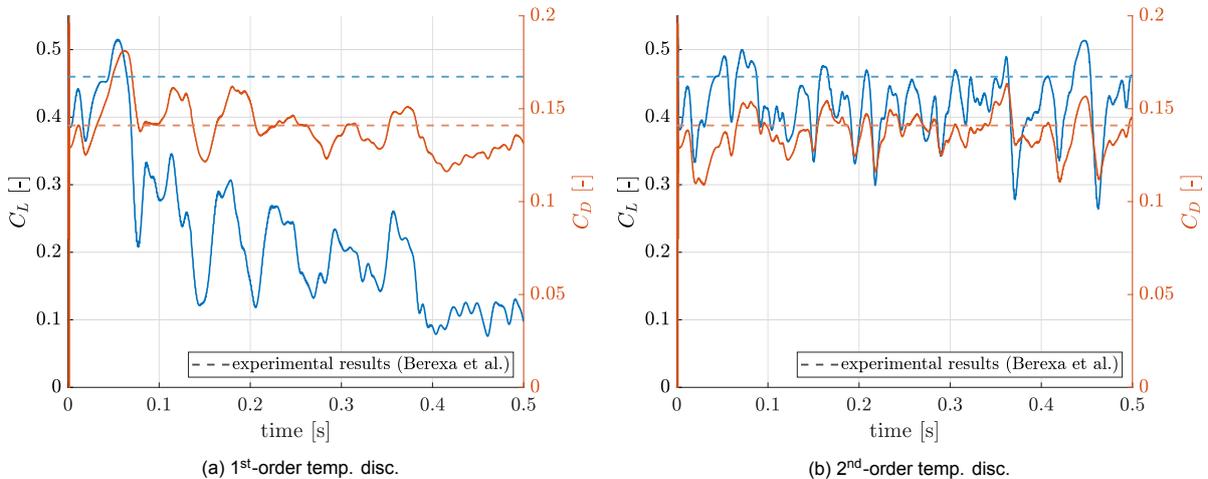
Figure 5.7: Comparison of C_L, C_D time-histories by temporal discretisation order; $\alpha = 5°$; $\Delta t = 1 \cdot 10^{-3}$ s; Model_4_24_

Under 2nd-order temporal discretisation, the resulting lift and drag coefficients show periodic oscillations around a rather constant mean, as expected. Note also that this mean was already predicted quite closely by the steady precursor simulation. Meanwhile, the aerodynamic forces integrated over the canopy decrease steadily over the course of the captured time-interval when a 1st-order discretisation scheme is applied.

The difference between the two solutions also becomes evident when taking a look at the resultant flow fields. Figure 5.8 shows the velocity field time-averaged over the final 0.1 s across a cut-section of the fluid-domain which slices through the bleed air spoiler.

(a) 1st-order temp. disc.   (b) 2nd-order temp. disc.

Figure 5.8: Visual comparison of time averaged flow fields; $\Delta t = 5 \cdot 10^{-3}$ s; Model_4_24_

At 1st-order discretisation, the solver predicts a significantly higher degree of flow separation in the wake of the BAS, which explains why the lift decreases so substantially. At 2nd-order discretisation, the separation bubble stays much more attached to the surface of the canopy, and air is visibly accelerated more d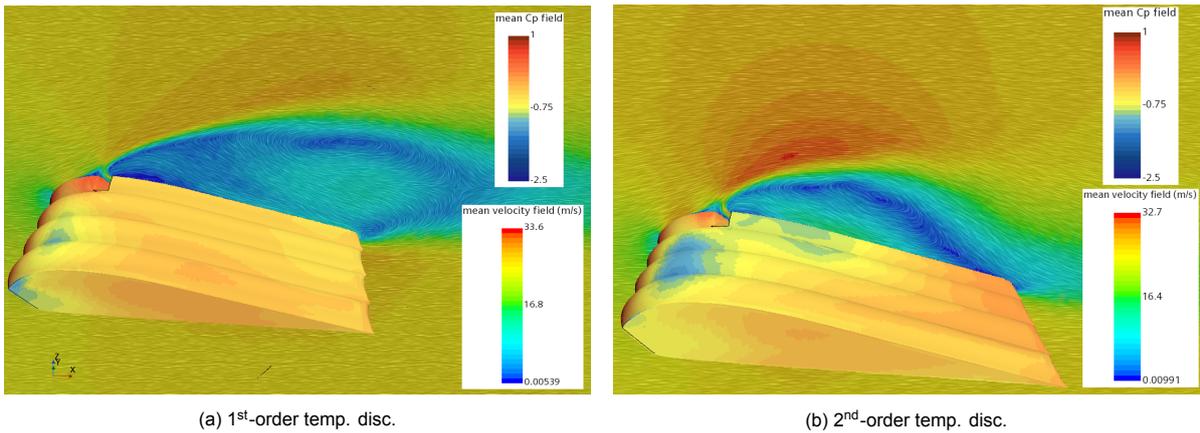ownwards in following the virtual profile as shown by the larger mean velocities over the separation bubble and also the direction of the tracer particles towards the rear of the vortex. The C_P distributions are also lower on the front surface to the sides of the BAS, indicating stronger suction. The same applies to the surface directly downstream of the BAS, where the unsteadiness of the vortex bubble that forms on the kite also seems to contribute some suction effects that are not present on the outer sides of the canopy where the flow remains more attached.

A look at the time-averaged pressure coefficient, C_p, distribution in Figure 5.9, sampled over the same cut-section of the canopy that is depicted in the flow field scenes from above, further supports these observations.

(a) 1st-order temp. disc.   (b) 2nd-order temp. disc.
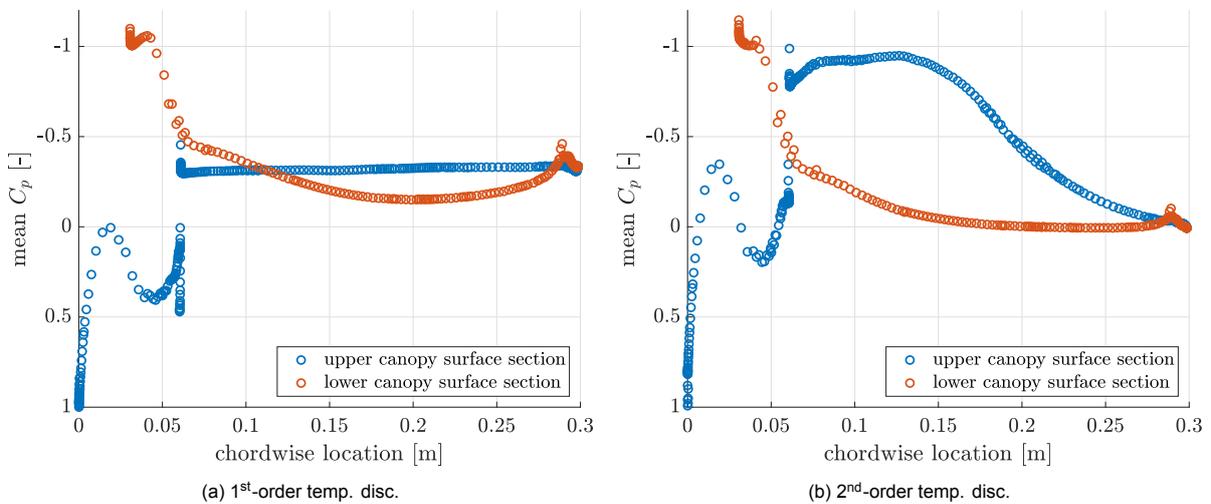
Figure 5.9: Comparison of mean chordwise C_p distributions

Under 2nd-order discretisation, also in front of but particularly behind the bleed air spoiler the average pressure is significantly lower than in the 1st-order discretised simulation, reaching C_p values of almost -1 in contrast to about -0.35. This matches with the stronger suction along the mid section of the kite that was already observed on the mean flow field visualisations. While the pressure distributions on the lower sides do not differ as much, they also reveal that higher pressures are predicted with the lower-order scheme. Nevertheless, when integrated, the area between the curves is evidently much larger with the higher order scheme which also coincides with the observed higher lift values.

Although a 2nd-order scheme entails higher computational expense for a given iteration, it evidently allows for the choice of much larger time-steps which altogether reduces the number of necessary

computations by a factor of 12 (compared to $\Delta t = 8 \cdot 10^{-5}$ s). All following simulations were hence carried out using the higher order scheme.

5.3.3. Comprehensive global mesh refinement study

Having established reasonable time-discretisation settings in the previous step, it is necessary to circle back to a more comprehensive global mesh refinement study to investigate whether the previously identified resolution remains adequate for unsteady simulations, even at larger angles of attack where a significant degree of flow separation is to be expected.
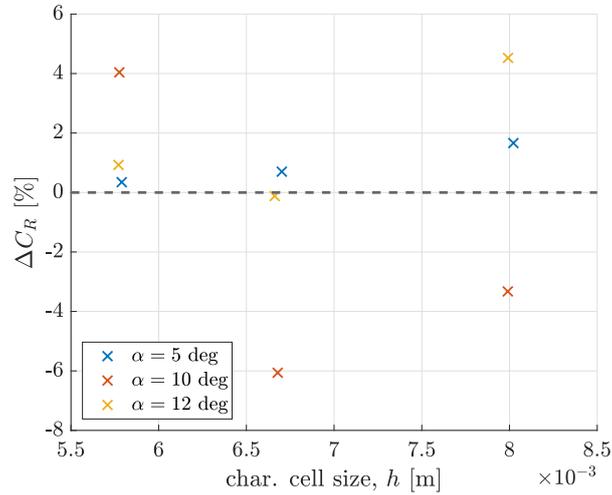
Figure 5.10: Deviation of C_R from the average extrapolated value, $\overline{c_{R,ext}^{ij}}$, with respect to the char. cell size, h

The corresponding simulations were run for either 0.5 s or 1 s, depending on how long it took for the solutions to show somewhat period oscillations around a plateaued mean, and time-averages were calculated using the final 0.3 s of the results. The resultant solution histories are depicted in Figure 5.11.

(a) $\alpha = 5°$

(b) $\alpha = 10°$

(c) α = 12°
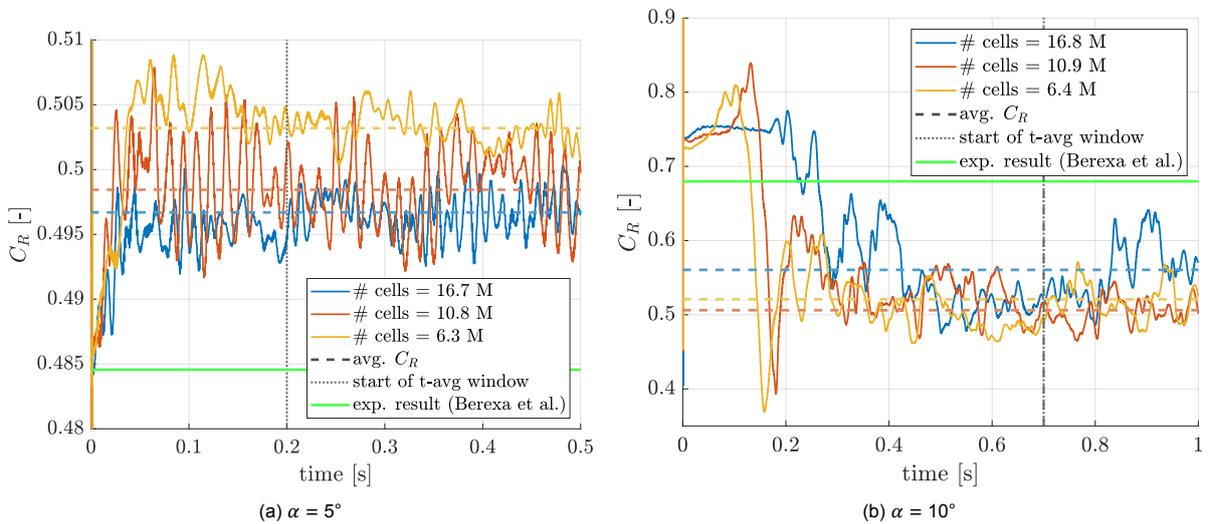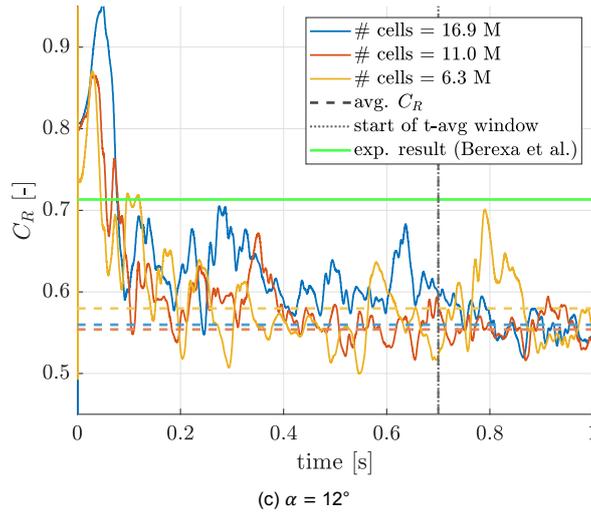
Figure 5.11: C_R time-histories for various grid resolutions at the selected αs; configuration: baseline

Whereas the magnitude of the oscillations is very small in comparison at small angles of attack, unsteady effects produce variations in C_R of up to 0.13 to the time-averaged values at 10° and 12°. If such an oscillation falls within the time-averaging window, results are quickly distorted. While it is acknowledged that even longer simulation times may be required for more dependable results, it is decided that the results are good enough at this point to move on with the evaluation. The corresponding convergence statistics are presented in Table 5.3

Table 5.3: Results of comprehensive global mesh convergence study on validation case

|	$\alpha = 5°$	$\alpha = 10°$	$\alpha = 12°$
$N_1, N_2, N_3 \ (\cdot 10^6)$	16.7,10.8,6.3	16.8, 10.9, 6.4	16.8, 11.0, 6.3
r_{21}, r_{32}	1.16, 1.20	1.16, 1.20	1.16, 1.20
$C_{R,1}, C_{R,2}, C_{R,3}$	0.4969, 0.4987, 0.5025	0.5605, 0.5061, 0.5208	0.5598, 0.5540, 0.5797
$\epsilon_a^{21}, \epsilon_a^{32}$	0.004, 0.008	0.097, 0.029	0.010, 0.047
$\epsilon_{ext}^{21}, \epsilon_{ext}^{32}$	0.006, 0.009	0.023, 0.005	0.004, 0.013
p	3.46	11.3	8.6
$\text{GCI}_{\text{fine}}^{21}, \text{GCI}_{\text{fine}}^{32}$	0.68%, 1.12%	2.93%, 0.56%	0.54%, 1.55%

Because the unsteady simulations paired with larger angles of attack produced larger oscillations in C_R, it comes as no surprise that the resultant convergence indices are slightly higher than after the steady RANS calculations. The only grid which truly stands out from the rest is the finest simulation at 10°, whose mean is estimated to differ from the predicted value by 2.3% due to one larger oscillation towards the end of the time series. However, this is still deemed to be in the tolerable range. Seeing no reason to change the base size parameter for following simulations, it is kept at 0.15 m (medium coarseness).

5.3.4. Investigation of increased bleed air spoiler wake resolution

For simulations on configurations with bleed air spoilers, it is assumed that the mesh resolution identified in Section 5.3.1 remains adequate for all those regions in the flow domain which are not or are little affected by the flow phenomena induced by the bleed air spoiler. This supplemental convergence study instead focuses on the wake region of the BAS which is primarily affected by the added turbulent phenomena to investigate whether further local volumetric resolution can enhance the accuracy of the solution, and to what extent it should be applied for the subsequent angle-of-attack study. To this end, a cone shaped region was defined in the wake of the BAS of Model_4_25_ as shown in Figure 5.3d, where different target cell sizes of 0.5 %, 1 % and 5 % of the chord length, c, were imposed. Simulations were initialised with the usual 500 steady iterations and then run under unsteady conditions

for a simulated time of 0.5 s. Figure 5.12 shows the corresponding C_R time-histories along with their time-averaged values and the experimental result. A summary of the calculated deviations along with an indication of the additional cell count between the simulations is given in Table 5.4.
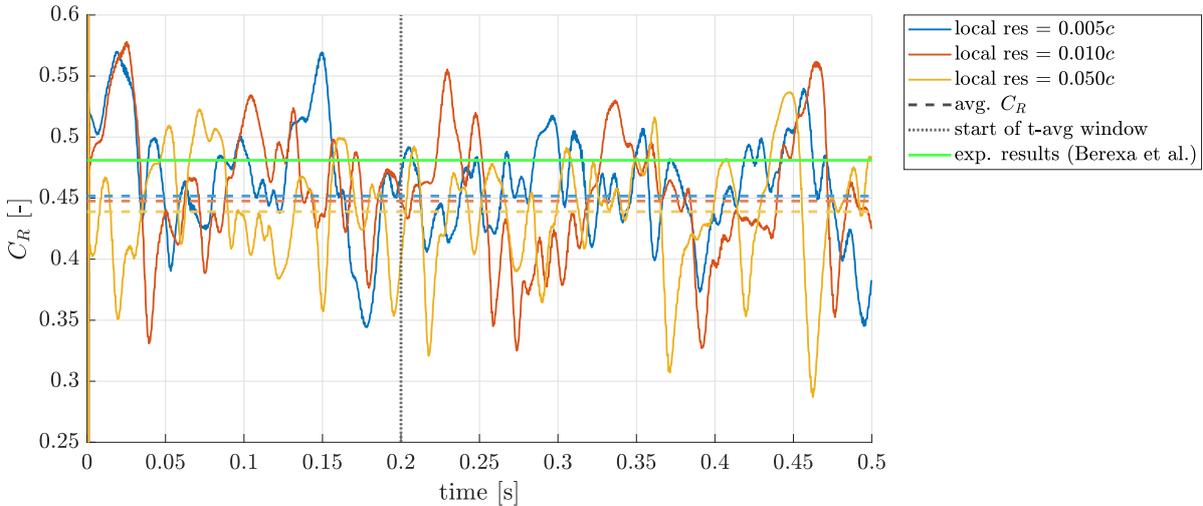

Figure 5.12: C_R time histories for various degrees of resolutions of the BAS wake; $\alpha = 5°$; Model_4_25_

Table 5.4: Convergence behaviour of C_R with increased BAS wake resolution

| cell resolution in BAS wake [m]	$0.05c$	$0.01c$	$0.005c$
cell count difference[1], $N_i - N_{i-1} \left(\cdot 10^6 \right)$	–	0.22	1.41
ΔC_R to finest mesh	-2.8%	-0.9%	–
ΔC_R to experiment	-8.8%	-7.0%	-6.1%

[1] $N_1 = 11.3 \cdot 10^6$

While the oscillations which are predicted by the simulations are rather chaotic, their time-averages again allow for insightful observations. Firstly, note that an increased resolution of the BAS' wake causes the predicted mean to move closer to the result expected from the experiment by Berexa et al. [3]. Meanwhile, whereas the first increase in resolution improves the result by 1.8% at an extra cost of only 220k cells, the next finer mesh already comes in at 1.4M more cells (+10% compared to the baseline) while only leading to a further 0.9% improvement. It is therefore decided that a target size of 1% of the chord length is a sound compromise between additional computational expense and adequate resolution of the unsteady flow field in the bleed air spoiler's wake region. This concludes the sensitivity investigations of the validation case with respect to spatial and temporal discretisation.

5.4. Comparison of CFD results to experiment

Having thoroughly established a proper foundation with respect to the spatial and the temporal discretisation of simulations on the validation case, an angle-of-attack study was prepared for the three selected configurations shown in Figure 5.2. To this end, the baseline case was analysed most extensively over 7 angles of attack in a range of $\alpha = [2°, 20°]$. Meanwhile, four were run in a range of $\alpha = [5°, 12°]$ for each of the two spoiler configurations. The resulting absolute values for C_L, C_D and L/D as well as their respective deviations to the values from the wind tunnel experiment are shown in Figure 5.13. All simulations at angles of attack below 10° were stopped after a simulated time of 0.5 s ($tv_\infty/c \approx 42$ flow overs); simulations at $\alpha \geq 10°$ took a little longer to plateau around steady mean solutions and were thus only stopped after 1 s ($tv_\infty/c \approx 83$ flow overs). Time-averages were taken over the last 0.25 s. The error bars shown in the plot relate to the maximum and minimum recorded oscillations during the time-averaging window.

(a) C_L polars

(b) ΔC_L from exp.

(c) C_D polars

(d) ΔC_D from exp.
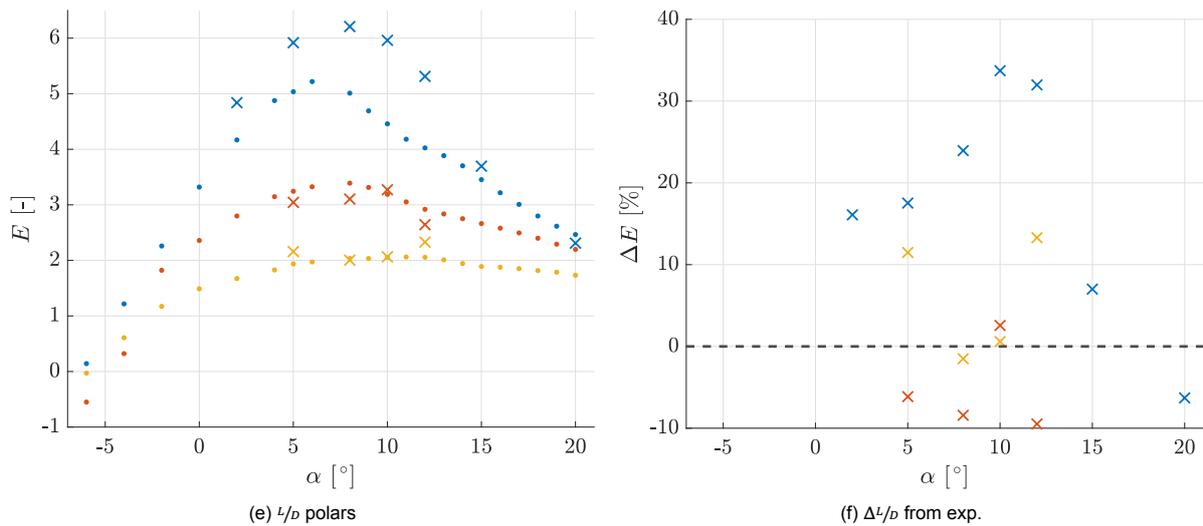
(e) L/D polars

(f) $\Delta L/D$ from exp.
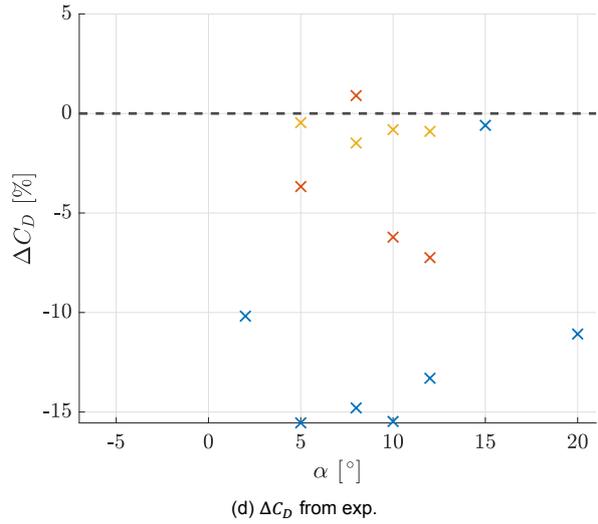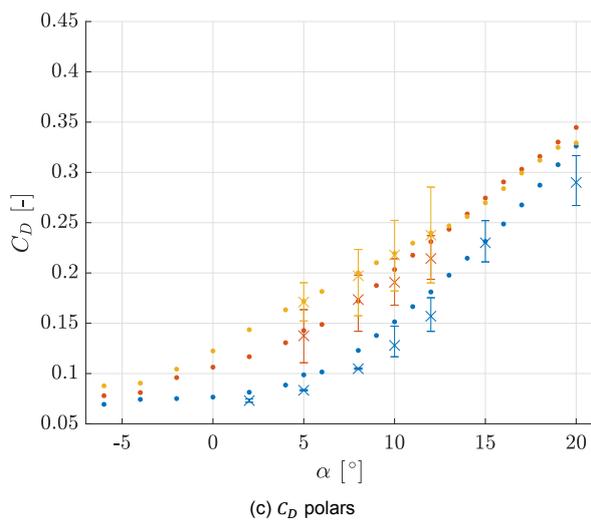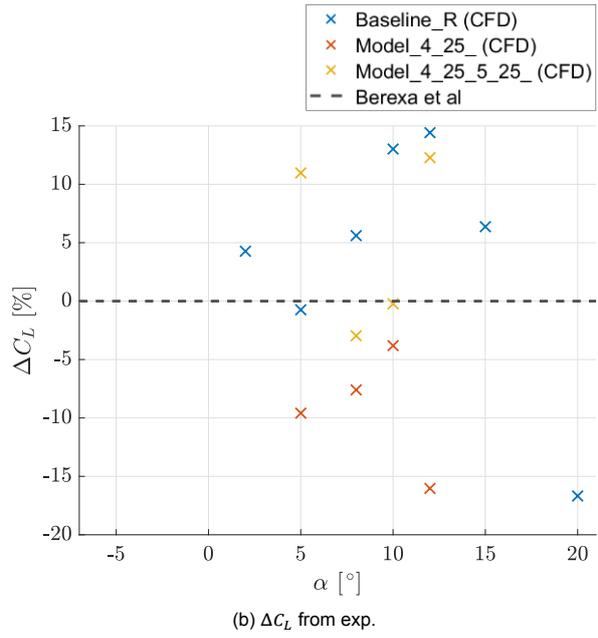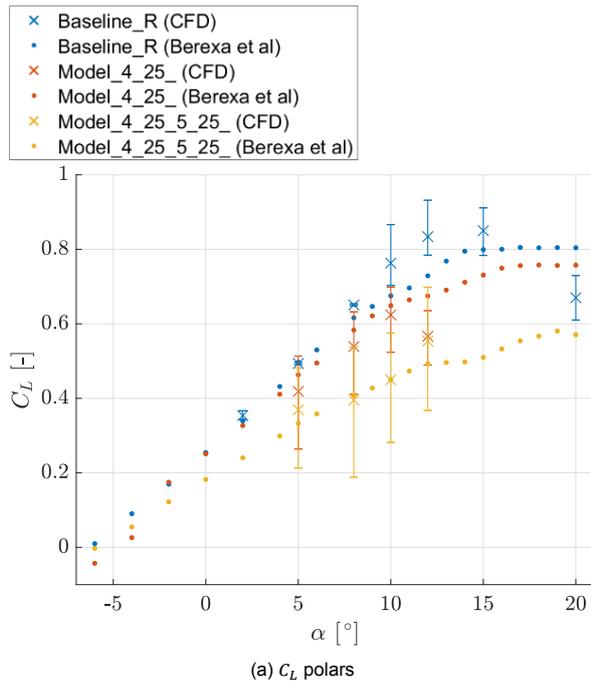
Figure 5.13: Comparison of aerodynamic polars with respect to the tested baseline and BAS case

A number of observations can be made on the basis of these findings. At low angles of attack of up to 8°, the simulations on the baseline case predict C_L values that are in very good to reasonable agreement with the findings from the wind-tunnel experiment ($|\Delta C_L|$ of up to 5.6%). Notice also how the error bars reveal that for $\alpha \geq 10°$, the flow around the canopy becomes increasingly unsteady. At larger α, the CFD simulations tend to predict larger lift forces than were recorded in the experiment for the baseline configuration. Two possible scenarios exist to explain this. The first, intuitive conclusion is that the CFD simulation overestimates the true aerodynamic forces at large α. However, the experimental polar appears to show some inconsistencies in this range as well. At angles >10°, the previously linear C_L polar begins to deviate from a straight line, and the plateau observed at >15° also seems a bit suspicious as no typical stall effect is noticeable. Both a clearly linear increase in C_L until stall and a characteristic stall behaviour are however visible in the CFD results. Although the absolute discrepancies in C_D are not as large, relative differences show that the drag is generally under-predicted by 10% to 15%.

Even though the flow dynamics become significantly more complex with the bleed air spoilers, the canopy performance of the spoiler configurations is predicted with almost the same level of accuracy as that of the baseline configuration. In fact, relative differences in C_L and C_D are often even smaller, although it is important to acknowledge that the latter is mostly due to larger reference values. As a result, the glide ratio is also predicted within an accuracy range of 13%.

Generally speaking, the CFD simulations yielded results that replicated the findings of the wind tunnel study to an acceptable degree of accuracy. The settings identified in the preceding sensitivity studies proved to be suitable for both the baseline configuration as well as the two bleed air spoiler configurations even at higher angles of attack and gave dependable indications of the aerodynamic forces acting on the canopy. In line with these observations, it is concluded that the developed simulation environment produces results with a sufficient degree of reliability and accuracy, which allows for a considerable degree of confidence in the results of the subsequent CFD study on bleed air spoiler effects on large ram-air kites despite the lack of similar experimental data to back them up.

5.5. Lessons learned

The simulations series carried out for the validation phase proved to be useful in a number of different ways. Next to providing safe waters to develop the fundamental framework of the simulation environment, they also helped to highlight its potential pitfalls and thus enabled some key guidelines in preparation for the simulations in chapter 6 on more unfamiliar geometries. The most important lessons learned are summarised below:

- Steady RANS simulations are not viable to obtain dependable indications of the aerodynamic effects induced by bleed air spoilers

- The temporal discretization order has a profound impact on the accuracy of the solution in unsteady simulations. All unsteady simulations should be run with 2^{nd}-order discretisation.

- Relatively large initial Courant numbers (up to 30) in the main flow regions may already yield close approximations of the impact of the BAS spoiler effects, and can thus be used to get a sound approximation of the time-averaged flow field.

- Supplemental local mesh refinement in the bleed air spoiler region is recommended to capture a majority of the added turbulent effects

- Up to 80 flow overs (i.e. $tv_\infty/c > 80$) may be necessary to reach sufficient periodicity in the solution that allows for proper time-averaging. For cases with less flow separation over the airfoil (moderate angles of attack) unsteady sampling over 40 flow overs may also be feasible.

6

Aerodynamic Investigation of Bleed Air
Spoilers on Power Kite

6.1. Introduction

The validation phase of the established CFD tool-chain already provided substantial insights into the requirements for dependable predictions of bleed air spoiler effects on the performance characteristics of parachutes. Meanwhile, the transferability of the observed effects to power kites for Airborne Wind Energy purposes is limited due to key differences between the geometries. Their soft, ram-air canopies have relatively large thickness-to-chord ratios and a significant arc-anhedral component (i.e. arch shape when viewed from the front), both of which are major contributors to the overall stiffness of the kite. Furthermore, whereas the geometry discussed in the validation had a separate intake for each canopy chamber leading to a fully open front, SkySails' kite, for instance, is fed by only one intake per side and ram-air pressure is distributed to the remainder of the kite through a series of so-called cross-ports. Consequently, a study of bleed air spoilers for AWE purposes would be incomplete without an evaluation of a typical geometry.

To the end of this second phase of this study, investigations are carried out on one of SkySails conventional geometries, model 'SK123'. Similar to the previous chapter, first the experimental setup will be more closely mapped out through descriptions of the generated fluid domain and a summary of key inflow and physics settings in the simulation. Then, the sensitivity studies are presented which formed the basis of the bleed air spoiler investigations that follow in the subsequent part of the report. Finally, a summary of the insights that were acquired through the investigations is given. In order to avoid any unnecessary disclosure of the kite's performance characteristics, all aerodynamic forces are generally normalised. Throughout the sensitivity studies, this is generally done on the basis of the respectively finest/most accurate solution. During the main evaluation phase, results are presented with respect to the performance of the baseline configuration without bleed air spoilers, unless otherwise stated in the text.

6.2. Experimental setup

This section aims to provide the reader with a comprehensive overview of the design choices that were made in creating the simulation environment for this case study, which complement descriptions of the general approach to mesh generation and physics and solver settings that were already provided in Chapter 4.

6.2.1. Description of the fluid domain

A key difference between the fluid domains of the present case and the previous validation study is that the kite is analysed on the basis of an undisturbed far-field, and not within a wind-tunnel-like environment. It was given a bullet shape with a radius and distance to the outlet of 30 and 40 chord lengths respectively (see Figure 6.1a). As with the validation case, only one side of the kite is modelled and

a symmetry plane cuts the fluid domain in half. As a result, the simulation of manoeuvres involving side-slip is not possible. The rear half-disk takes the role of the pressure outlet. The remaining walls serve as inlet planes allowing for the test of a wide range of angles of attack.
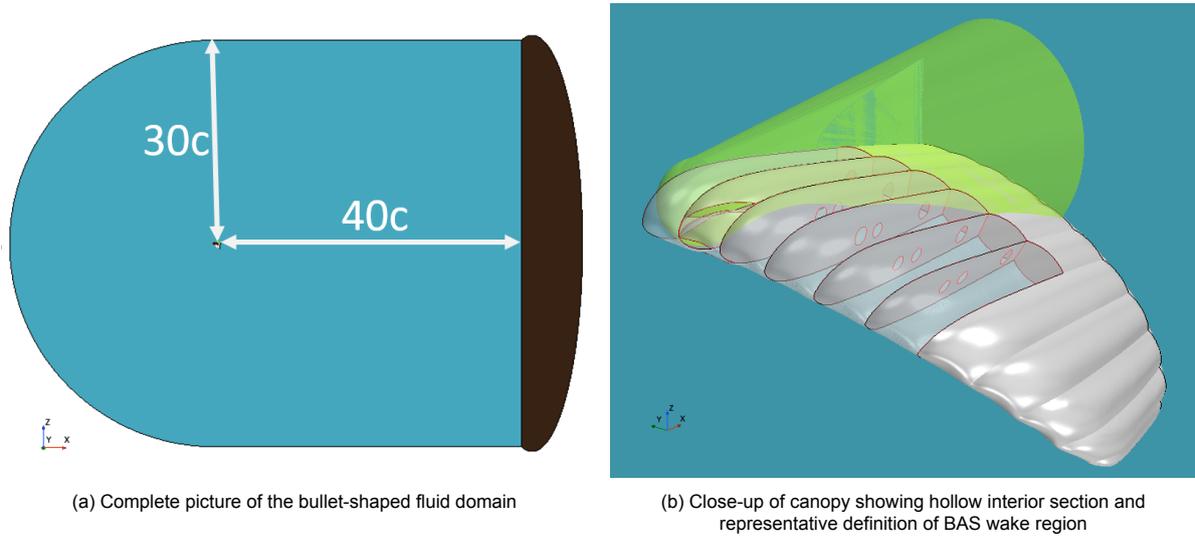

(a) Complete picture of the bullet-shaped fluid domain

(b) Close-up of canopy showing hollow interior section and representative definition of BAS wake region

Figure 6.1: Impressions of the fluid domain

Description of the kite geometry

Model 'SK 123' is one of the several standard configurations currently employed by SkySails with an area of roughly $120\,m^2$ and a maximum chord length of $7.5\,m$. In order to get a realistic approximation of the ballooned shape which it takes under the influence of aerodynamic loading and internal pressure in flight, its baseline geometry was obtained through careful post-processing of a provided tessellated structure file which itself stems from an independent, panel-code based, fluid-structure-coupled simulation by Paul Thedens[1]. It is always important to keep in mind that the rigidity of the kite under all flight conditions is one of the central assumptions of the study.

A large portion of the kite is hollowed out as shown in Figure 6.1b in order to enable the implementation of the intake and bleed air spoilers as well as the analysis of inner secondary flows and pressure levels that develop in the individual chambers. Unlike the geometry used for the validation case, which had walls of finite thickness to match the rigid structure used in the experiment, the enclosing canopy surface and contained ribs are modelled as sheets of infinitesimal thickness. This entails a number of key benefits but also some compromises for the setup. A key advantage is of course a near-ideal replication of the thin fabric of which the kite consists. Also, changes to the geometry such as the implementation of the intake and modifications with bleed air spoilers are easier to handle. On the downside, sheet surfaces inherently require identical surface meshes on either side which of course has a negative effect on the cell count as the interior is modelled better than would be necessary. This is the chief reason why the rear 40% of the canopy – which is not of much interest for bleed air spoiler use anyhow as found by [17] and described in Section 1.2 – and the outer part of the kite beyond the 6^{th} rib are filled out and subtracted from the fluid-domain.

Description of the computational mesh

The same meshing principles already applied for the validation case and described in more detail in Section 4.1.2 were applied to generate the surface and volume meshes for the simulations treated in this chapter. Figure 6.2 provides a series of impressions of the resultant grid.

[1]While the corresponding dissertation is not yet published at the time of writing this report, the working principles of the model are well documented in [2].

(a) bleed air spoiler and canopy wake

(b) the forward edge of bleed air spoiler
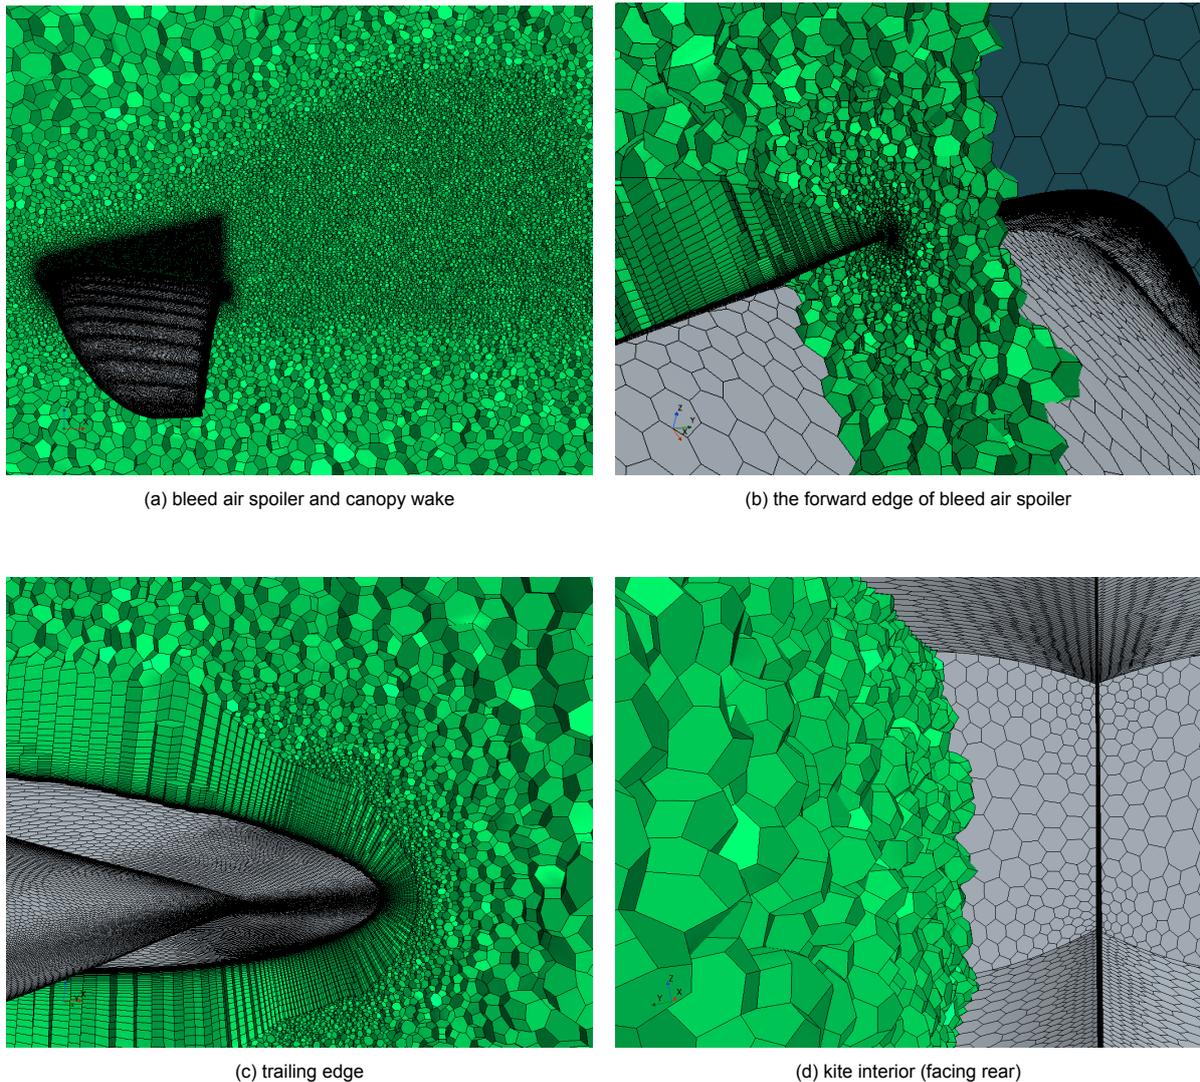

(c) trailing edge

(d) kite interior (facing rear)

Figure 6.2: Impressions of the computational mesh showing cells on a cutting plane that slices through a bleed air spoiler

Figure 6.2a shows cut sections of the imposed volumetric refinement in the bleed air spoiler wake, of which the defined extents are shown clearer in Figure 6.1b, and the main wake which is aligned with the incoming flow and reaches about two-chord lengths behind the canopy. Figure 6.2b gives an impression of a typical rim of the prism layer mesh and how its cell layers gradually transition to the unstructured core mesh before curving around into the also unstructured inner mesh. The scene also again demonstrates the cost with respect to cell count of resolving unbounded prism layer edges with a grid spacing of $y^+ = 1$ ($\approx 4 \cdot 10^{-5}$ m). Figure 6.2c further gives a look at the trailing edge, showing a case where prism layers are retracted in order to maintain maximum aspect ratios of 1.5 in the top layer. Lastly, Figure 6.2d shows parts of the interior of chambers 2 and 3 (counted from the symmetry plane) and illustrates the generally high degree of resolution already mentioned before as well as the usual lack of prism layers.

6.2.2. Environmental conditions

The definition of environmental conditions was nearly identical to those already applied in the validation case. The central difference is the inflow velocity of 10 m/s which roughly reflects the expected relative velocity during the reel-in phase. The ballooned kite geometry described before stems from a simulation of the same velocity. Pivotal flow parameters and kite dimensions are summarised below in Table 6.1.

Table 6.1: Key parameters of the kite simulation environment

Parameter	Definition	Value
Reynolds number	$Re = \frac{U_\infty c}{\nu}$	$5.2 \cdot 10^6$
Dynamic pressure	$q = \frac{1}{2}\rho U_\infty^2$	$61.25\,\text{Pa}$
Projected surface area of kite	S	$94\,\text{m}^2$
max. chord length	c	$7.5\,\text{m}^2$

6.3. Sensitivity Studies

Again, a thorough sensitivity study of the designed simulation environment is critical for the dependability of the subsequent investigations. As for the validation case, effort was put into understanding the limitations and establishing proper control of the spatial and temporal discretisation of the aerodynamic phenomena.

6.3.1. Global mesh sensitivity study on baseline case

Following the approach from the previous chapter, an initial global mesh sensitivity study was carried out using steady RANS computations on the baseline geometry (with an intake but without bleed air spoilers) under standard conditions. After proper time-discretisation settings were identified as described in Section 6.3.2, it was decided that the three meshes and their resulting solutions are also suitable as initialisation points for a mesh convergence study with unsteady calculations. For the sake of better comparability they are therefore discussed here together, but just note that the time-discretisation choices already follow from Section 6.3.2. Figure 6.3 depicts the solution histories by iteration and simulated time for the respective convergence studies.
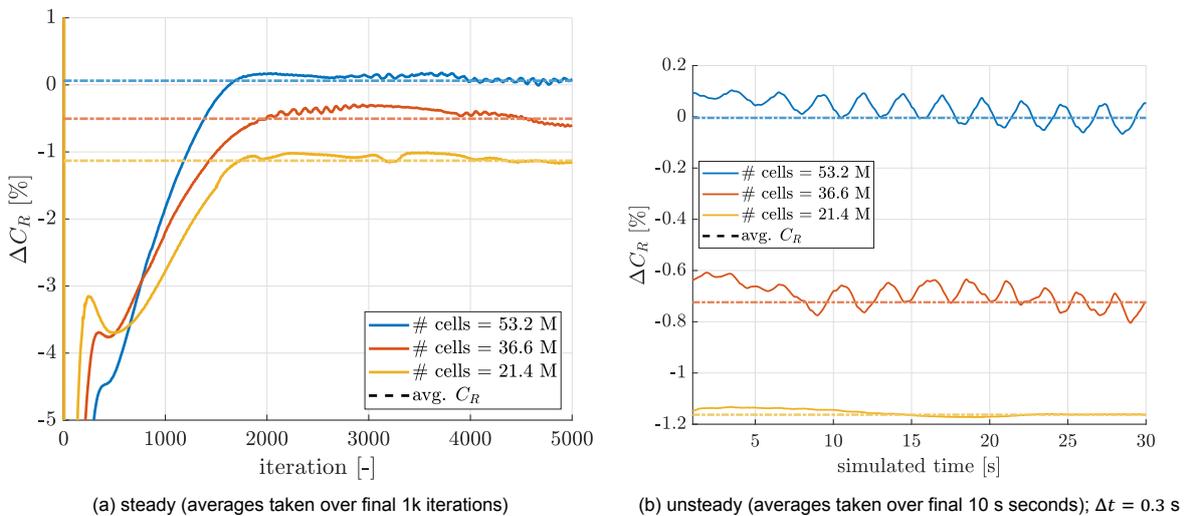

(a) steady (averages taken over final 1k iterations)

(b) unsteady (averages taken over final 10 s seconds); $\Delta t = 0.3$ s

Figure 6.3: Solution histories of the steady and unsteady mesh convergence simulations (ΔC_R denotes the deviation from finest unsteady time-average)

The three steady simulations reached rather well-converged solutions with a low level of remaining nonphysical oscillations that stay within a range of 0.3% to the average results. It is presumed that these oscillations are at least in part due to the intake, where large velocity and pressure gradients are present around the edges that are challenging to time-average, calling for unsteady analyses to reproduce them further. The fact that all three results fall within a range of just over 1% to each other however already suggests that the meshes are once again fairly well resolved.

A look at their unsteady continuations up to a simulated time of 30 s shows that while some unsteady periodic behaviour persists in the simulations, their time-averages are very close (within ~1%) to their steady counterparts, as expected from proper steady RANS solutions. The average solutions are also compared in Figure 6.4 with respect to the characteristic cell sizes.
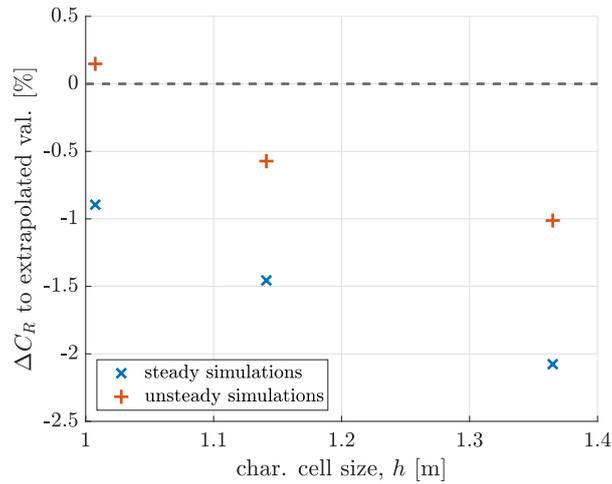
Figure 6.4: Solution convergence with the characteristic cell size, h

The direct comparison above shows that both sets of solutions follow very similar trends, with the unsteady simulations consistently yielding about 1% larger values for C_R than the steady cases. The increase in C_R seems to be roughly inversely proportional to the decrease in h. Table 6.2 provides a summary of the processed grid convergence characteristics.

Table 6.2: Results of initial global mesh convergence study on baseline case; $\alpha = 15°$

	mesh index	1	2	3
	base size [m]	$0.5c$	$1c$	$2c$
	$N_i \left(\cdot 10^6\right)$	53.2	36.6	21.4
	$r_{i+1,i}$	1.13	1.20	–
steady	ΔC_R to finest mesh	–	-0.57%	-1.19%
	$\epsilon_a^{i+1,i}$	0.0057	0.0120	–
	$\epsilon_{ext}^{i+1,i}$	0.0137	0.0193	–
	$\mathrm{GCI}_{fine}^{i+1,i}$	1.73%	2.36%	–
unsteady	ΔC_R to finest mesh	–	-0.72%	-1.16%
	$\epsilon_a^{i+1,i}$	0.0072	0.0117	–
	$\epsilon_{ext}^{i+1,i}$	0.0032	0.0104	–
	$\mathrm{GCI}_{fine}^{i+1,i}$	0.40%	0.33%	–

The results support the previous assumption that the meshes have a high resolution due to which the individual results do not vary much. The steady simulations score a little 'worse' in terms of the grid convergence indices of the two finer meshes of 2.4% and 1.7%, which is however also approximately the order of magnitude of the results contained in the editorial example of the mesh convergence guideline [8]. The computed order comes out to 2.75, which is not exactly the expected 2nd order that was imposed on the simulations but also not very far from it either.

The unsteady simulations perform a little better, by comparison, scoring GCIs of 0.3% and 0.4% with their finest meshes. Note that the finest mesh scores worse than the medium one since the slope between solutions 1 and 2 is higher than that between 2 and 3. Due to the susceptibility of the method for very close results, the estimated order of the scheme of 9.4 is quite far from the expectation. All in all, the convergence study gives no grounds to assume that any of the investigated base sizes produce insufficient meshes. Once again, the medium coarseness with a base size equal to the chord length is chosen for subsequent simulations knowing that there is still some leeway to smaller cell counts.

6.3.2. Establishing an efficient time-discretisation method for BAS simulations

The combined complexity of the individual kite simulation environments in terms of cell count and the required time windows necessary to extract reliable averages together with the aspiration to cover as many different spoiler configurations as possible under the time limits of the thesis called for a very efficient time-discretisation method.

Proper choice of Δt in overcoming the transient phase
In gaining experience with the new simulation environment, it became clear that there are substantial transient effects to be overcome at the start of all unsteady simulations. Unlike in the validation case, where steady RANS computations generally proved to be good initialisation points for unsteady simulations on the same geometry, canopy configurations with bleed air spoilers were prone to diverge after about 500 iterations. After 100 to 200 iterations, which was generally sufficient to reduce the residuals to an order of magnitude of 10^{-5}, they were therefore switched to more stable, unsteady computations. At this point, however, the flow field is still rather arbitrary. As a result, the unsteady solution takes a substantial time until constant time averages can be identified.

Fortunately, it was discovered that the simulations continued to produce stable results even at relatively large time-steps. Figure 6.5 depicts the resultant time-histories of simulations using different Δt.
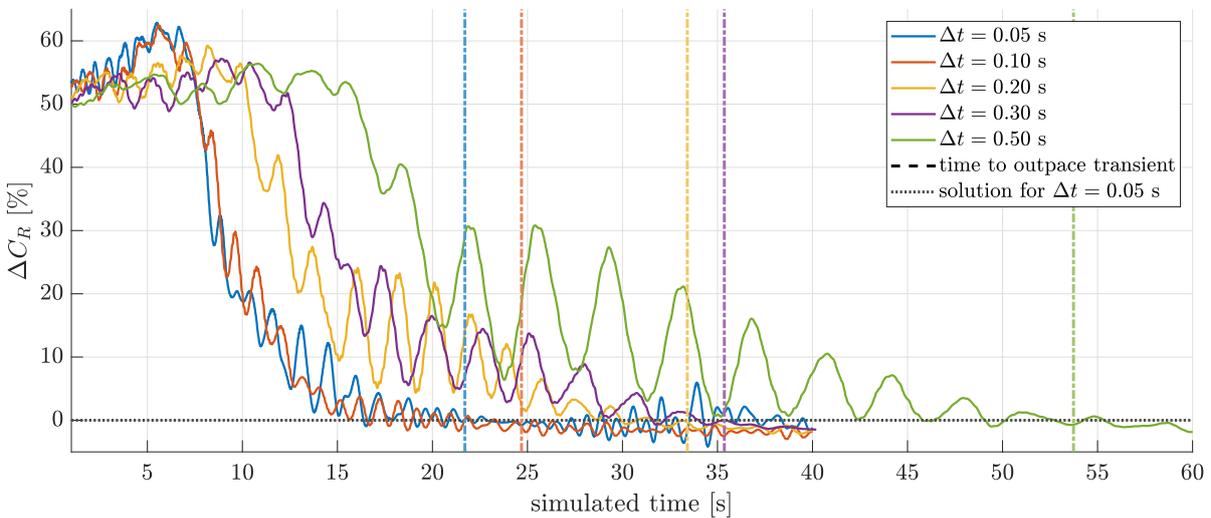

Figure 6.5: Solution histories under various Δt; $\alpha = 15°$; standard bleed air spoiler

The plot above allows for a number of valuable insights. A first, soothing observation is the fact that all simulations eventually plateau and oscillate about the same time-averaged value of C_R, as also supported by Figure 6.6 which shows that all final time-averages are within a range of 2% to each other. Notice also that larger Δt go along with smaller oscillations about their respective means after having outpaced the transient, which is expected as unsteadiness is smeared out over a given time-step (see Section 2.2.3). In this sense, a strong degree of unsteadiness seems to persist in the $\Delta t = 0.05\,\text{s}$ simulation, for instance, whereas oscillations are almost reduced to 0% under $\Delta t = 0.5\,\text{s}$. Rather than having to time-accurately resolve the nonphysical unsteadiness at the start of the simulations with small time-steps (which would require a Δt of 0.01 s as explained in the following section), the results show that it is possible to outpace the transient phase using time-steps that are up to 50 times larger.

To identify the time-step at which the simulation is deemed to have outpaced its transient, moving top-hat time-averages with a 5 s-window of C_R are calculated, and the time-step at which this moving average first falls within 2% of the final $\overline{C_R}$ is determined as the point of convergence. Evidently, simulations under large Δt require a longer time window to reach this point of convergence, however, the relationship between the two is clearly not linear and thus an optimum choice of Δt is expected that entails the smallest number of necessary time-steps, and thus computational power, to outpace the transient. For easier analysis, the required time-steps with respect to Δt are plotted below.
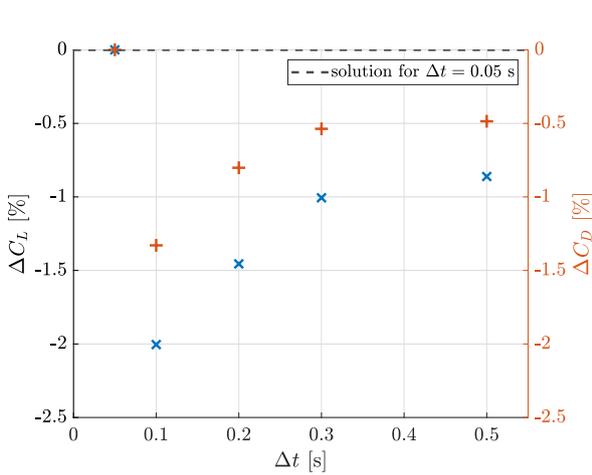
Figure 6.6: time-averaged solutions after the simulation is said to have outpaced the transient
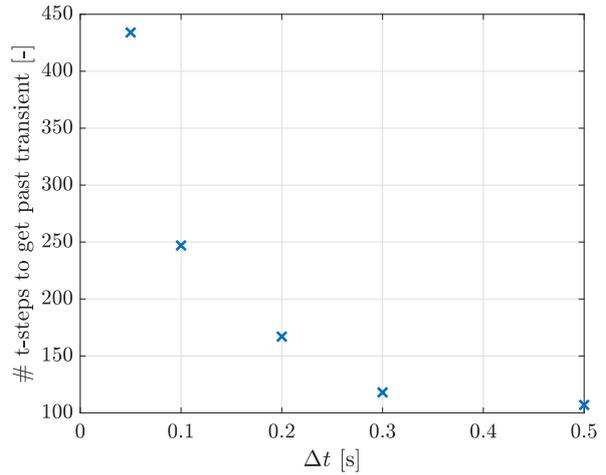

Figure 6.7: Time-steps required to outpace transient

Figure 6.7 shows that there is indeed a rather systematic dependence of the number of time-steps on the chosen Δt. However, while the time-step increment from 0.05 s to 0.1 s more than halved the required computational expense to outpace the transient, there is clearly a diminishing return as the next doubling already only reduces the number of required time-steps by an additional 30%, and barely any benefit is obtained by the last increment from 0.3 s to 0.5 s. It is finally concluded that a $\Delta t = 0.3$ s is well suited for this part of the simulations.

Δt for time-accurate reproduction of unsteady phenomena

The choice of Δt for the time-accurate reproduction of unsteady phenomena caused by the bleed air spoilers was again made upon careful consideration of convective Courant number scenes of several important flow regions. Eventually, $\Delta t = 0.01$ s was identified as a suitable candidate, of which a corresponding Co-scene is depicted in Figure 6.8 which shows a chordwise plane section that slices through the bleed air spoiler.
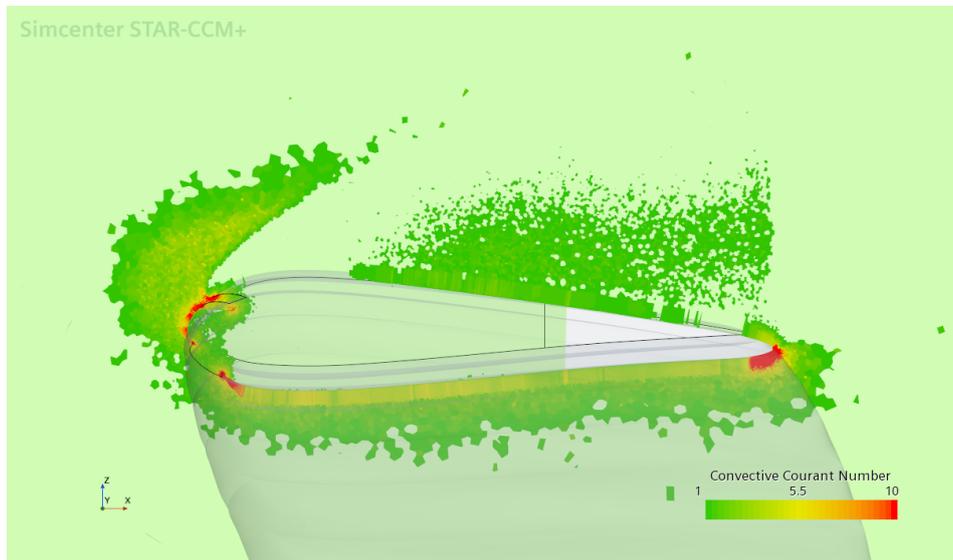

Figure 6.8: Convective Courant Number scene; $\Delta t = 0.01$ s

As for the validation case, it is tolerated that Cfl numbers much larger than 1 are reached in some parts of the flow. As the scene above illustrates, this applies mostly to the boundary layer in front of the bleed air spoiler, the boundary layer on the lower side of the canopy and the lower trailing edge. However, these are areas where quite steady flow behaviour is expected and thus a smearing effect

should not have a significant impact on the final solution. Important unsteady behaviour in the wake of the bleed air spoiler, where velocities are generally lower than in the surrounding flow, a relatively accurate reproduction of unsteadiness is expected.

6.3.3. Investigation of increased BAS wake resolution

Since it proved to be beneficial to apply additional refinement to the bleed air spoilers immediate wake in the validation phase, a similar wake region was defined for the kite simulations. A representative example of this volume is pictured in Figure 6.1b.

With the experience from the validation case and a set of three simulations of varying refinement of the BAS wake (see Figure 6.9, it soon became clear that very extensive simulations would be required to reliably identify differences between their mean results. However, the required time investment of approximately two days with 60 CPUs for a given 60 s simulation is quite large. Meanwhile, at baseline cell counts for bleed air spoiler simulations of about 50 M, a further resolution of the BAS wake to 0.5% (finer than the validation case) only inflicts a 3% increase in the number of cells. It was therefore decided to simply impose the finest investigated resolution of 0.005 on the following simulations. The state of the simulations at this point, depicted in Figure 6.9, gave no grounds to assume that further refinement would lead to substantial differences in the solution after all resolutions showed similar trends.
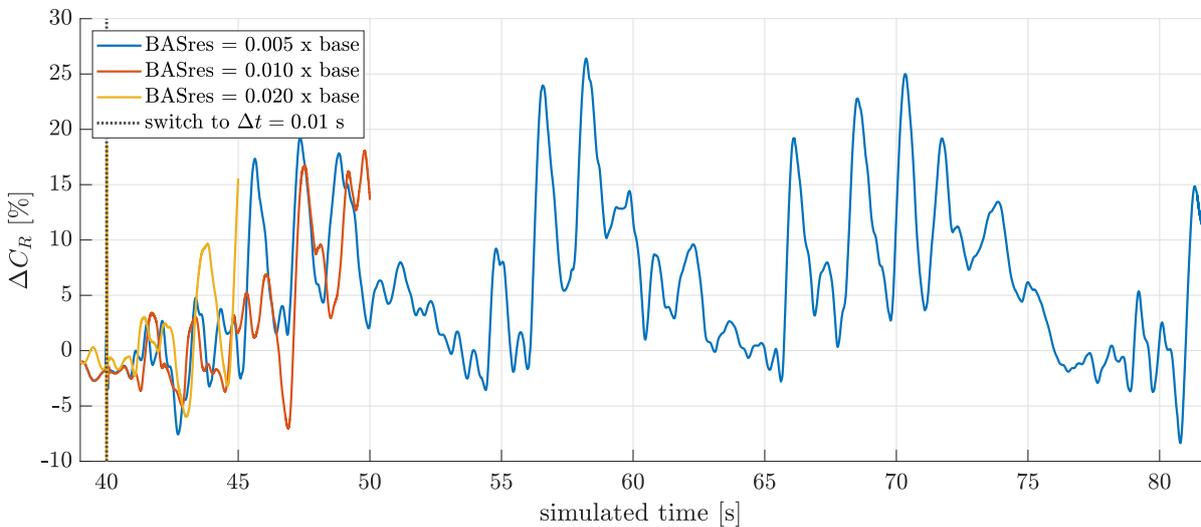

Figure 6.9: Time-history of status quo BAS simulation; $\alpha = 15°$; $\Delta t = 0.01$ s after 40 s

6.4. Evaluation of Results

The sensitivity studies discussed above proved to be valuable both for establishing proper spatial and temporal discretisation settings as well as in identifying and evading potential pitfalls, such as initialisation issues. With this backing, a series of CFD investigations was launched to gain insights into the performance of the kite under the influence of various bleed air spoiler configurations. In order to establish a proper benchmark to compare the performance characteristics to, the baseline case without bleed air spoilers was evaluated first. Then, thorough parameter studies on the spoiler sizing and the chordwise positioning were carried out. Key performance indicators that were used to compare the different configurations are the time-averaged C_R values, the glide ratio, the pitching moment around the control pod, and average ram-air pressure across the six resolved chambers of the kite and predicted changes to the average power output of the system.

6.4.1. Comparison of baseline results to panel code

For the evaluation of the baseline configuration, the kite's performance is analysed under the same inflow conditions that were also imposed on the panel-code based FSI simulation from which the ballooned geometry was obtained. The resultant, time-averaged flow field is depicted in Figure 6.10.
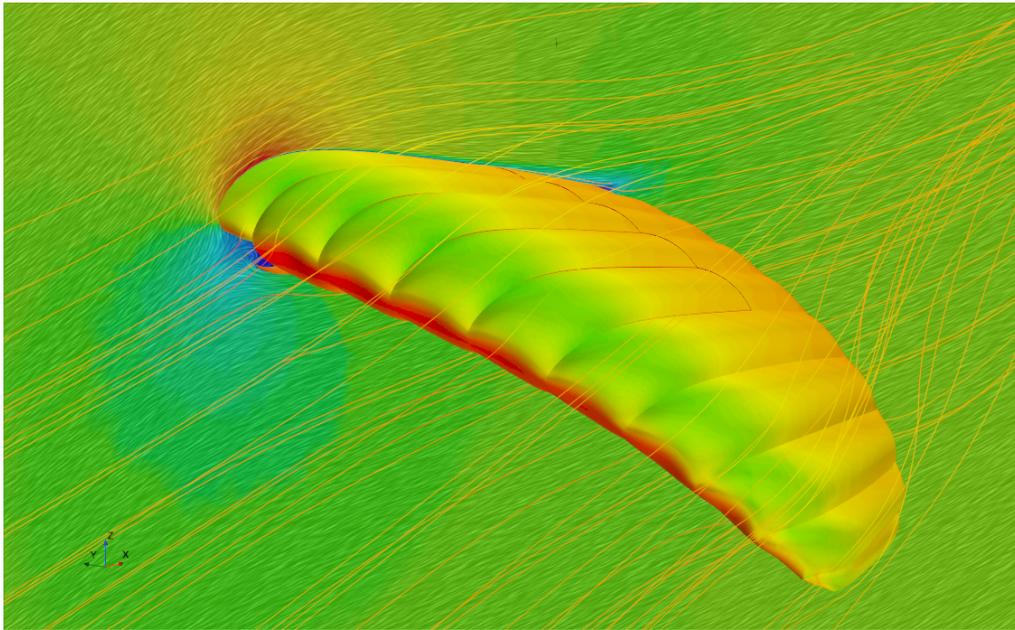
Figure 6.10: SkySails power kite snapshot of steady RANS simulation (velocity magnitude on plane section; C_P field on canopy surface); $\alpha = 15°$

The scene above nicely illustrates the typical flow field that develops around the kite, including a glimpse at the stagnated velocities inside the canopy, the ram pressure on the leading edge (i.e. $C_P = 1$) and the acceleration of the surrounding flow as it follows the profile of the kite. Note also that little to no flow separation is present on the cutting plane which slices through chamber 2, indicating good operating conditions for the canopy.

While no adequate experimental data is available to compare the results to, the aerodynamic forces which were computed with the panel code simulation by Paul Thedens allow for a quick double-check. The deviations between the CFD and panel code findings are summarised in Table 6.3. Note that in the FSI simulation the canopy pitches freely around the control pod, and thus the moment around the pod must be equal to zero for a converged solution (see Figure 6.11 for an illustration). When the resultant structure was adapted for the CFD simulation, the coordinate system was left untouched and thus ideally the moment around the pod should still be equal to zero. The angle of 15° thus reflects the so-called 'trim angle of attack', α_{trim}, of the kite configuration. Furthermore, it is important to mention that the panel code also approximates the lift and drag contributions of the bridal system. Its influence was removed from the lift and drag values before comparing them to the CFD results, but they also caused a positive moment (upward pitching) around the pod which entails some degree of expected mismatch in the moments.
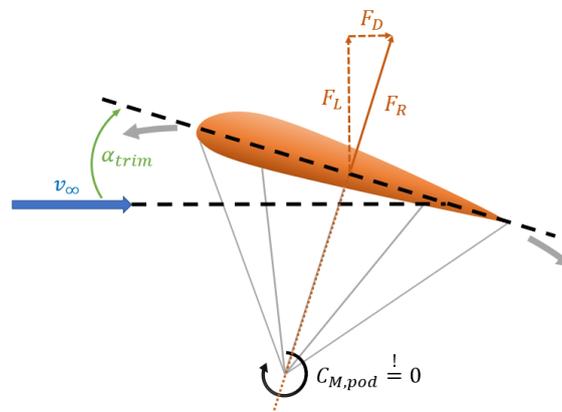

Figure 6.11: Illustration of $C_{M,pod}$ and its importance for the trim angle of attack, α_{trim}

Table 6.3: Percentage deviations between CFD and in-house panel-code results, $\alpha = 15°$

Parameter	Δ to panel code
C_L	-2.78%
C_D	+0.51%
C_R	-2.73%
$C_{M,\text{pod}}$	-0.04
L/D	-3.27%

Clearly, the CFD simulation yields very similar results to the panel-code computations. The predicted lift is slightly lower and the drag slightly higher. Because the panel code is based on pure potential flow analysis and does not include models to take into account separation or viscous boundary layer effects, it is no surprise that the CFD simulation predicts a slightly worse, but likely more realistic performance. This also reflects in the glide ratio, which is also slightly lower. Note that the CFD solver predicts a slight downward pitching moment around the pod, which is likely in part due to the missing influence of the bridal system that was explained before.

6.4.2. Evaluation of bleed air spoiler sizing

For a first investigation of the application of bleed air spoilers on the kite, a parameter study is carried out with respect to the chordwise sizing of the outlets. The nominal basis for the bleed air spoilers (size: 100%) matches the dimensions and position of the current BAS that is built implemented in the SK123. As shown in Figure 6.12a, it was designed with an orientation angle of 45° under the expectation that a BAS right under the suction point would have the ideal effect (an assumption that will be examined in Section 6.4.3), and an angled outlet would be most likely to eventually cross this point along its span. By means of this size-focused investigation, a number of questions shall be answered, including the following:

1. What is the impact of the current bleed air spoiler design on the performance of the kite?

2. Can a similar effect be achieved with smaller outlets?

3. How do bleed air spoilers affect the trim angle of attack, α_{trim}, of the canopy?

4. What are the limits to the bleed air spoiler opening in terms of available ram-air pressure?

To this end, three bleed air spoiler configurations were generated, which are depicted in Figure 6.12.
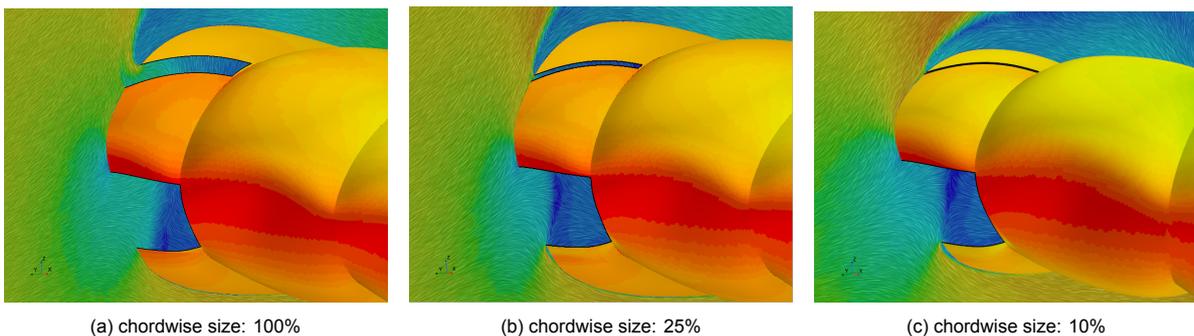

(a) chordwise size: 100%   (b) chordwise size: 25%   (c) chordwise size: 10%

Figure 6.12: Close-ups of intake and BAS openings; velocity magnitudes depicted on cutting plane, C_P distribution on canopy surface; $\alpha = 15°$

Simulations were first run for all configurations under the trim angle of attack of the baseline case (15°). As expected, the drag of the canopy and hence the resultant moment about the control pod increased, indicating that the canopy would rotate backwards around it under unrestricted conditions (as demonstrated in Figure 6.11). In an effort to find the new trim angle of attack for the respective spoiler cases, additional simulations were run at higher angles of attack. Once a negative slope could

be identified, the anticipated α at which C_M should be zero was extrapolated and another corresponding simulation was executed. The simulations were run for a simulated time ranging between 15 to 32 s under time-accurate time-stepping (after Δt is switched to 0.01 s), depending on how long the solution for C_R took to reveal periodic oscillations around a sufficiently constant mean. The time-histories of C_R are not further discussed here but have been included in Appendix A.1 for further reference. The appendix also holds snapshots[2] of the final flow fields for all cases. A total of at least four simulations were run in this way for each spoiler size.

The final simulations were further processed by time-averaging the results over the final 5 s of the respective solution times. The resulting aerodynamic force and moment polars are shown in Figures 6.13 to 6.15. The included error bars indicate the respective maximum and minimum value that was contained in the time-averaging window. All figures include an additional plot showing the performance of the spoiler configurations over their size under the simulation for α predicted to be closest to the respective trim state. These trim states have also been emphasised in the other plots with circles.

(a) $C_{M,pod}$ polars for the evaluated spoiler sizes (trim states are circled)

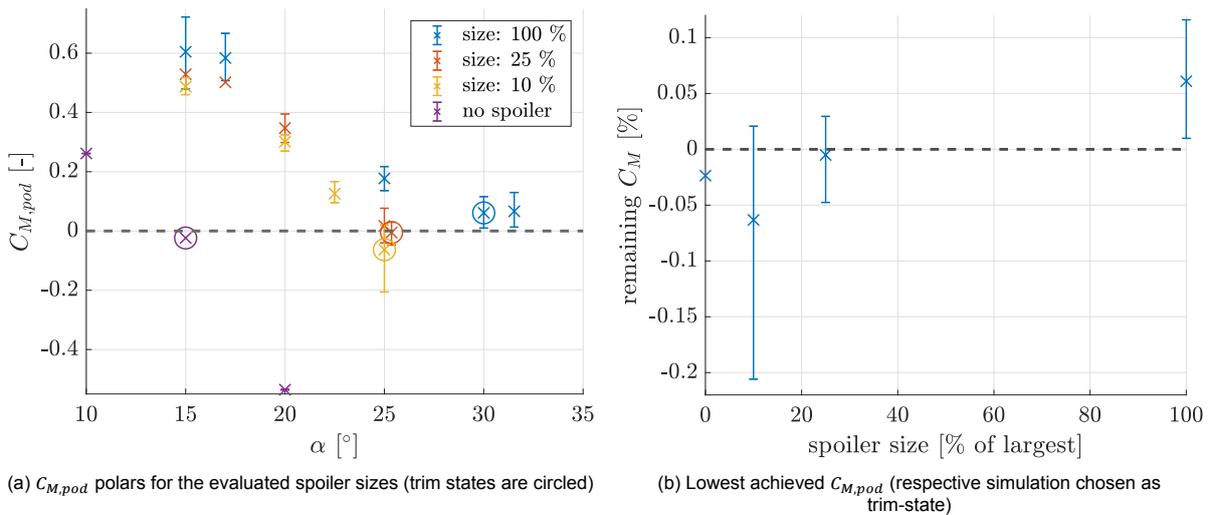(b) Lowest achieved $C_{M,pod}$ (respective simulation chosen as trim-state)

Figure 6.13: pitch moment around control pod, $C_{M,pod}$

The C_M polars again clearly demonstrate that the spoilers cause a strong backwards pitching moment on the kite. Note that this moment is strongest for a given angle of attack under the largest spoiler. Similarly, the 25% cases also consistently predict higher moments (ΔC_M of 0.04 to 0.06) than the 10% cases. The fact that the C_M polars have negative trends is an important feature to note, as it shows that the kite remains stable (i.e. if the moment increased even further, α would diverge). Trim angles of attack are eventually identified for all spoiler configurations in a range of 25° to 30°. This already touches on one of the most important findings of this parameter study: Clearly, smaller changes in the trim angle of attack can be achieved with smaller spoiler sizes (as also later highlighted by Figure 6.17).

It is acknowledged that another simulation should have ideally been carried out for the 10% case at 24.2°, which would likely be closer to the true α_{trim} by interpolation. Likewise, it was only possible to match α_{trim} for the larger spoilers by $\Delta C_{M,pod} \approx 0.06$, as highlighted by Figure 6.13b. Further simulations were omitted due to computational resource constraints. It is therefore important to keep in mind that when speaking of a 'trim state', what is really meant is the result of the simulation which came closest to it. For the purpose of this evaluation, all simulations are considered to be sufficiently in range of $C_{M,pod} = 0$ to allow for this assumption. After all, as indicated by the error bars, all simulations at least scratch on this state, and due to the unsteadiness of the flow, no perfectly trimmed condition could be found with the chosen simulation approach anyhow. The only truly appropriate remedy to take the free movement of the kite into account would be a fluid-structure coupled simulation. In fact, this would not even necessarily have to involve a flexible canopy, but at least enable free movement of the rigid kite configuration around the control pod. While this would naturally increase the computational expense required for a given iteration, it would save the trouble of having to run a number of

---

[2]While time-averaged flow fields would arguably be more informative than snapshots, the required buffering of previous solutions would have multiplied the already large file size of a given simulation and was therefore omitted.

independent simulations back to back. Within the time frame allotted to this thesis, the development of a fluid-structure coupled simulation environment would however have been beyond the scope of the realistically feasible, and hence it is only mentioned as a possible improvement for subsequent studies.

(a) ΔC_L polars for the evaluated spoiler sizes

(b) ΔC_D polars for the evaluated spoiler sizes

(c) ΔC_R polars for the evaluated spoiler sizes
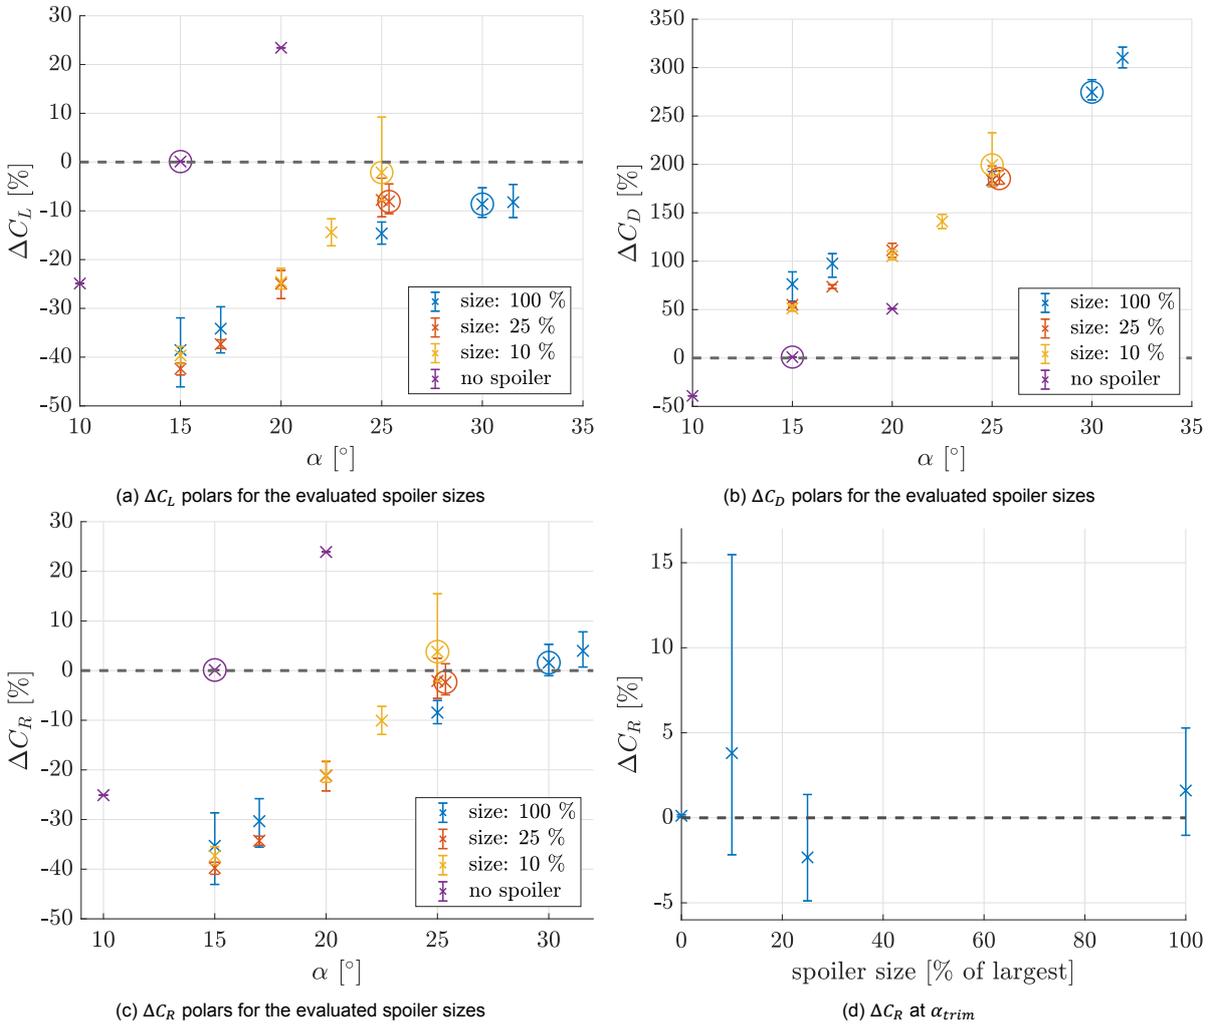
(d) ΔC_R at α_{trim}

Figure 6.14: ΔC_R with respect to baseline case (trim states are circled)

A look at the resulting C_L polars shows that all configurations with spoilers had a remarkably similar impact on the lift acting on the canopy for a given smaller (untrimmed) angle of attack in relation to the baseline case. In fact, the results show that the largest evaluated spoiler did not achieve as much of a decrease in C_L at $\alpha = 15°$ and $17°$ as the smaller two. It also appears that under the influence of spoilers, C_L still increases more or less linearly with α. Meanwhile, it appears that the slope of $C_L(\alpha)$ is more positive for smaller spoilers than for larger ones, which manages to retain smaller values of C_L even under large α. The unsteadiness of the flow also clearly shows in the oscillations of the lift forces, as indicated by the error bars, which were basically non-existent under the baseline case. Generally, it seems to oscillate by around 5 to 8% in either direction, with the exception of the $\alpha = 25°$ degree result for the 10% case where the time-averaging window clearly captured a larger variation. The trim lift is largest for the smallest bleed air spoiler (almost identical to the baseline) and approximately equal with the 25% and 100% sized ones.

Similar linear trends are observed with respect to the C_D polars in Figure 6.14b. The data show that the largest spoiler, which naturally also causes the biggest disturbance to the flow (see Figure A.2 for reference), causes the aerodynamic drag to increase the most. At its $\alpha_{trim} = 30°$, it is even 275% larger than that of the baseline. Meanwhile, the results for the two smaller spoilers almost coincide, showing a drag increase of approximately 200% each.

Jointly evaluated in the form of C_R, the linear trends persist, as expected. Note that the deviations in C_R almost match that in C_L which, as it is larger, has a much stronger impact on $\sqrt{C_L^2 + C_D^2}$ than C_D. Had the trim angle of attack not changed, then the spoilers would have clearly caused a very significant decrease in the resultant aerodynamic loading on the kite. However, the results show that due to the increased α_{trim}, the actual value of C_R is once again very close to the baseline. As a matter of fact, both the 10% and 100% cases finally yield loadings that exceed that of the baseline, as highlighted by Figure 6.14d, even if only by 4%. It, therefore, appears that the investigated configurations at least cannot achieve the desired decrease in C_R. Rather, the decrease in lift and the decrease in drag have virtually compensated each other in this regard.

(a) ΔE polars for the evaluated spoiler sizes (trim states are circled)
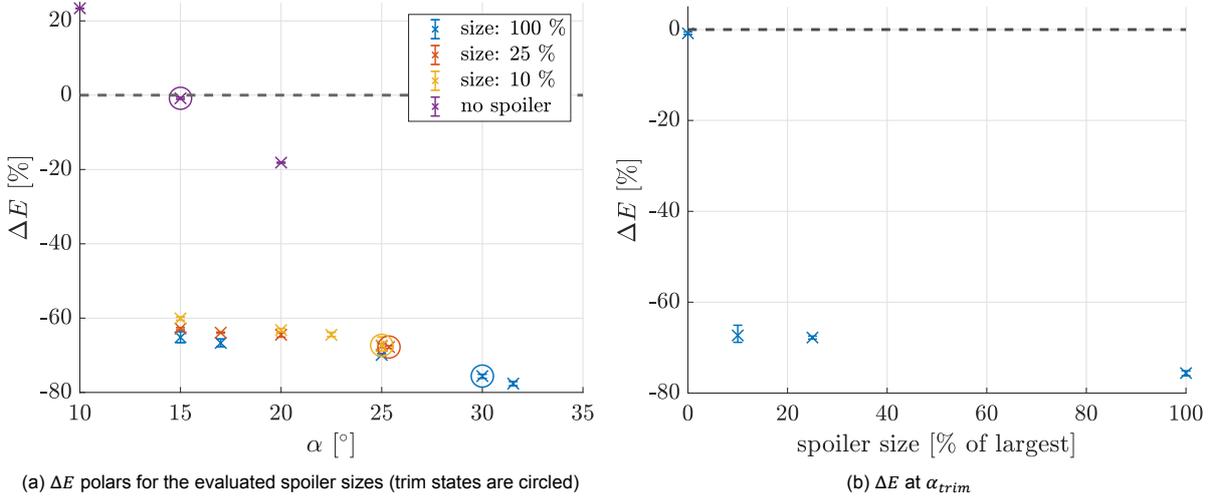
(b) ΔE at α_{trim}

Figure 6.15: ΔE with respect to the baseline case

Fortunately, this is not the case with respect to the glide ratio, depicted above. The combined general decrease in the lift polars and the increase in the drag polars lead to an immense drop in E, which amounts to -60 to -80%. Note that for a given angle of attack, the size of the spoiler has an almost insignificant impact on the value of E. Only upon closer review does one notice that there is a small but consistent added benefit associated with a larger spoiler at a given α. With respect to α_{trim}, the difference is more noticeable, but also here does the smallest spoiler still produce just 8% less than the decrease of 75% caused by the configuration currently implemented in SkySails' kites.

It is a known problem in the company that the current bleed air spoilers tend to deflate the canopy too much, causing it to lose a significant degree of its stiffness. Fluid-structure-coupled simulations using an in-house panel-code tool-chain showed that the kite is prone to collapse at internal pressures of $C_p < 0.9$. In order to keep track of the 'cost' in terms of the internal pressure drop, pressure sensors where thus defined towards the rear of the individual chambers. The resultant average pressures that were recorded over the course of the simulations are plotted in Figure 6.16.

(a) avg. chamber C_p polars for the evaluated spoiler sizes (trim states are circled)

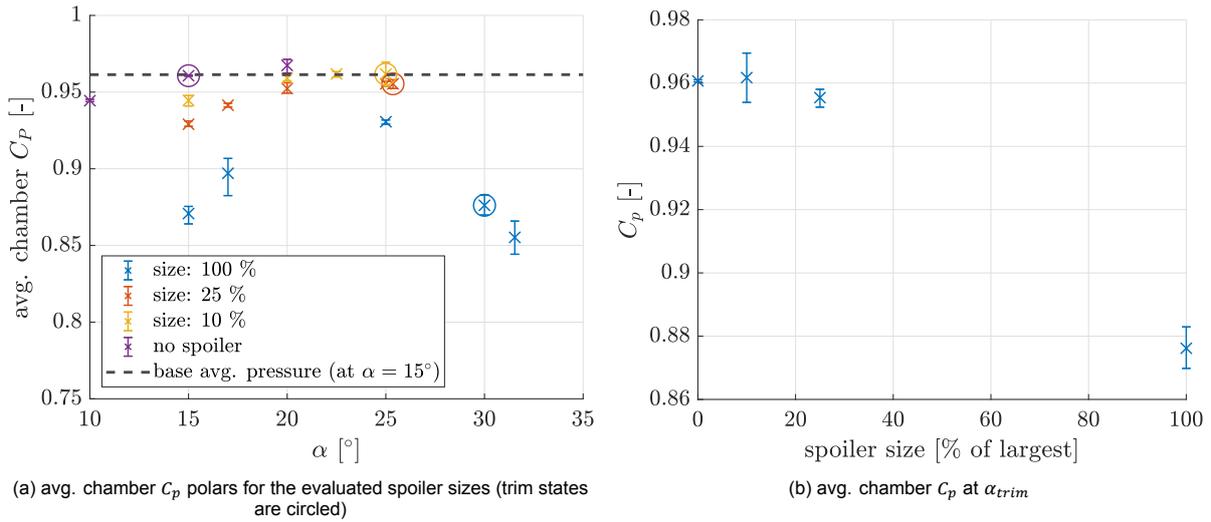(b) avg. chamber C_p at α_{trim}

Figure 6.16: average chamber pressure coefficients

The data allow for a number of important observations. Interestingly, some of the lowest pressures are recorded at $\alpha = 15°$. This seems counter-intuitive, as one would expect the canopy's intake to perform best at its design angle of attack. However, a second look at Figure 6.12 demonstrates that the intakes are actually located a bit too low as they only to cover a part of the stagnation point pressure at 15°. When the kite's angle of attack increases, the stagnation point will move downwards, putting the intake in a better position to capture ram-air pressure, which presumably explains the overall positive trend for small values of α. At 30°, the stagnation point has moved past the intake once again and pressure levels are lower accordingly.

It is important to note that the dimensions and location of the intake are also just design choice that can always be re-dimensioned appropriately. Therefore, a direct comparison between the spoiler sizes for given angles of attack is more instructive. The plot shows that for any given α, the smaller bleed air spoiler requires a significant smaller pressure drop while achieving a comparable performance in terms of a change in the glide ratio, as discussed before. In fact, none of its recorded pressures fall into the risk zone of $C_p < 0.9$ at which diminishing stiffness is expected, and at the trim-angle of attack pressure levels barely differ from the baseline configuration. Of course, this is also in part because of a non-optimal intake for $\alpha = 15°$. Although none of the configurations have to be ruled out because of these results, they still show that smaller bleed air spoilers may be of advantage when one does not want to design intakes to cover an excessive range of attacks.

Using the tool for power cycle evaluations developed in the context of chapter 3, the predicted average cycle power of the system was computed for all sets of C_R and E that were found in this study.
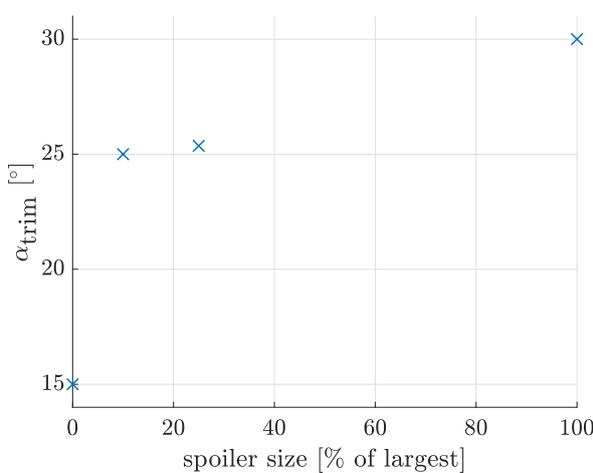

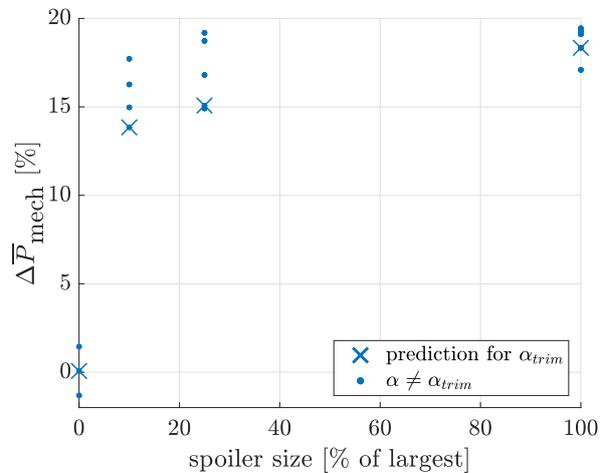Figure 6.17: trim angle of attack over spoiler size

Figure 6.18: max. predicted $\Delta \overline{P}_{\text{mech}}$ over spoiler size

The evaluations show that that the three examined spoiler configurations could lead to an increase of the average cycle power of 14 to 18%, depending on the spoiler size. As noted before, this is mostly due to a reduction of the glide slope. Reviewing our conclusions from Section 3.5, decreases in E are generally preferable to equal changes in C_R, and hence a notable improvement of the system's average power output could still be achieved despite the rather invariant resultant loading. Note also that if the angle of attack were fixed at 15°, which is associated with the topmost data points in Figure 6.18, further improvements to a maximum of $\Delta \overline{P}_{\text{mech}} = 19\%$ would have been possible.

A final account of the key cycle performance characteristics of the individual configurations is summarised below:

Table 6.4: Summary of optimum cycle performances using bleed air spoilers of different size
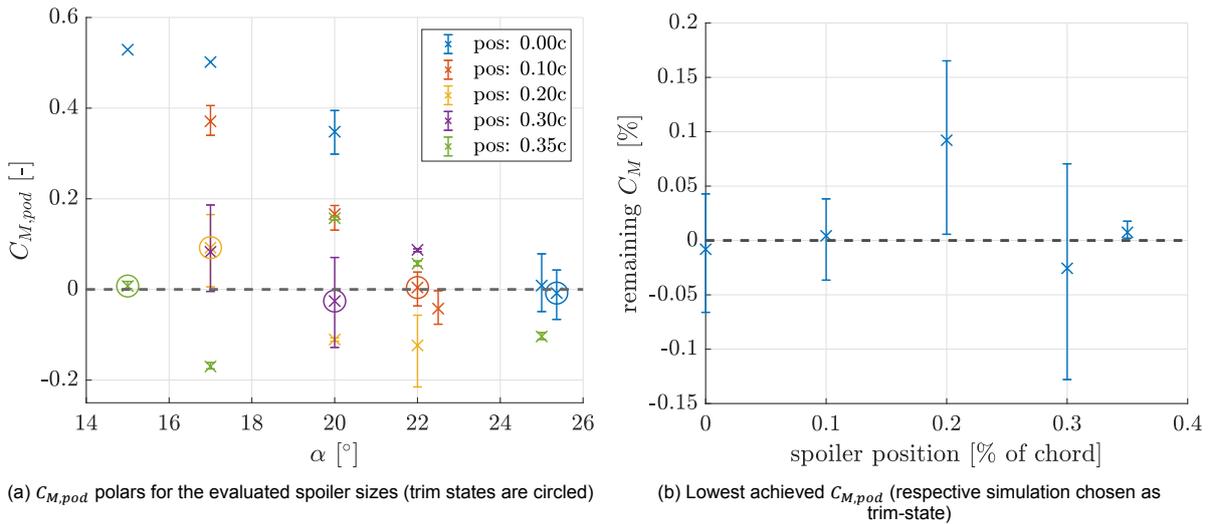
Relative size	$\frac{E_{in}}{E_{base}}$ [-]	$\frac{C_{R,in}}{C_{R,base}}$ [-]	μ_{max} [-]	$\sum\{t_{transfer}, t_{return}\}$ [s]	$\Delta \overline{P}_{\text{mech}}$ [%]
no spoiler	1	1	0.52	194	0
10%	0.33	1.04	0.91 (+73%)	150 (-23%)	+13.8
25%	0.32	0.98	0.92 (+75%)	150 (-23%)	+15.1
100%	0.24	1.02	1.10 (+110%)	140 (-28%)	+18.3

The results further emphasise the strong benefit which the evaluated spoilers could have with respect to the system's cycle performance. The possible reduction in E enabled much larger reel-in speeds, $\dot{l} = \mu v_\infty$, which drastically cuts down on the time required for the transfer and return phase by 23 to 28%. Therefore, although the resultant loading coefficient barely changed, substantial improvements in the average mechanical power output could be achieved.

6.4.3. Evaluation of chordwise bleed air spoiler positioning

It was decided to focus the second parameter study on the chordwise positioning of the bleed air spoiler. The main reason for this is to put the past assumption to the test, that the spoiler is best positioned under the main suction point of the canopy. To this end, the spoilers matching the 25% case from the previous study in size were positioned in a chordwise range of $0c$ to $0.35c$ on the second chamber, where the $0c$ case is identical to the one treated in Section 6.4.2. Based on the observations by Bergeron et al. in [5] (discussed in Section 1.2), positions further to the rear of the kite were not tested. As will become clearer through the analysis of the results, this also proved to be a sound choice.

Simulations with the newly generated geometries were first run for αs of 20° and 22° in order to allow for a first inter-/extrapolation of the trim angle of attack. Between three to five simulations at different angles of attack were eventually carried out for each spoiler position until trim conditions could be identified. Using the same tool-chain for the evaluation as before, the respective solutions were then time-averaged over the final 10 s of the simulated time, and force and moment polars were generated for each case. The results are plotted in Figures 6.20 to 6.21.

(a) $C_{M,pod}$ polars for the evaluated spoiler sizes (trim states are circled)

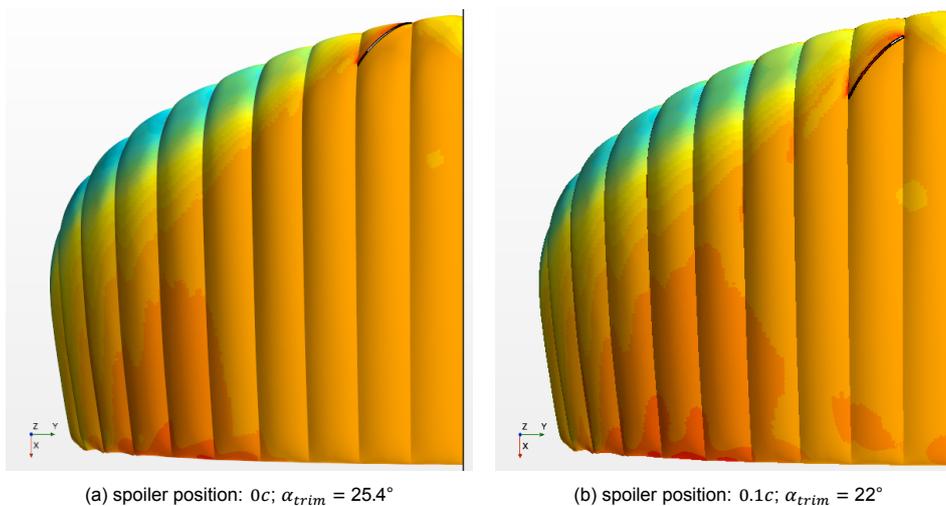(b) Lowest achieved $C_{M,pod}$ (respective simulation chosen as trim-state)

Figure 6.19: pitch moment around control pod, $C_{M,pod}$

The $C_{M,pod}$ polars in Figure 6.19 already allow for a number of interesting insights. Looking at the data points for $0c$, $0.1c$ and $0.2c$, it appears that while the slope of the polars remains approximately the same, the curves shift downwards as the spoiler is positioned further to the rear. This also clearly shows in the angles of attack identified as the respective α_{trim}, which decreases from 25.3° over 22° to 17° in the process. Note that the time-averaged results for $0.2c$ already captured quite a high degree of unsteadiness at 17° and 22°.

The subsequent results for $0.3c$ and $0.35c$ are not as conclusive. No linear trend in $C_{M,pod}$ could be identified for these cases, as values seem to be scattered without an obvious pattern around $C_{M,pod} = 0$. As a consequence, the trim angles of attack which were identified should be treated with caution, as the true result could realistically differ from it by a significant margin. With respect to the $0.35c$-configuration, for instance, it seems that a second zero-crossing could occur at $\sim 22.5°$.

Figure 6.19b further highlights that the average pitch moments associated with the trim angle of attack amount to absolute values of up to 0.1. While this specific value (for $0.2c$) could arguably call for further simulations, the results also demonstrate that there is a lot of unsteadiness present for $0.2c$ and $0.3c$ which makes the cases slightly ambiguous anyhow, and thus the reason for this unsteadiness is thus rather investigated.

Before indulging further in numbers, a series of qualitative observations can be made on the basis of the C_p-surface scenes presented in Figure 6.20 to get a clearer picture of the mechanisms that might lead to the unsteadiness and curious trends of $C_{M,pod}(\alpha)$ observed in Figure 6.19.
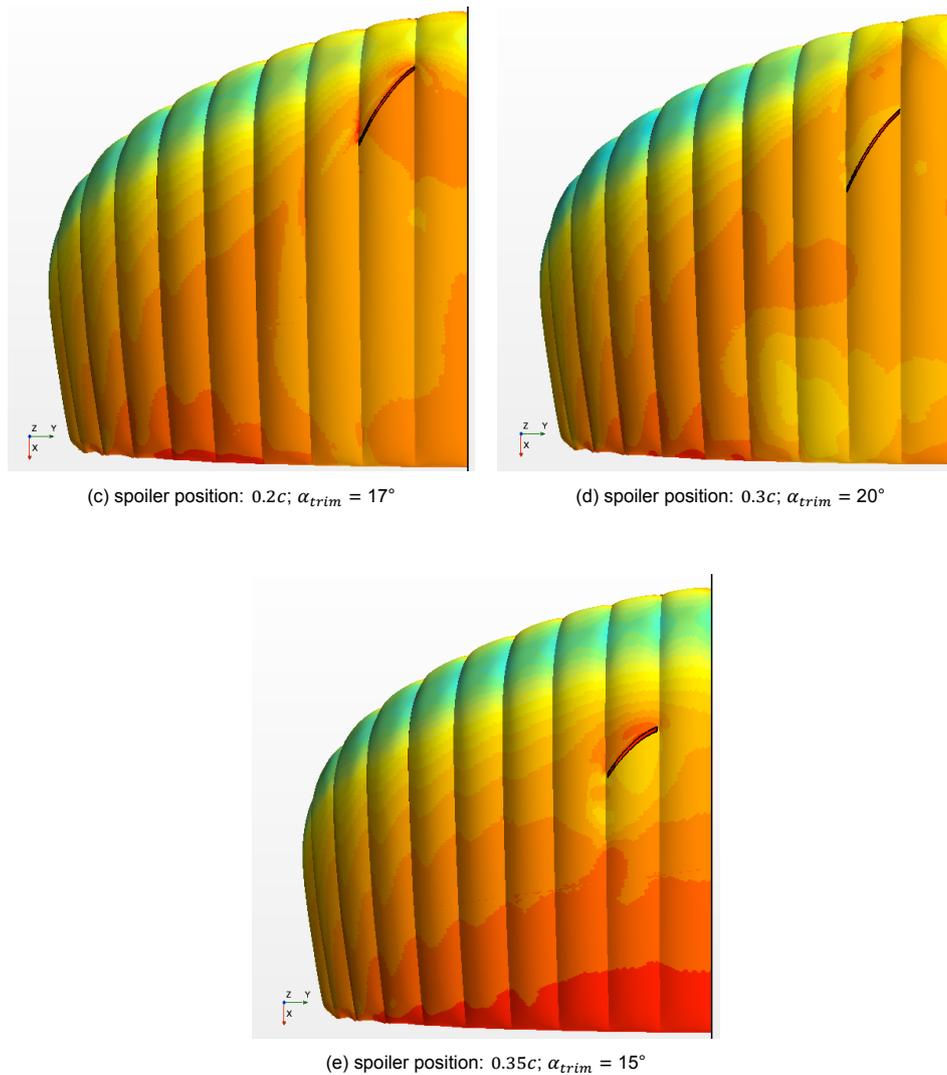

(a) spoiler position: $0c$; $\alpha_{trim} = 25.4°$

(b) spoiler position: $0.1c$; $\alpha_{trim} = 22°$

(c) spoiler position: $0.2c$; $\alpha_{trim} = 17°$

(d) spoiler position: $0.3c$; $\alpha_{trim} = 20°$

(e) spoiler position: $0.35c$; $\alpha_{trim} = 15°$

Figure 6.20: Top view of pressure coefficient, C_p, field on canopy surface for various chordwise spoiler positions

The scenes above nicely demonstrate the effect which the spoilers have on the pressure distribution which develops on the kite's surface. Where a suction peak would usually occur at the front of the kite, the vortices which the spoiler inflicts immediately cause the flow to detach and hence the static pressures in its wake are generally notably higher. This effect is clearly not only limited to the chordwise direction, but rather leads to a visible decrease in suction on the neighbouring chambers as well under all cases except for $0.35c$.

Another key insight is the fact that the spoiler only lies in the region of the suction peak at $0c$ and $0.1c$. Meanwhile, at $0.2c$ it only barely scratches it, and by $0.3c$ it is already to its rear although the vortices induced by it still cause the suction on chambers 1 to 3 to be lower than under the baseline. This observation offers a potential explanation to the trends noted previously with respect to the $C_{M,pod}$ polars. It seems that as long as the spoiler crosses the suction peak for all relevant values of α, a given configuration is stable in the sense that its $C_{M,pod}(\alpha)$ polar maintains a consistently negative slope. When it is not permanently in the reach of the suction peak, the flow will periodically detach and reattach. In a physical, free-moving system, it would not be surprising to see it oscillate back and forth when the spoilers are actuated, as the slits periodically move in and out of the suction peak. Once again, the best way to further investigate this phenomenon would be a fluid-structure coupled CFD simulation, which could probably replicate this unsteadiness better. For the purpose of this investigation, the results for $0.3c$ and $0.35c$ are further evaluated, but only those for $0c$ to $0.2c$ are regarded as credible.

The influence of the spoilers on the aerodynamic forces acting on the kite are depicted below.

(a) ΔC_L polars for the evaluated spoiler sizes

(b) ΔC_D polars for the evaluated spoiler sizes

(c) ΔC_R polars for the evaluated spoiler sizes

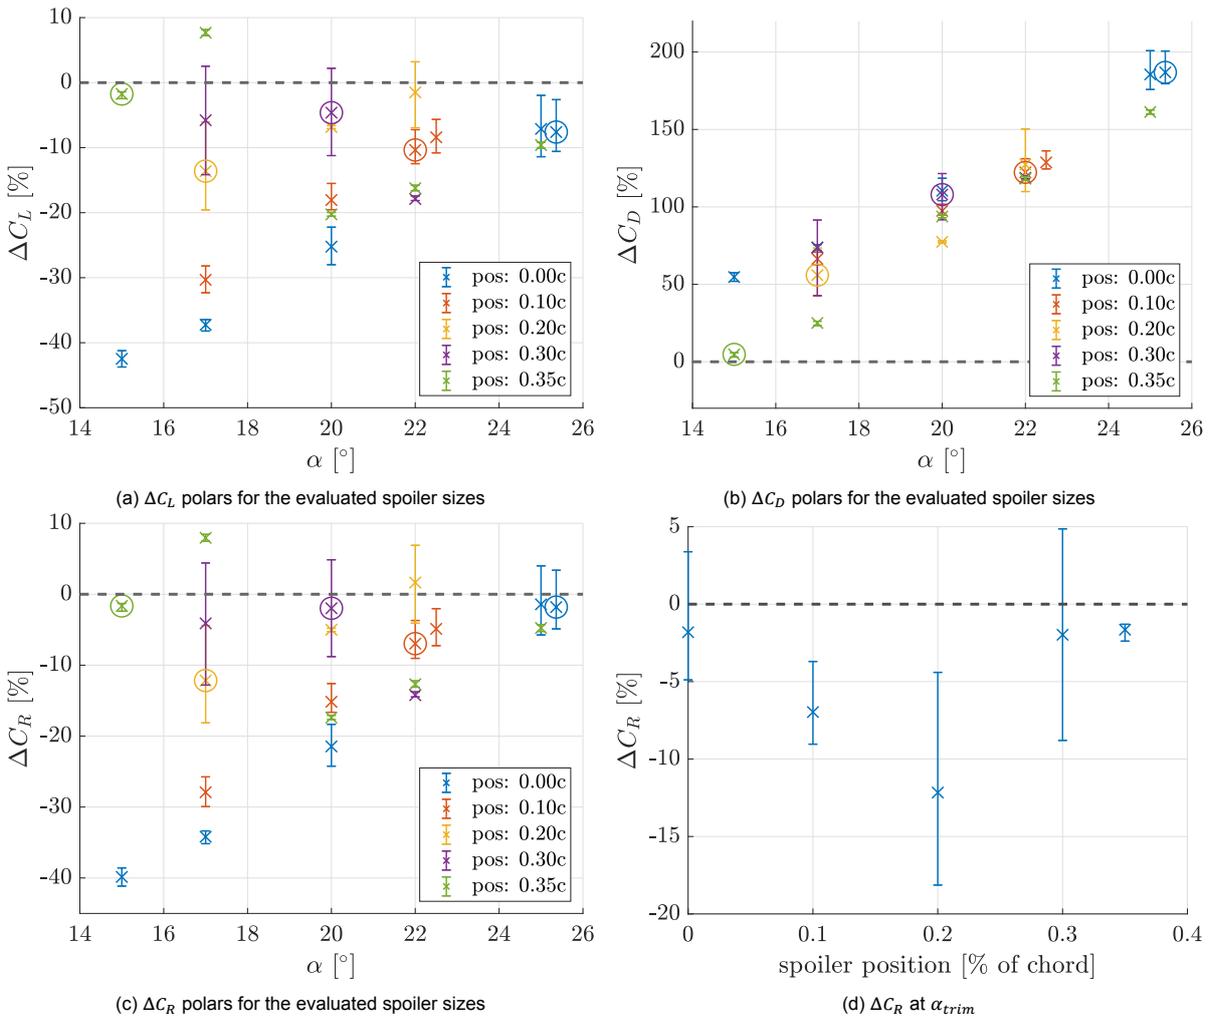(d) ΔC_R at α_{trim}

Figure 6.21: ΔC_R with respect to baseline case (trim states are circled)

The results for C_L and C_D show overall similar trends to those observed in the parameter study on the spoiler sizing. However, a number of distinctions stand out. Whereas the tested spoilers at $0c$ in Section 6.4.2 were able to achieve at most a 9% drop in C_L, the data points for $0c$ to $0.2c$ indicate a somewhat linear dependence of C_L on the chordwise position, with $0.2c$ yielding the lowest lift at -13%. Meanwhile, the aerodynamic drag also appears to decrease linearly as the spoiler is moved aft. These trends are readily explained by the approximately linear dependence of α_{trim} on the chordwise position. A review of Figure 6.21 reveals that in contrast to all configurations from Section 6.4.2 with spoilers located at the leading edge, substantial decreases in C_R of up to 12% are in fact achievable when the slits are moved further to the rear.

The corresponding impacts of the spoilers on the glide slope are shown in Figure 6.22.

(a) ΔE polars for the evaluated spoiler sizes (trim states are circled)
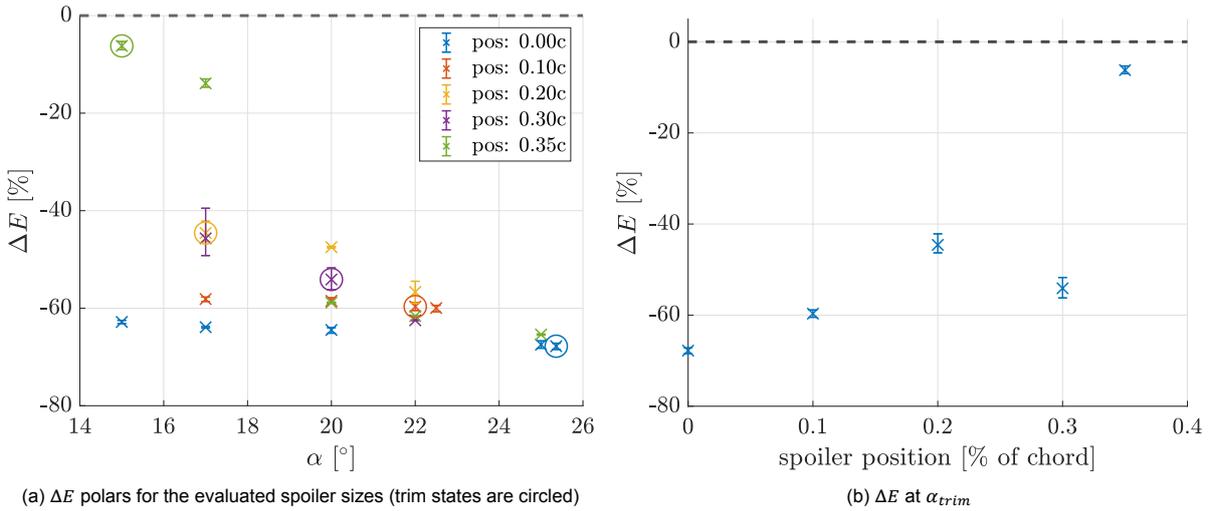
(b) ΔE at α_{trim}

Figure 6.22: ΔE with respect to baseline case

Note that the respective $E(\alpha)$ curves become steeper as the spoiler is moved to the rear. Whereas the $0c$ and $0.1c$ case (and all other cases from Section 6.4.2 for the matter) entail a virtually negligible dependence of E on α, the glide slope polar almost reassumes its baseline behaviour as the spoiler leaves the vicinity of the suction peak. At $0.35c$, the original form is more or less reestablished where $\Delta E_{0.35c}(\alpha = \alpha_{trim})$ returns to values close to 0. This indicates that the spoilers become less effective in decreasing the glide-slope of the canopy, as also demonstrated by Figure 6.22b.

Given the resultant aerodynamic loading coefficient and glide slopes for the five treated cases, the predicted average cycle powers are again computed based on the method described in Chapter 3.
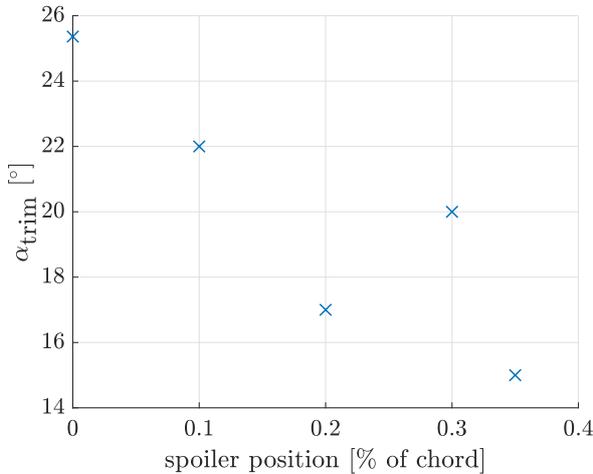

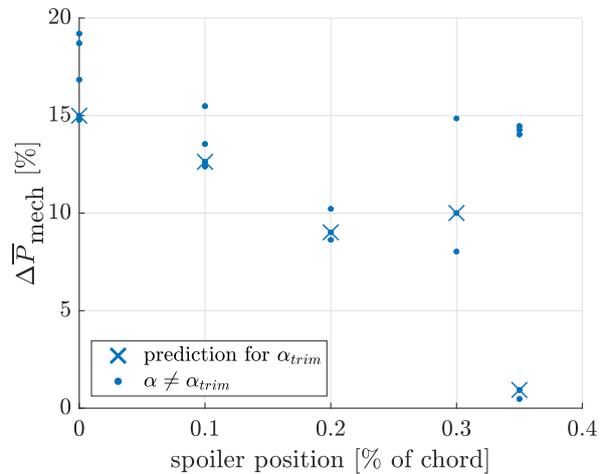Figure 6.23: trim angle of attack over spoiler size

Figure 6.24: max. predicted $\Delta \overline{P}_{\mathrm{mech}}$ over spoiler size

Although lower values of C_R were reached by moving the spoiler further to the rear, the parallelly diminishing influence on the glide slope leads to consistently lower average cycle powers. Although a performance increase of about 9% can still be achieved by a BAS positioned at $0.2c$, the results clearly indicate that there is no inherent reason to move the spoiler away from the leading edge. As such, the assumption laid out by the kite engineering team at SkySails that it is best positioned under the suction point seems to be confirmed.

As before, the performance characteristics corresponding to the maximum average cycle power that can be achieved under the trim-conditions of the evaluated configurations are summarised below:

Table 6.5: Summary of optimum cycle performances using bleed air spoilers of varying chordwise location

Chordwise position	$\frac{E_{in}}{E_{base}}$ [-]	$\frac{C_{R,in}}{C_{R,base}}$ [-]	μ_{max} [-]	$\Sigma\{t_{transfer}, t_{return}\}$ [s]	$\Delta\overline{P}_{\mathrm{mech}}$ [%]
$0c$	0.32	0.98	0.92 (+75%)	150 (-23%)	+15.1
$0.1c$	0.40	0.93	0.80 (+52%)	158 (-18%)	+12.6
$0.2c$	0.55	0.88	0.67 (+29%)	171 (-12%)	+9.0
$0.3c$	0.46	0.98	0.74 (+42%)	163 (-16%)	+10.0
$0.35c$	0.94	0.98	0.54 (+2%)	192 (-1%)	+0.9

The data align with the observations already made with respect to the plots. As such, they show that as the chordwise position is increased, E_{in} generally increases accordingly and C_R decreases (for $0c$ to $0.2c$). As a result of the diminishing effect on the glide slope, the kite must be reeled in at lower speeds which increases the time required for the transfer and return phase and thus lowers the average power output of the system.

7

Conclusions

This master thesis discusses the investigation of bleed air spoilers in their potential for improving the average cycle performance of ground-generator based kite power systems. For this purpose, the characteristic power cycle was first evaluated in a general sense to determine how the aerodynamic behaviour of the canopy should ideally change during the transfer and return phase to achieve an overall benefit. Through a number of sensitivity studies and thorough validation on the basis of experimental results by an independent comparative study, a tool-chain was developed to investigate arbitrary bleed air spoiler configurations by means of computational fluid dynamics simulations. Based on a standard power kite geometry employed by SkySails Power GmbH, two parameter studies were finally carried out to understand how such spoilers are best used to reach the desired aerodynamic effects.

The investigations presented in Chapter 3 revealed the importance of the kite's aerodynamic characteristics during the return phase on the power plant's general performance. It was demonstrated that these can be summarised through two pivotal metrics, namely the glide ratio, E_{in}, and the resultant aerodynamic loading coefficient, $C_{R,in}$. Sample calculations showed that by lowering E_{in}, much faster reel-in speeds can be achieved as the kite is pushed further into the wind window and becomes less prone to overshoot to large angles of ϑ. Meanwhile, a reduction of C_R decreases the tether loads for a given reel-in speed, and thus less power is wasted during the return phase. An initial evaluation of arbitrary combinations of these metrics showed that they could realistically improve the average performance of a representative kite-power system by around 30%.

Next, a tool-chain centred around the commercial CFD suite Star-CCM+ was developed which allows for the aerodynamic investigation of arbitrary bleed air spoiler configurations. In order to validate the simulation environment, the setup of an independent experimental study on bleed air spoilers by Berexa et al. ([3]) was replicated as closely as possible using the original, provided canopy geometry. Based on several sensitivity studies, dependable grid resolution and time discretisation settings were established. With final cell counts of around 11M and time-accurate resolution of approximately 40 flow overs, the experimental data could be reproduced to a satisfactory degree. Results were generally agreed best at angles of attack with moderate flow separation. The application of bleed air spoilers did not lead to larger relative errors in comparison to the baseline configuration with no spoilers. Some discrepancies were noted between the respective baseline results at large angles of attack amounting to up to 15%. However, these data points are inconclusive and various wind tunnel effects could also be the reason. It was finally concluded that the simulation environment yields sufficiently dependable results.

For the second CFD study, a realistically billowed power kite geometry used by SkySails was adapted for simulations appropriately following the insights acquired during the validation phase. Once again, several sensitivity studies were carried out which indicate that cell counts at the order of 50M cells are required to accurately resolve the flow field. After establishing a method to obtain dependable time-averaged statistics of the forces and moments acting on the kite, two elaborate parameter studies were performed on the implementation of bleed air spoilers. In the first series, the chordwise size of the spoiler was varied. Results showed that all of the tested configurations are predicted to lead to average power output improvements of 13.8 to 18.3%. This was mostly thanks to substantial decreases in E between 67 and 76%, while the magnitude of C_R remained approximately the same with decreases in

C_L and increases in C_D compensating each other in their effect. Meanwhile, smaller spoiler sizes were effective in limiting the trim angle of attack to a smaller range. The results also highlighted an important limitation of the study, namely that one cannot capture the true unsteady nature of the system using a fixed angle of attack. Ideally, free movement of the canopy around the control pod would be possible, as enabled by fluid-structure coupled simulations. In the second parameter study, the influence of a spoiler's chordwise positioning was investigated. Although no improvement could be established with respect to the leading-edge positioned spoilers from the previous study, as the trim glide slope was consistently higher, the analysis showed that by moving the spoiler further to the rear, decreases of C_R can be achieved. The results also suggest that spoilers positioned at 20% of the local chord length or more lead to significant fore and aft pitching motions of the kite and that it is advisable to apply the spoiler within the reach of the suction peak.

All in all, the performed investigations shed light on the questions introduced at the start of this thesis in a number of ways. Besides establishing the grounds for dependable CFD simulations on canopies with various bleed air spoiler configurations, it was found that these features can be a powerful tool to change the reel-in conditions of kite power systems in favourable ways. Next to demonstrating their potential, the simulations also highlighted a number of secondary observations that should be kept in mind if some of the designs are turned into first prototypes. As such, the recorded pressures showed large variations in the trim angle of attack must be taken into account with respect to the design of the canopy's ram-air intakes. Also, relatively small bleed air spoilers may already entail a substantial benefit without affecting the internal pressure too much.

## 7.1. Recommendations for Future Work

Despite a large number of newly gained insights, the investigations which were carried out also leave a lot of new questions to be answered, opening a wide range of possibilities for future works. Some of the most important gaps are as follows:

- As addressed before, the assumption of a fixed angle of attack and a rigid canopy entails important limitations on the produced results. Future studies could pick up on the general CFD setup and best practices developed in this thesis and extend it with capacities for fluid-structure coupled simulations. These could drastically cut down on the simulation overhead for a given case because only one simulation would be necessary to obtain trim conditions, whilst enabling the physical capture of the kite's free movement.

- In part because of the large number of test cases that were treated in this thesis, the consideration of the detailed flow dynamics induced by bleed air spoilers shifted a bit out of focus. Instead, configurations were intentionally evaluated mostly on the basis of their time-averaged results to allow for tangible comparisons between the different designs. Future studies could focus on the further tuning of bleed air spoilers by spending more time scrutinising the aerodynamic phenomena at work in the simulations.

- By the end of this thesis, one parameter study required about two weeks of computations on 160 CPUs in a multi-node cluster environment. Clearly, the availability of such resources is a luxury that not all research groups can tap into. Future works could thus try to find ways to reduce the computational expense for a given simulation further. For instance, wall modelled CFD environments could be tested instead of the wall-resolved meshes used for the purpose of this study to see if they can replicate the results to a sufficient degree.

- Clearly, the parameter studies carried out in this work only scratch the surface of the solution space that exists with respect to different bleed air spoiler configurations. In addition to the investigation of dependencies on the size and chordwise positioning, the concurrent use of several spoilers could be further tested. It is not unlikely that an additional spoiler located further outboard would decrease the glide ratio even more and maybe decrease the resultant loading as well. Likewise, different spoiler geometries could be tested (e.g. perpendicular to the flow instead of at an angle).

Bibliography

[1] Yacine Addad, Ulka Gaitonde, Dominique Laurence, and Stefano Rolfo. Optimal unstructured meshing for large eddy simulations. *ERCOFTAC Series*, 12(January):93–103, 2008. ISSN 22151826. doi: 10.1007/978-1-4020-8578-9{_}8.

[2] Michael R. Barnes. Form finding and analysis of tension structures by dynamic relaxation. *International Journal of Space Structures*, 1999. ISSN 02663511. doi: 10.1260/0266351991494722.

[3] Emily Berexa, Robert Warner, Jurgen Seidel, and Keith Bergeron. Effect of Spoilers on Ram Air Parachute Aerodynamic Force and Moment Coefficients. In *AIAA Aviation 2019 Forum*, number June, Dallas, 2019. doi: 10.2514/6.2019-3370.

[4] K. Bergeron, M. Ward, M. Costello, and S. Tavan. AccuGlide 100 and Bleed-Air actuator airdrop testing. In *AIAA Aerodynamic Decelerator Systems (ADS) Conference*, number March, Daytona Beach, 2013. ISBN 9781624102028. doi: 10.2514/6.2013-1378.

[5] Keith Bergeron, Michael Ward, and Mark Costello. Aerodynamic effects of parafoil upper surface bleed air actuation. In *AIAA Atmospheric Flight Mechanics Conference 2012*, number August, Minneapolis, 2012. AIAA AVIATION Forum. ISBN 9781624101847. doi: 10.2514/6.2012-4737.

[6] Nedeleg Bigi, Alain Nême, Kostia Roncin, Jean Baptiste Leroux, Guilhem Bles, Christian Jochum, and Yves Parlier. Analytical tether model for static kite flight. *Green Energy and Technology*, 0 (9789811019463):57–78, 2018. ISSN 18653537. doi: 10.1007/978-981-10-1947-0{_}3.

[7] Michael Casey and Torsten Wintergerste. Quality and trust in industrial CFD - Best practice guidelines, 2000.

[8] Ismail B. Celik, Urmila Ghia, Patrick J. Roache, Christopher J. Freitas, Hugh Coleman, and Peter E. Raad. Procedure for estimation and reporting of uncertainty due to discretization in CFD applications. *Journal of Fluids Engineering, Transactions of the ASME*, 130(7):0780011–0780014, 2008. ISSN 00982202. doi: 10.1115/1.2960953.

[9] T Donné. *European Research Roadmap to the Realisation of Fusion Energy*. EUROfusion, Garching, 2018. ISBN 978-3-00-061152-0.

[10] P.A. Durbin and T.I-P. Shih. An overview of turbulence modeling. *Modelling and Simulation of Turbulent Heat Transfer*, 15:1–29, 2005. doi: 10.2495/978-1-85312-956-8/01.

[11] Michael Erhard and Hans Strauch. Flight control of tethered kites in autonomous pumping cycles for airborne wind energy. *Control Engineering Practice*, 40:13–26, 2015. ISSN 09670661. doi: 10.1016/j.conengprac.2015.03.001. URL http://dx.doi.org/10.1016/j.conengprac.2015.03.001.

[12] Michael Erhard and Hans Strauch. Automatic control of pumping cycles for the skysails prototype in airborne wind energy. *Green Energy and Technology*, 0(9789811019463):189–213, 2018. ISSN 18653537. doi: 10.1007/978-981-10-1947-0{_}9.

[13] Joel H. Ferziger and Milovan Peric. *Computational Methods for Fluid Dynamics*. Springer, Berlin, 3 edition, 2002. ISBN 3-540-42074-6.

[14] N. Fogell, L. Iannucci, and K. Bergeron. Fluid-structure interaction simulations of the inflated shape and associated flowfield of the MC4/5 parafoil during steady gliding flight. In *24th AIAA Aerodynamic Decelerator Systems Technology Conference*, number June, Denver, 2017. AIAA AVIATION Forum. ISBN 9781624105036. doi: 10.2514/6.2017-3543.

[15] Mikko Folkersma, Roland Schmehl, and Axelle Viré. Steady-state aeroelasticity of a ram-air wing for airborne wind energy applications. *Journal of Physics: Conference Series*, 1618(3), 2020. ISSN 17426596. doi: 10.1088/1742-6596/1618/3/032018.

[16] Alek Gavrilovski, Michael Ward, and Mark Costello. Parafoil glide slope control using canopy spoilers. *21st AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar 2011*, pages 1–11, 2011.

[17] Alek Gavrilovski, Michael Ward, and Mark Costello. Parafoil control authority with upper-surface canopy spoilers. *Journal of Aircraft*, 49:1391–1397, 2012. ISSN 15333868. doi: 10.2514/1.C031685.

[18] Mehdi Ghoreyshi, Keith Bergeron, Jürgen Seidel, Andrew J. Lofthouse, and Russell M. Cummings. Grid quality and resolution effects on the aerodynamic modeling of ram-air parachute canopies. *AIAA Aerospace Sciences Meeting*, 53(January), 2015. doi: 10.2514/6.2015-0407.

[19] Mehdi Ghoreyshi, Keith Bergeron, Jurgen Seidel, Adam Jirásek, Andrew J. Lofthouse, and Russell M. Cummings. Prediction of aerodynamic characteristics of ram-air parachutes. *Journal of Aircraft*, 53(6):1802–1820, 2016. ISSN 15333868. doi: 10.2514/1.C033763.

[20] M Higgins. Control System for Ram Air Gliding Parachutes, 1979.

[21] Ivan Komusanac, Guy Brindley, and Daniel Fraile. Wind energy in Europe in 2019, 2019.

[22] Eckart Laurien and Herbert Oertel. *Numerische Strömungsmechanik*. Springer, Wiesbaden, 5 edition, 2009. ISBN 978-3-658-03144-2 ISBN. doi: 10.1007/978-3-658-03145-9.

[23] Miles L. Loyd. Crosswind Kite Power. *Journal of energy*, 4(3):106–111, 1980. ISSN 01460412. doi: 10.2514/3.48021.

[24] Florian R Menter. AIAA 93 · 2906 Zonal Two Equation k · co Turbulence Models for Aerodynamic Flows . Mailing Address : 24th Fluid Dynamics Conference FOR AERODYNAMIC FLOWS. *Fluid Dynamics*, 1993.

[25] Stephen B. Pope. *Turbulent Flows*. Cambridge University Press, New York, 2000. ISBN 9780511840531. doi: 10.1017.

[26] RWE GmbH and SkySails GmbH. RWE Renewables and SkySails Power harness high-altitude winds for innovative power generation [Press Release], 2021. URL https://www.rwe.com/-/media/RWE/documents/07-presse/rwe-renewables-gmbh/2021-01-26-rwe-renewables-and-skysails-power-harness-winds-for-innovative-power-generation.pdf.

[27] Roland Schmehl, editor. *Airborne Wind Energy - Advances in Technology Development and Research*. Springer, Singapore, 2018. ISBN 978-981-10-1947-0. doi: 10.1007/978-981-10-1947-0{_}26.

[28] Robert Schorr. *Numerical Methods For Parabolic-Elliptic Interface Problems*. PhD thesis, TU Darmstadt, 2019.

[29] SkySails GmbH. Wind Power 2.0: Revolutionary Airborne Wind Energy System to Provide Green Power to the Republic of Mauritius, 2020. URL https://skysails-power.com/kite-power-for-mauritius/.

[30] SkySails GmbH. SkySails Power Imagebrochure, 2021. URL https://skysails-group.com/wp-content/uploads/2021/05/SkySails_Power_Imagebrochure_28-02-2021-1.pdf.

[31] United Nations. The 2030 Agenda for Sustainable Development, 2016.

[32] Versteeg H.K and Malalasekera W. *Introduction To Computational Fluid Dynamics*. Longman Scientific & Techincal, Essex, 1995. ISBN 0-582-21884-5.

[33] Frank M. White. *Fluid mechanics (seventh ed.)*. McGraw-Hill, New York, 2011. ISBN 978–0–07–352934–9.

A
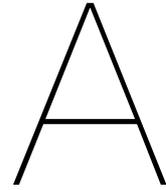
Supplements to Chapter 6

The depicted flow scenes have identical scaling for the velocity field (cut section) and pressure coefficient distribution (canopy surface) to allow for relative comparisons. All scenes show an instantaneous (not-time averaged) flow field at the trim angle of attack. Colour bars were intentionally omitted for confidentiality reasons.

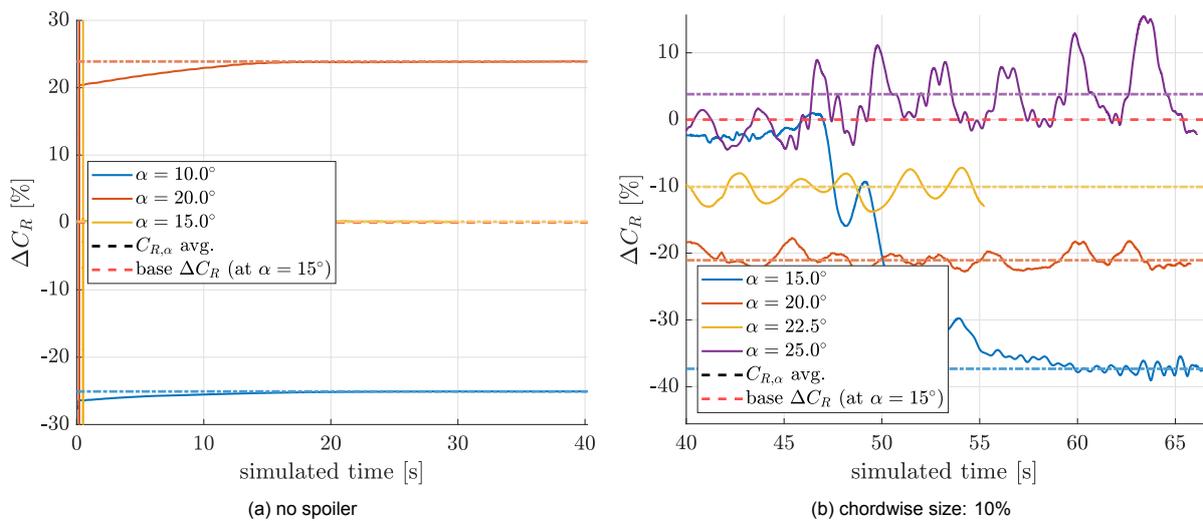## A.1. Parameter Study on Spoiler Sizing

Time histories

(a) no spoiler

(b) chordwise size: 10%

Figure A.1: C_R histories at various α with differently sized bleed-air spoilers (continued on the next page)

(c) chordwise size: 25%

(d) chordwise size: 100%
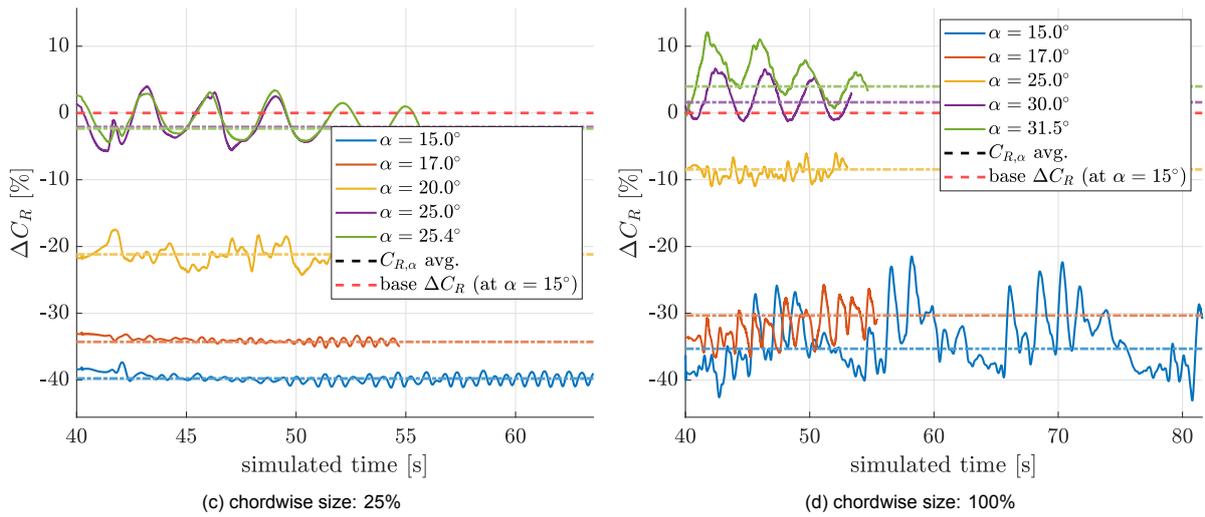
Figure A.1: C_R histories at various α with differently sized bleed-air spoilers (continued)

Velocity field views
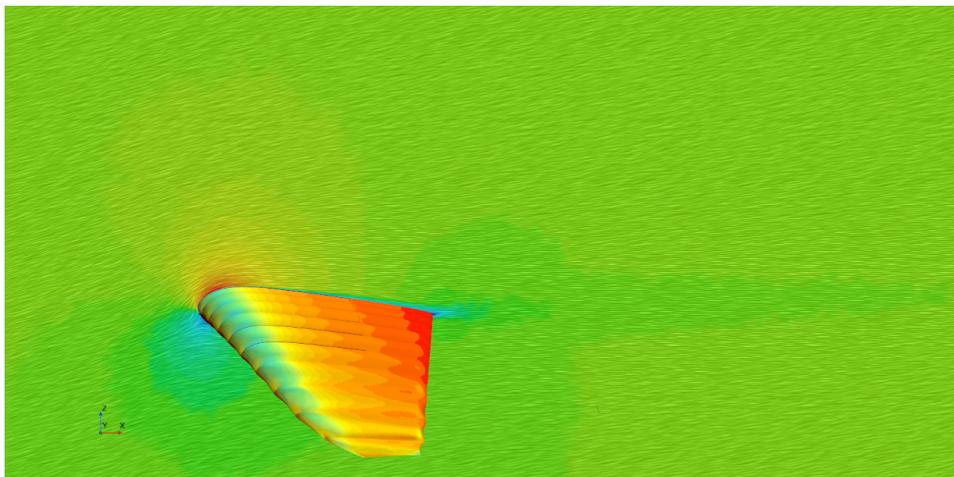
(a) no bleed-air spoiler; α_{trim} = 15°

Figure A.2: Velocity fields on vertical cutting plane that slices through bleed-air spoiler; C_p field on canopy surface (continued on the next page)
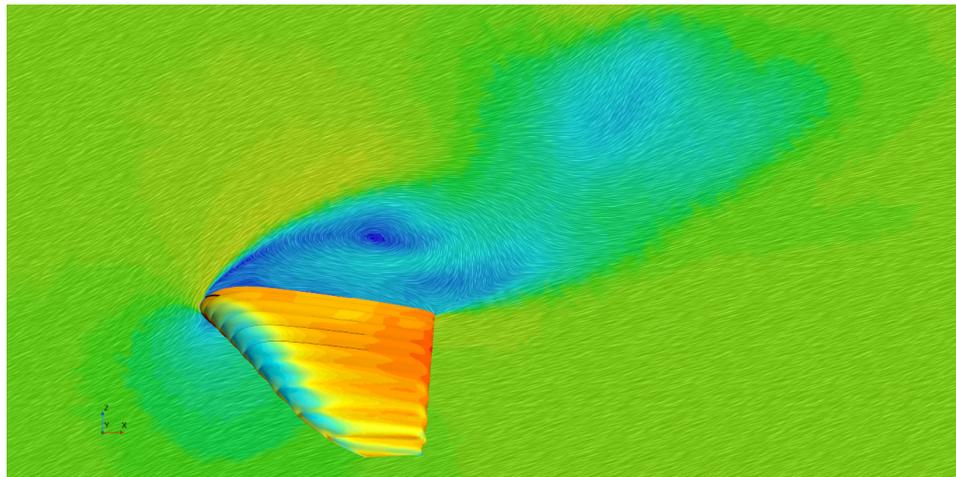
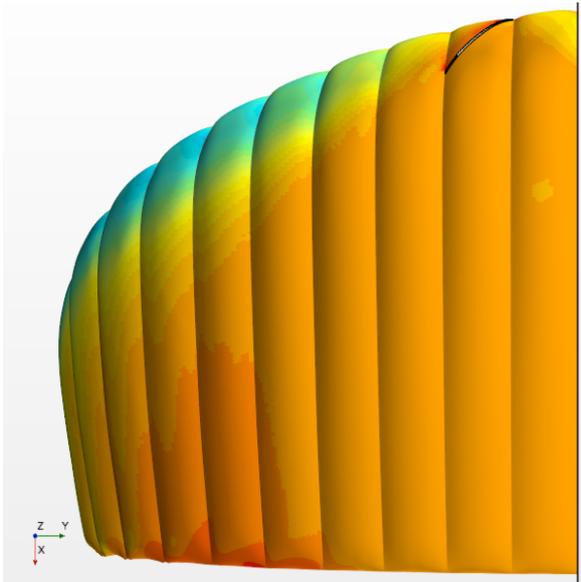(b) bleed-air spoiler size: 10% of reference; $\alpha_{trim} = 25°$
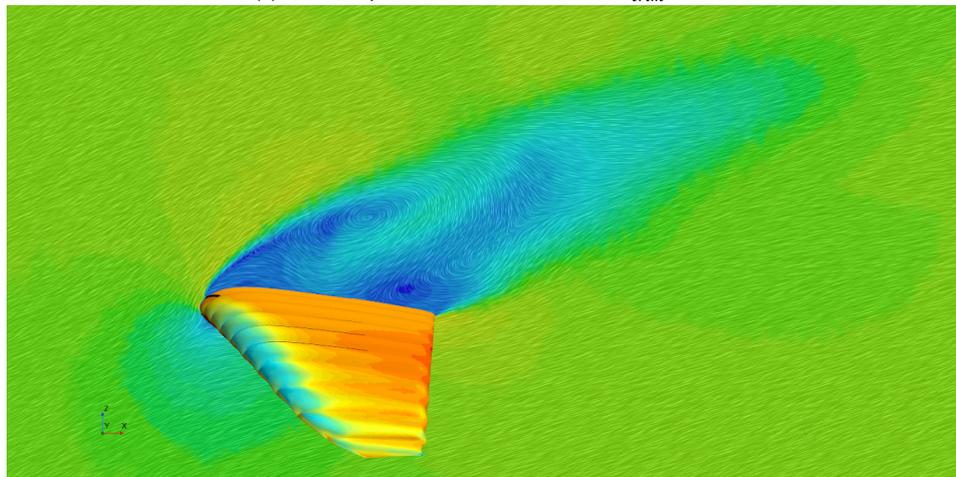
(c) bleed-air spoiler size: 25% of reference; $\alpha_{trim} = 25.4°$
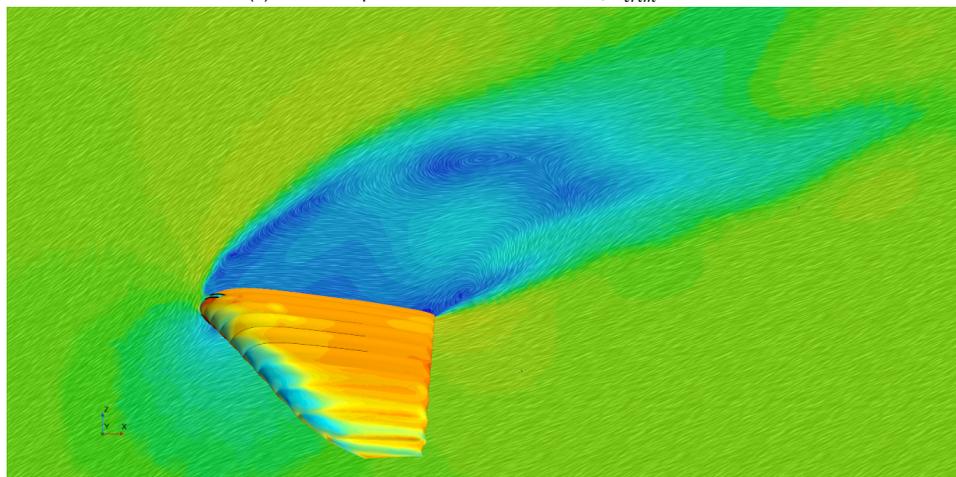
(d) bleed-air spoiler size: 100% of reference; $\alpha_{trim} = 30°$

Figure A.2: Spoiler sizing; Velocity fields on vertical cutting plane that slices through bleed-air spoiler; pressure coefficient, C_p, field on canopy surface (continued)

Canopy surface C_p views
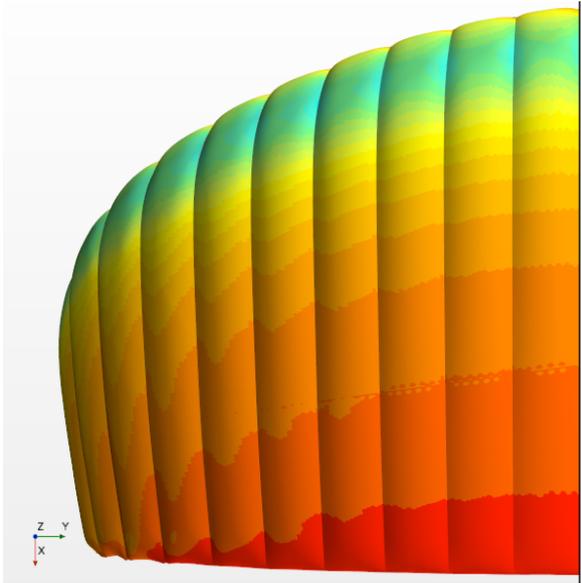
(a) no bleed-air spoiler; $\alpha_{trim} = 15°$
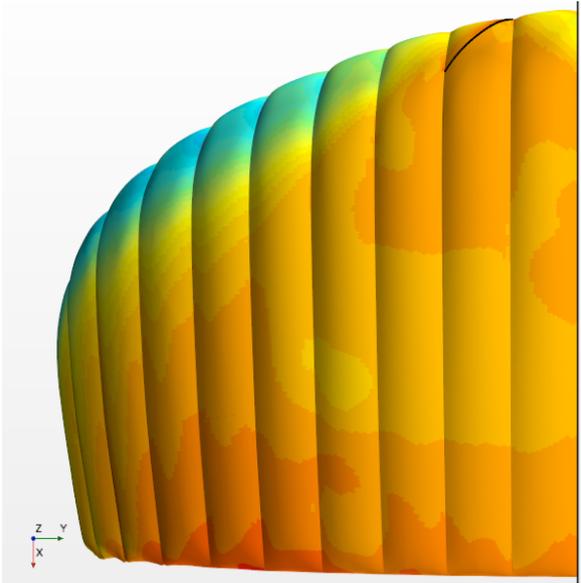
(b) bleed-air spoiler size: 10% of reference; $\alpha_{trim} = 25°$

(c) bleed-air spoiler size: 25% of reference; $\alpha_{trim} = 25.4°$
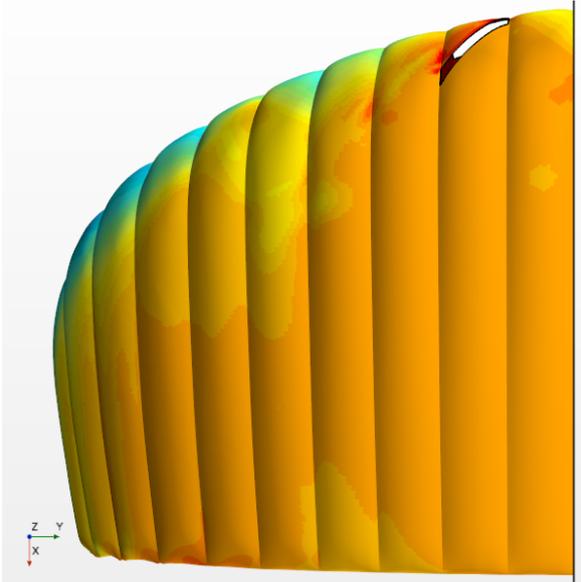
(d) bleed-air spoiler size: 100% of reference; $\alpha_{trim} = 30°$

Figure A.3: Spoiler sizing; Top view of pressure coefficient, C_p, field on canopy surface

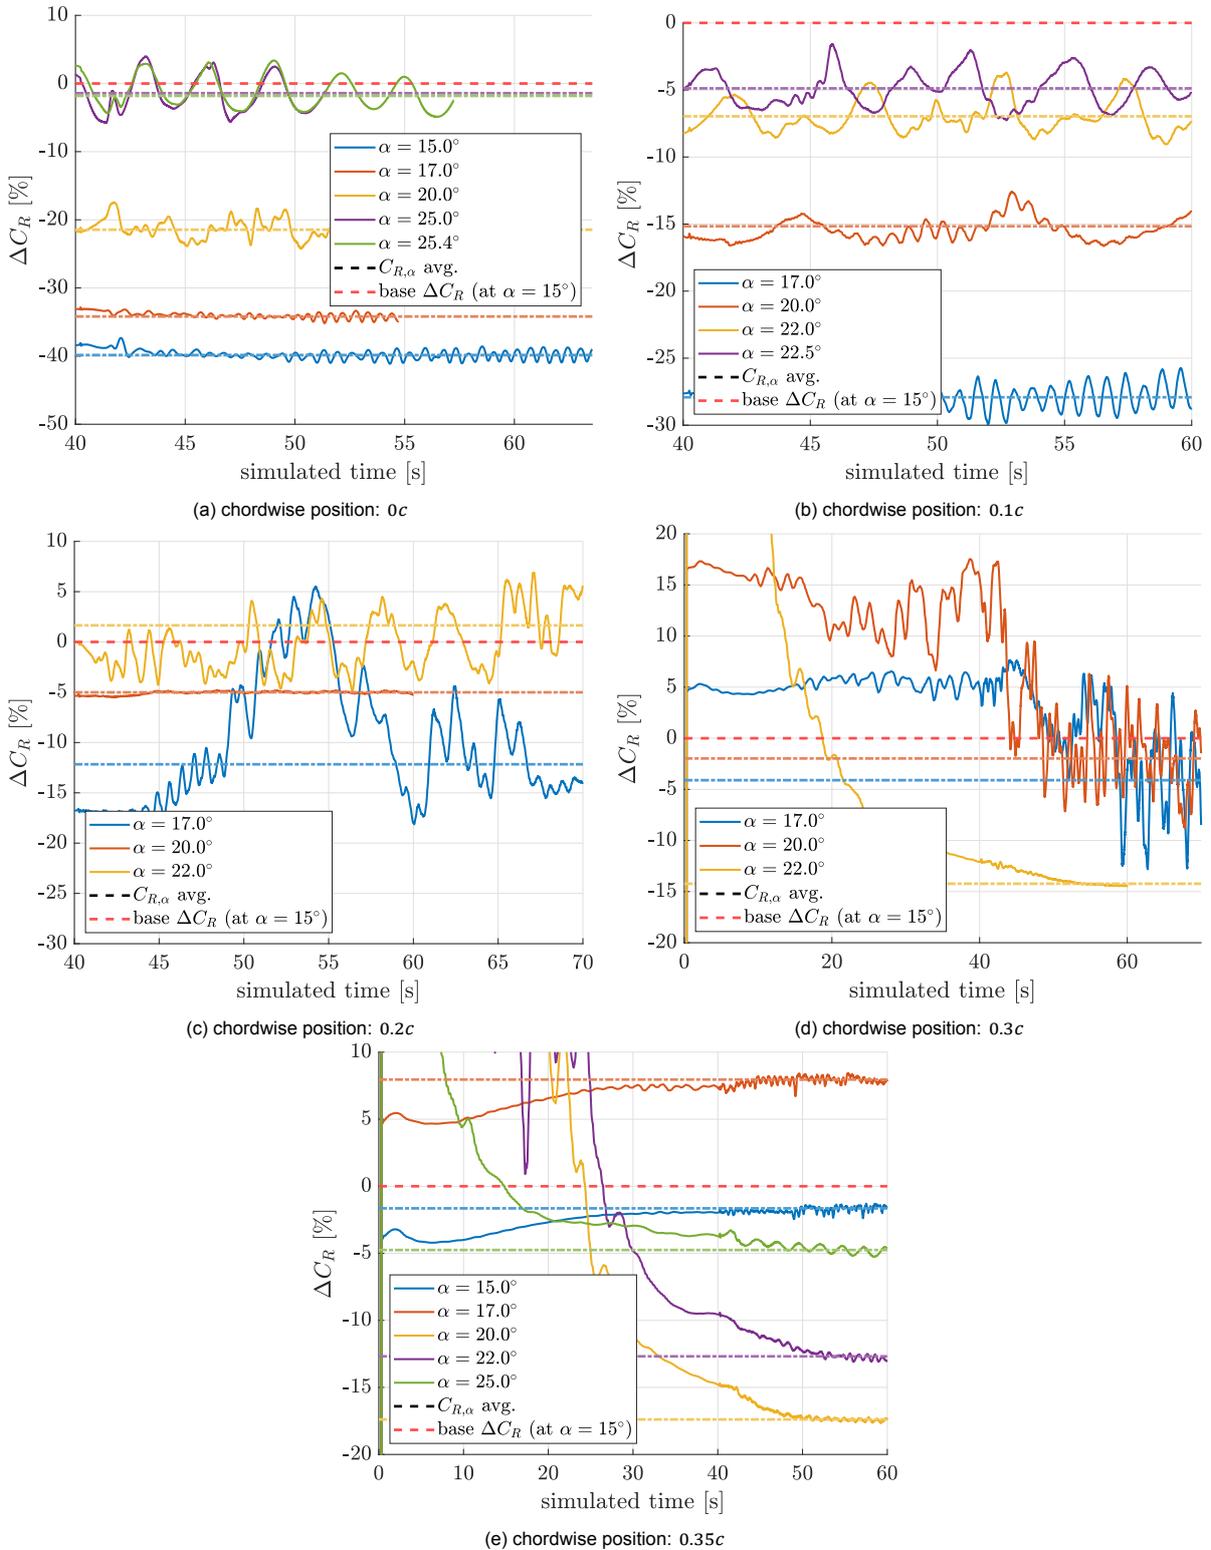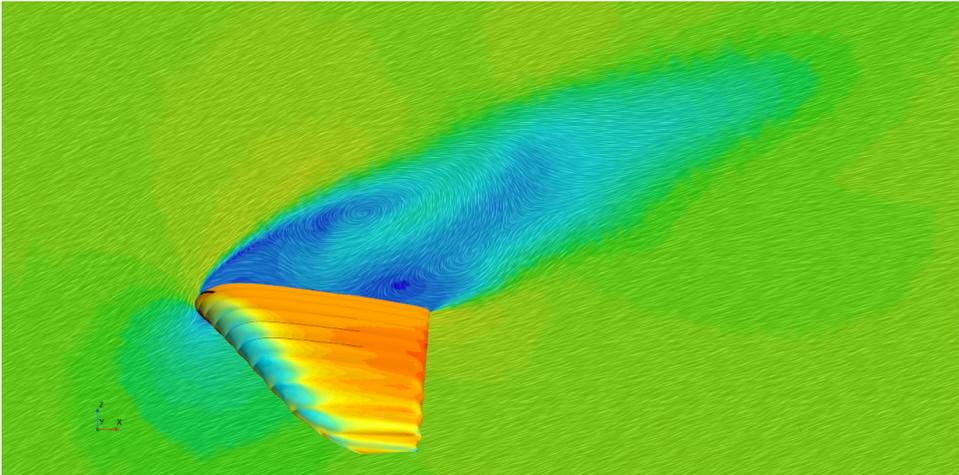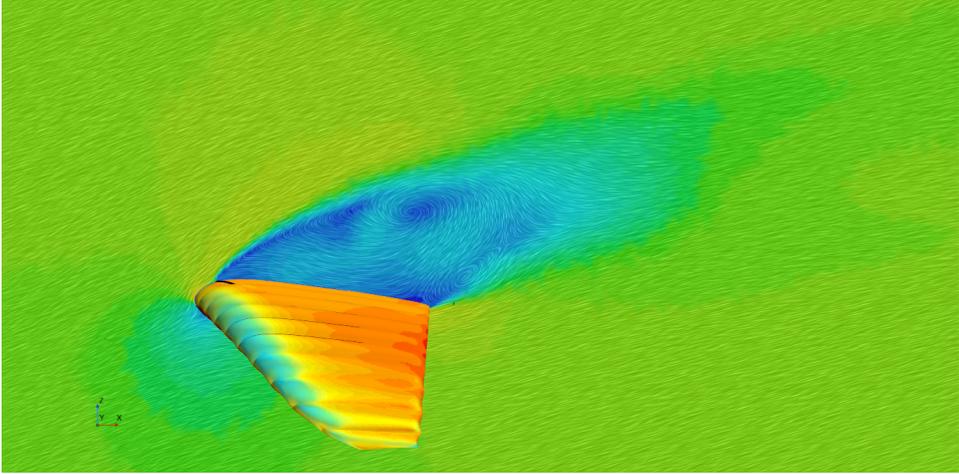## A.2. Parameter Study on Chordwise Positioning of Spoilers

Time histories

(a) chordwise position: $0c$

(b) chordwise position: $0.1c$

(c) chordwise position: $0.2c$

(d) chordwise position: $0.3c$

(e) chordwise position: $0.35c$

Figure A.4: C_R histories at various α with bleed-air spoilers of varying chordwise location

Velocity field views

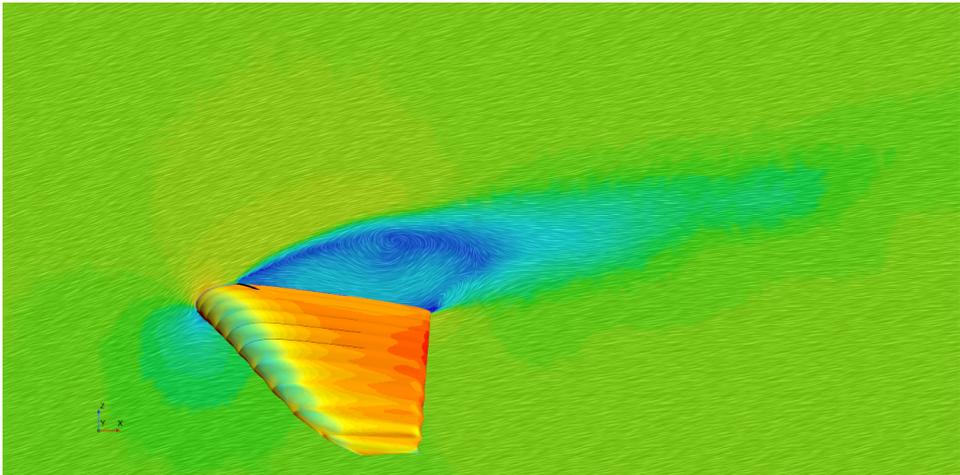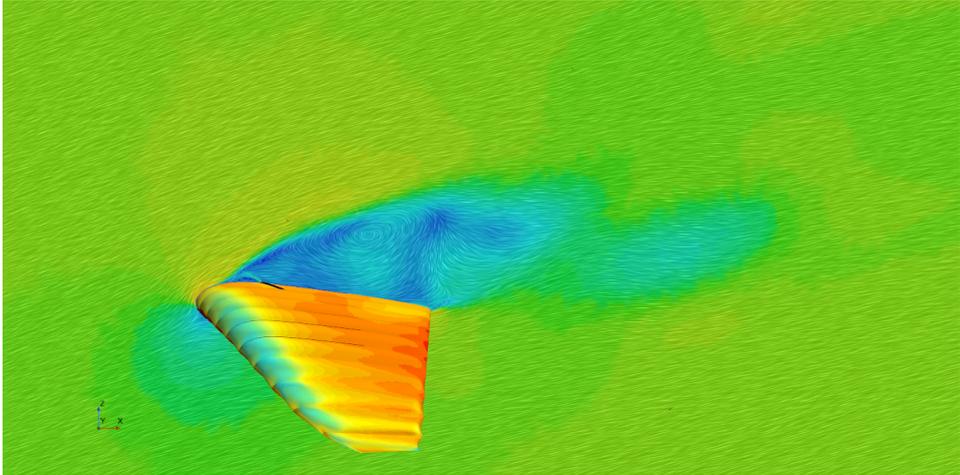(a) spoiler position: $0c$; $\alpha_{trim} = 25.4°$

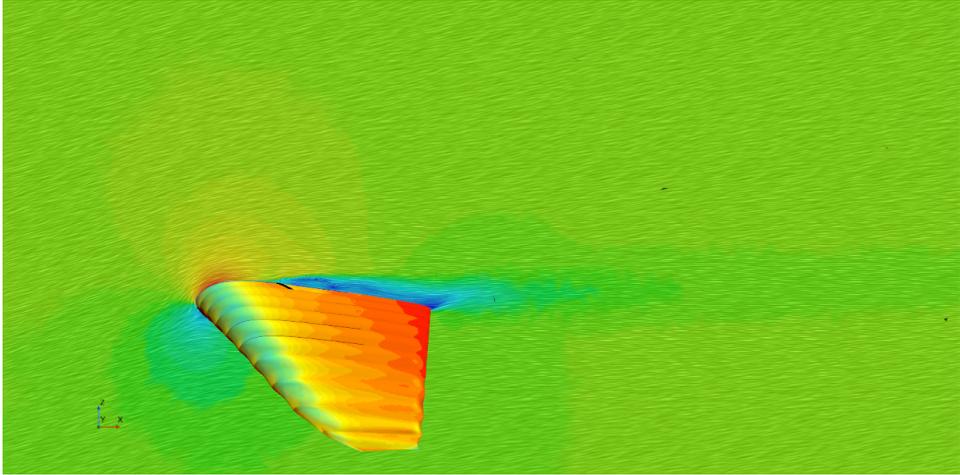

(b) spoiler position: $0.1c$; $\alpha_{trim} = 22°$

Figure A.5: Velocity fields on vertical cutting plane that slices through bleed-air spoiler; C_p field on canopy surface (continued on next page)

(c) spoiler position: $0.2c$; $\alpha_{trim} = 17°$

(d) spoiler position: $0.3c$; $\alpha_{trim} = 20°$

(e) spoiler position: $0.35c$; $\alpha_{trim} = 15°$

Figure A.5: Chordwise spoiler positioning; Velocity fields on vertical cutting plane that slices through bleed-air spoiler; pressure coefficient, C_p, field on canopy surface (continued)