Power to the airborne wind energy performance model

Estimating long-term energy production with an emphasis on pumping flexible-kite systems

Schelbergen, M.

DOI
10.4233/uuid:353d390a-9b79-44f1-9847-136a6b880e12

Publication date
2024

Document Version
Final published version

Citation (APA)
Schelbergen, M. (2024). *Power to the airborne wind energy performance model: Estimating long-term energy production with an emphasis on pumping flexible-kite systems*. [Dissertation (TU Delft), Delft University of Technology]. https://doi.org/10.4233/uuid:353d390a-9b79-44f1-9847-136a6b880e12

POWER TO THE AIRBORNE WIND ENERGY PERFORMANCE MODEL

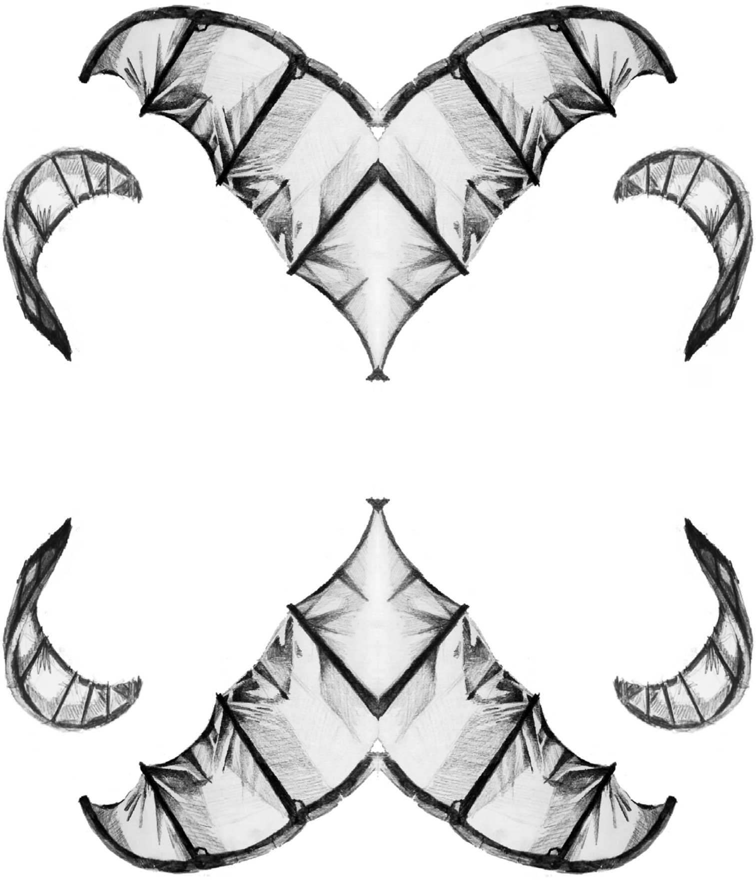

ESTIMATING LONG-TERM ENERGY PRODUCTION WITH AN EMPHASIS ON PUMPING FLEXIBLE-KITE SYSTEMS

M. SCHELBERGEN

Power to the airborne wind energy performance model

Estimating long-term energy production with an emphasis on pumping flexible-kite systems

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. dr. ir. T.H.J.J. van der Hagen,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op maandag 18 maart 2024 om 15:00 uur

door

MARK SCHELBERGEN

Master of Science in Aerospace Engineering,
Delft University of Technology, Delft, Nederland,
geboren te Terneuzen, Nederland.

Dit proefschrift is goedgekeurd door de promotoren.

Samenstelling promotiecommissie bestaat uit:

Rector Magnificus,	voorzitter
Prof. dr. S. J. Watson,	Technische Universiteit Delft, promotor
Dr.-ing. R. Schmehl,	Technische Universiteit Delft, promotor

Onafhankelijke leden:

Prof. dr. ir. J.W. van Wingerden	Technische Universiteit Delft
Prof. dr. C.L. Archer	University of Delaware
Prof. dr. P.W. Cheng	Universität Stuttgart
Prof. dr. G. Sánchez-Arriaga	Universidad Carlos III de Madrid
Prof. dr. D.A. von Terzi	Technische Universiteit Delft, reservelid

Funded by: REACH (H2020-FTIPilot-691173) funded by the European Union's Horizon 2020 research and innovation programme under grant agreement No. 691173

Keywords: Airborne wind energy, flexible-kite system, performance modelling and optimisation, energy production estimation, power production characterisation, vertical wind profile characterisation, test flight data analysis

Printed by: Ipskamp Printing

Cover design: Niels Bierens

ISBN 978-94-6384-549-6

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

Contents

Summary

Only in a scenario in which utility-scale airborne wind energy (AWE) systems are deployed on a large scale, can AWE make a significant contribution to the energy transition. The viability of large-scale deployment of AWE closely depends on the cost-competitiveness and complementarity of AWE with respect to tower-based wind turbines. The long-term energy production is a key aspect of these two criteria. It is often claimed that AWE systems have better power generation characteristics than wind turbines as they can adjust their operating height to the altitude with the most favourable wind conditions. This optimal operating height can be derived from the vertical wind profile, which describes how wind speed changes with height. The variation of this wind profile is not adequately reflected in efficient long-term energy production estimations and, thus, introduces uncertainty to these estimations.

This thesis aims to answer the following research question: how significant is the wind profile variability to the annual energy production estimation?

A systemic approach is required to evaluate the sensitivity of the modelled energy production to the wind profile variability. This is achieved by using a newly developed energy production estimation framework that expands existing characterisations of the wind climate and power production. This framework aims to incorporate wind profile variation with an adjustable level of detail.

The first part of this thesis establishes the climatology of the vertical wind profile. It starts with a spatial analysis of the potential benefits of flexible-height operation. Subsequently, the characteristics of prevalent wind profiles are studied for an on- and offshore location in the Netherlands. To this end, hourly wind profiles from wind atlas data are used. Every wind profile is decomposed into a shape function and magnitude component. Next, clustering is used to derive eight prevalent wind profile shapes for each location. The resulting sets of mean-cluster wind profile shapes, with approximately an equal share of logarithmic and jet-like shapes, are employed to characterise the wind climate. A unique aspect of the derived characterisation is that it incorporates jet-like wind profiles, which cannot be achieved using conventional wind profile relationships. The good agreement between the identified clusters and local weather patterns in terms of their temporal occurrence and their relationship to atmospheric stability suggests that an adequate wind profile climatology is derived.

The second part of this thesis comprises the modelling of the flight operation. To characterise the flight behaviour of a 25 m^2 flexible kite, a three-hour test flight dataset of the development system of Kitepower B.V. is utilised, in which 87 pumping cycles are flown. The measured kite attitude is employed to decompose the steady-state aerodynamic force for the identification of the lift and drag coefficients of the kite. Due to ineffective measurements, no clear relationships with the angle of attack could be identified. Instead, a workable model structure has been derived, which solely depends on the depower signal.

After a preliminary flight trajectory reconstruction, the swinging motion of the kite is studied with two approaches: by approximating the motion as a transition through steady-rotation states and by solving the motion dynamically. The kite is modelled with two rigidly linked point masses representing the control unit and wing. Despite its simple form, this two-point kite model conveniently extends a discretised tether model and computes the roll and pitch of the kite with sufficient accuracy to allow adequate modelling of the turning mechanism and kite aerodynamics.

The validation of the pumping flight operation model necessitates detailed information about the instantaneous vertical wind profile, for which the wind velocity measured near the ground with an anemometer is insufficient. However, information about the wind profile can be gained by employing the kite as a flying probe to measure the relative wind. The validation shows that the quasi-steady model poorly explains the measured mean cycle power trend with wind speed. The discrepancy between the modelled and measured trends also emerges when assessing the performance estimation on the level of the figure-of-eight manoeuvres. This mismatch is predominantly explained by a deficiency in the wind speed reconstruction, yielding a distorted view of the model validity.

The last part of this thesis synthesises the findings from the previous parts in a novel framework for estimating the annual energy production for AWE systems. The basic pumping flight operation model is employed to generate one power curve for every mean-cluster wind profile shape by conducting 50 optimisations for increasing magnitudes, expressed with the wind speed at 200 m height. The optimisation effectively explores the operational decision space to find the pumping cycle operations that maximise power output. The resulting power curves are compared against the 5^{th}–95^{th} percentile power band and the mean curve obtained by pumping cycle optimisations for every hourly wind profile in one year of wind atlas data. The observed differences between the characterisations are explained by the loss of detailed wind profile features due to the clustering.

To evaluate the impact of particular wind profile shapes on the annual energy production estimation, a sequence of estimations is conducted with the developed framework in which progressively more wind profile shapes are considered. Including the jet-like wind profile shapes substantially impacts the estimate and underlines the importance of these wind profiles to the energy production.

The significant sensitivity of the energy production estimation to the wind profile variability found in this work underlines that wind resource assessments for AWE systems need to move away from using the calculation methods borrowed from conventional wind energy. The newly developed framework provides a more accurate alternative. It enables the assessment of the operational benefits of AWE systems compared to tower-based wind turbines and, thus, can be utilised to further improve the understanding of the viability of large-scale deployment of AWE systems. Further research is still needed to validate this framework.

Samenvatting

Alleen in een scenario waarin op grote schaal vliegersystemen worden ingezet, kunnen deze een significante bijdrage leveren aan de energietransitie. De levensvatbaarheid van grootschalige inzet van vliegersystemen hangt sterk af van hun aanvullend en concurrentievermogen ten opzichte van conventionele windturbines. De lange termijn energieproductie is van groot belang voor deze twee criteria. Vaak wordt beweerd dat vliegersystemen efficiënter zijn dan windturbines, omdat ze hun vlieghoogte kunnen aanpassen naar de hoogte met de gunstigste windomstandigheden. Deze optimale vlieghoogte kan worden afgeleid uit het verticale windprofiel, dat beschrijft hoe de windsnelheid verandert met de hoogte. De veranderlijkheid van dit windprofiel wordt niet adequaat weerspiegeld in schattingen van de lange termijn energieproductie en introduceert daarmee onzekerheid in deze schattingen.

Dit proefschrift heeft als doel om de volgende onderzoeksvraag te beantwoorden: hoe significant is de veranderlijkheid van het windprofiel voor de schatting van de jaarlijkse energieproductie?

Een systemische aanpak is nodig om de gevoeligheid van de gemodelleerde energieproductie voor de veranderlijkheid van het windprofiel te evalueren. Dit wordt bereikt door een nieuw ontwikkeld raamwerk voor het schatten van de energieproductie te gebruiken die verder bouwt op bestaande karakteriseringen van het windklimaat en de energieproductie. Dit raamwerk kan de windprofielvariatie meenemen met een instelbaar detailniveau.

Het eerste deel van dit proefschrift onderzoekt het klimaat van het verticale windprofiel. Na het analyseren van de potentiële voordelen van een aanpasbare vlieghoogte, worden kenmerken van veelvoorkomende windprofielen bestudeerd voor een on- en offshore locatie in Nederland. Hiervoor wordt gebruik gemaakt van uurlijkse windprofielen uit windatlasdata. Elk windprofiel wordt ontleed in een vormfunctie en een groottecomponent. Vervolgens worden er voor elke locatie acht veelvoorkomende windprofielvormen afgeleid met behulp van clustering. In de resulterende sets windprofielvormen heeft de helft een logaritmische vorm en de andere helft een laag maximum. Deze sets worden gebruikt om het windklimaat te karakteriseren. Een uniek aspect van deze karakterisering is dat het niet-logaritmische windprofielen meeneemt. De overeenkomst tussen de geïdentificeerde clusters en lokale weerpatronen betreft hun timing en relatie met atmosferische stabiliteit suggereert dat het afgeleide windprofielklimaat adequaat is.

Het tweede deel van dit proefschrift omvat het modelleren van de pompcyclus vlucht van de vlieger. Om het vlieggedrag van een 25 m^2 flexibele vlieger te karakteriseren, wordt gebruik gemaakt van een drie uur durende testvlucht dataset van het ontwikkelsysteem van Kitepower B.V., waarin 87 pompcycli worden gevlogen. De gemeten vliegeroriëntatie wordt gebruikt om de aerodynamische kracht in stationaire toestand te ontleden en daarmee de lift- en weerstandscoëfficiënten van de vlieger te identificeren. Door

ineffectieve metingen konden geen duidelijke relaties met de invalshoek worden vastgesteld. Een werkbare modelstructuur is afgeleid welke uitsluitend afhankelijk is van het depower signaal.

Na de reconstructie van het vliegtraject is de schommelbeweging van de vlieger op twee manieren bestudeerd: door de beweging te benaderen als een overgang door stationaire-rotatie toestanden en door de beweging dynamisch op te lossen. De vlieger is gemodelleerd met één puntmassa voor de besturingsunit en één voor de vleugel, welke onderling star zijn verbonden. Dit vliegermodel breidt op handige wijze een gediscretiseerd kabelmodel uit. Ondanks de eenvoudige modelstructuur, wordt het rollen en stampen van de vlieger met voldoende nauwkeurigheid berekend om het stuurmechanisme en de aerodynamica van de vlieger te modelleren.

De validatie van het rekenmodel vereist gedetailleerde informatie over het momentane verticale windprofiel. De windsnelheid gemeten vlak boven de grond met een anemometer is hiervoor ontoereikend. Deze informatie kan echter worden afgeleid uit de relatieve windmetingen van de vlieger. De validatie laat zien dat de trend van het gemeten opgewekte vermogen als functie van de windsnelheid niet goed wordt verklaard door het rekenmodel. De discrepantie tussen de gemodelleerde en gemeten trends komt ook naar voren bij het bepalen van het opgewekte vermogen van de achtvormige manoeuvres. Deze mismatch wordt voornamelijk verklaard door een fout in de windsnelheidsreconstructie, waardoor een vertekend beeld ontstaat van de validiteit van het model.

Het laatste deel van dit proefschrift brengt de bevindingen uit de eerdere delen onder in een nieuw raamwerk voor het schatten van de jaarlijkse energieproductie van vliegersystemen. Het rekenmodel van de vlucht wordt gebruikt om één vermogenscurve te genereren voor elk veelvoorkomende windprofielvorm door 50 optimalisaties uit te voeren voor toenemende windsterktes. De optimalisatie is effectief in het vinden van de pompcyclus die de vermogensopbrengst maximaliseert. De resulterende vermogenscurven worden vergeleken met de 5^{e}–95^{e} percentielvermogensband en de gemiddelde curve, verkregen met optimalisaties voor elk uurlijks windprofiel in één jaar aan windatlasdata. De waargenomen verschillen kunnen worden verklaard door het verlies van gedetailleerde windprofielkenmerken als gevolg van de clustering.

Om de impact van bepaalde windprofielvormen op de schatting van de jaarlijkse energieproductie te evalueren wordt het ontwikkelde raamwerk gebruikt om een reeks schattingen uit te voeren, waarbij elke schatting meer windprofielvormen meeneemt. Het meenemen van de windprofielvormen met een laag maximum heeft een aanzienlijke invloed op de schatting en onderstreept het belang van deze windprofielen voor de energieproductie.

De aangetoonde gevoeligheid van de energieproductieschatting voor de veranderlijkheid van het windprofiel onderstreept dat studies naar vliegersystemen moeten afstappen van het gebruik van rekenmethodes die zijn geleend van conventionele windturbines. Het nieuw ontwikkelde raamwerk biedt een nauwkeuriger alternatief. Dit raamwerk maakt het mogelijk de operationele voordelen van vliegersystemen in vergelijking tot conventionele windturbines mee te nemen en kan worden gebruikt om het inzicht in de levensvatbaarheid van grootschalige inzet van vliegersystemen te verbeteren. Verder onderzoek is nodig om dit raamwerk te valideren.

Preface

Airborne wind energy technology is coloured by the enthusiastic people who work in the field. Many of them are (aerospace) engineers working on robust solutions for controlling the tethered kite. We engineers love to simplify parts of a problem, for example, the wind conditions, to be able to wrap our heads around other parts, such as the control system. A common simplification involves assuming a logarithmic shape of the vertical wind profile, which describes how the wind speed changes with height. Depending on the application, meteorologic scientists, who are fewer in number within the field, are less likely to approve this assumption. I hope this work encourages engineers to expand their knowledge of meteorological science with the help of the engineering solutions provided in this thesis.

Working at the intersection of meteorology and engineering during my PhD has been very exciting for me. It has been a great experience to work on such a complex technology in a small community. Even as a researcher, it felt like I was right between the people making it happen. I am curious to witness the further development of airborne wind energy technology, and I hope that when I look back in a few years, I can say that I made a small contribution to its success.

Delft, February 2024
Mark Schelbergen

1

Introduction

The potential contribution of airborne wind energy to the energy transition

To mitigate climate risks, net-zero CO_2 emissions and massive reductions of other greenhouse gasses need to be achieved around 2050 [1]. Only then can global warming be kept below 1.5°C as set out in the Paris Agreement [2]. In 2017, global warming induced by human activities reached the 1°C mark compared to pre-industrial levels. Still, global temperatures are rising at an alarming rate of 0.2°C per decade. This has dire consequences on natural and human systems, e.g., coral bleaching at the Great Barrier Reef induced by marine heatwaves jeopardises the larger marine ecosystem, and more severe droughts in the Horn of Africa have led to a higher risk of famines.

Decarbonising our energy system is an important driver for reaching net-zero CO_2 emissions by 2050 with a leading role for solar and wind energy. The energy system comprises the power, heat and transport sectors and is presently responsible for 73% of global greenhouse gas emissions [3]. Although the decarbonisation of the power sector is underway, the current trends do not meet the targeted yearly CO_2 emission reduction of 7.6% [4]. On the one hand, decarbonisation requires a reduction of global energy consumption by improving the efficiency of the energy systems. On the other hand, the composition of the energy mix needs to change rapidly by deploying low-carbon energy sources and moving away from fossil fuels [5]. To this end, the electrification of heat, transport, and industry is an important driver. The exact composition of the future energy mix depends not only on the decarbonisation policy but also on technological advancements. Ultimately, wind energy may even provide up to half of the future electricity generation [6].

The energy mix transition towards a larger share of renewables requires new innovations across all levels of the power system, including wind energy technology. The generation intermittency inherent to wind and solar power imposes new requirements on the demand and supply side of the power system related to long-term capacity adequacy and short-term power balancing. For wind energy technology, this translates to developing even larger turbines and improving the synergy between individual turbines, wind parks, and the rest of the power system [6]. Additionally, emerging wind technologies, such as airborne wind energy (AWE) and offshore floating concepts, could potentially contribute to the energy transition.

AWE technology is an appealing alternative to conventional tower-based wind technology as it has potential social and technical benefits but still needs more fundamental research to become cost competitive [7]. AWE systems employ flying devices tethered to the ground, referred to as kites, to harvest energy from the wind. Most AWE systems employ fast cross-wind flight, in which the kite has a similar function as the blade of a conventional wind turbine. The tether connects the kite to the ground; therefore, no material-intensive tower structure is needed. As the visibility of the kite in the landscape is less than that of a tower-based turbine, it is claimed that the resistance of local residents to the deployment of AWE systems will be less. However, no scientific proof exists yet to support this claim [8]. Concerning the technology, it is often claimed that AWE systems have better power generation characteristics as they are highly flexible in terms of operation and can harness untapped strong winds at higher altitudes. This flexibility comes at the price of increased dependency on the control system with associated stringent safety requirements [9].

Before AWE can serve the utility market, the technology must be scaled up, capable of operating in farms, and reach a similar autonomy level as tower-based wind turbines [10]. To reach these goals, the technology still needs fundamental research to, among others, increase system robustness and fault tolerance. Despite not being fully autonomous, the first products are already being commercialised. Instead of directly serving the utility market, these are expected to serve niche markets, such as remote areas and islands [11, 12], in which minimal supervision is still acceptable. However, to reach the technical readiness needed for serving the utility market, a wide 'valley of death' of the product development cycle still needs to be bridged, requiring large private and public investments [13].

The uncertainty of large-scale deployment

To attract investments and accelerate the ongoing product development, the cost-competitiveness and complementarity of AWE with respect to tower-based wind turbines should be better understood. The viability of large-scale deployment of AWE is closely related to these two aspects. As the current claims about the cost of energy of future machines are rough, they introduce a high uncertainty in the viability assessment of AWE. This was one of the reasons why AWE was still labelled as a high-risk technology in 2018 [13]. As investments are hard to justify due to associated risks, only limited funds are available for further product development, slowing down progress.

Even without evaluating an AWE system at a much larger scale than currently existing systems, an estimate of the cost of energy is subject to a degree of uncertainty depending on how adequately a computational model reflects reality. The cost of energy is estimated based on energy production, cost (capital, operation, and maintenance), and the lifetime of the system. This thesis focuses on the energy production estimation provided by a long-term performance model. The ultimate check for evaluating the uncertainty of the long-term performance model is validation with operational data covering an extended period. However, this data is not yet available as the existing development systems have only just started operating more continuously.

The long-term performance model connects the wind conditions at the deployment site with the power production characteristics of the system. An efficient model is obtained by separating these two aspects and statistically summarising the wind climate. By using a site-independent power production characterisation, the power production only needs to be characterised once for a system. On the contrary, the wind conditions are site-dependent and should be characterised for every deployment site.

Typically, long-term performance assessments for AWE systems borrow the wind climate characterisation from conventional wind technology. This characterisation describes the wind speed statistics at a single height. Such characterisation is sufficient for tower-based wind turbines because their long-term performance correlates well with the wind conditions at hub height. However, AWE systems may operate at different heights depending on the vertical wind profile, i.e., the wind speed variation with height. Therefore, it is more relevant for AWE systems to consider a larger height range in the wind climate characterisation.

Complementary to the single-height wind climate description, the conventional power curve is typically used to describe the power production characteristics of the AWE system. The power curve only describes the relationship between power output and the wind speed at a given height without covering any details about the wind profile. Therefore, it cannot fully reflect the operational flexibility of AWE systems, which can be adjusted to the specifics of the wind profile to maximise power output.

Employing conventional characterisations in energy production estimation for AWE systems introduces structural uncertainties. This is because the wind climate over a large height range is not adequately reflected, and the effect of the flexible-height operation of AWE systems is not well incorporated. By identifying the effect of wind profile variability on energy production, the impact of these uncertainties on the energy production estimation will be better understood. Moreover, it will provide valuable information about the level of detail required from the long-term performance model to meet accuracy requirements.

Increasing confidence in performance models

This thesis aims to assess how significant the wind profile variability is to the annual energy production estimation. Due to limited resources, conducting new experiments is not considered. Instead, the aim is pursued foremost with a model-based approach. An existing AWE development system is repeatedly studied throughout this thesis. Even though only one particular system is investigated, many of the findings concerning energy production estimation can be generalised, and the methodologies developed in this work can be applied to different types of up-scaled systems.

A systemic approach is adopted to enable evaluating the sensitivity of the modelled energy production to the wind profile variability. This approach includes expanding existing wind climate and power production characterisations to incorporate the wind profile variation with an adjustable detail level. These can then be used in a convergence study in which the energy production is estimated with varying wind profile variation detail. The resulting convergence trend supports making well-informed decisions regarding the detail level when trading off the accuracy and complexity of the energy production estimation.

For the derivation of the wind profile climatology and the development of the flight operation model, intensive use is made of data. The wind profile climatology is derived using a data-driven approach applied to wind atlas data. The flight operation model, which enables characterising the power production, is assessed with the support of existing test flight data before using the model to explore performance for previously unseen wind conditions.

Since validation is crucial for ensuring confident energy production estimates, this thesis places considerable emphasis on validating the AWE flight operation model. Lower-fidelity flight operation models are expected to be suitable candidates for estimating long-term energy production since the estimation does not require time-resolved information about the power production of the system. This thesis prioritises determining the accuracy of lower-fidelity models over refining models. The existing test flight data is used for this purpose and to identify potential limitations of the lower-fidelity models.

The final goal of the REACH project [14, 15], which funded this research, was to commercialise a 100 kW flexible-kite system [16]. Kitepower B.V., a leading start-up in AWE technology, led the consortium. Due to the unavailability of flight data of a 100 kW system, the lower-capacity development system of Kitepower is repeatedly investigated throughout this thesis instead. The findings concerning flight behaviour can be generalised to the 100 kW system but do not apply to different types of AWE systems. The close collaboration with Kitepower ensures the practical value of this work.

Thesis outline

In line with the aim of this thesis, the research questions below have been formulated. These research questions are derived in the literature review in Chapter 2, which provides additional context.

How significant is the wind profile variability to the annual energy production estimation? (main research question)

- How does flexible-height operation change the accessible wind resource? (Chapter 3)
- How can the wind profile variability be integrated into a statistical description of the wind climate? (Chapter 4)
- How can the pumping flight operation be modelled efficiently with due regard to the influence of the wind profile variability?
  - How should the aerodynamics of a flexible kite be modelled? (Chapter 5)
  - How should the turning of a flexible kite be modelled? (Chapter 6)
  - How should the flight path be modelled? (Chapter 7)
- How sensitive is the power output to the wind profile? (Chapter 8)

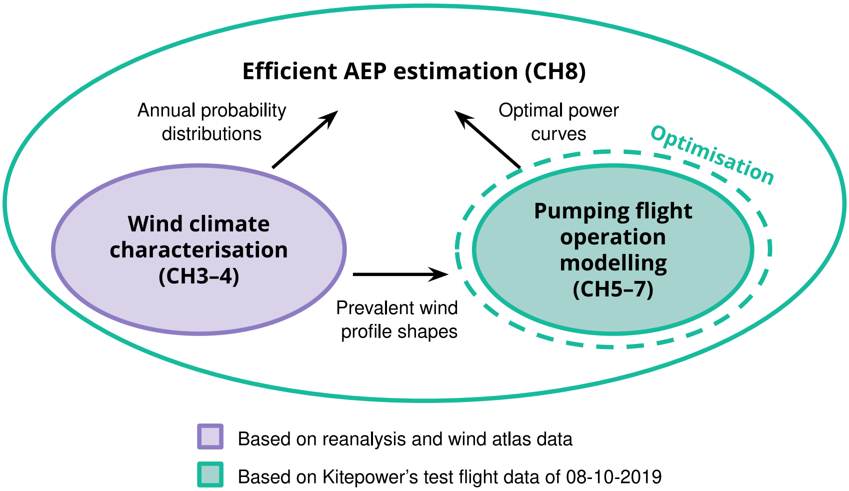

Figure 1.1: Visualisation of the structure of this thesis with the information flow between different parts and their dependency on data. The flight operation model is wrapped in an optimisation layer.

Figure 1.1 illustrates the structure of this thesis, which closely follows the aforementioned research questions. After the literature review in Chapter 2 follows the first part of this thesis with Chapters 3 and 4, focusing purely on the wind resource with an emphasis on the vertical wind profile and building up to the wind climate characterisation. The second part of this thesis comprises the modelling of the flight operation. Chapters 5 and 6 study how to model the pumping flight operation to estimate long-term performance, specifically delving into the aerodynamics and the turning mechanism of the kite. Chapter 7 evaluates the accuracy of the model. Chapter 8 synthesises the findings from both parts of the thesis, integrating the wind profile climatology with the modelling of the pumping flight operation. The energy production estimation framework developed in this chapter applies the model in optimisations to assess how energy production is affected by wind profile variability. Finally, Chapter 9 answers the research questions and discusses future work and practical implications.

2

Literature review

This chapter provides background knowledge and discusses previous works related to the aim of identifying the effect of wind profile variability on energy production. Moreover, it introduces the test flight dataset that has been used repeatedly throughout this thesis. Complementary to the literature presented in this chapter, every subsequent chapter briefly lists relevant literature.

2.1. Introduction

A wide variety of concepts fall under the umbrella of airborne wind energy (AWE). The common denominator is that all employ one or multiple flying devices that are connected to the ground via one or multiple ropes, or tethers, to harness energy from the wind. The flying devices are also referred to as kites because they are tethered to the ground, even though they may appear more like, e.g., a sailplane or a drone.

AWE has several potential advantages over conventional wind turbines due to which the technology has gained interest in the last two decades [7, 17]:

1. Higher altitudes can be reached where the wind is typically stronger and more persistent;
2. A higher power-to-mass ratio and, thus, less material required for the same power capacity;
3. Faster and temporary deployment reduces commissioning costs and enables the servicing of niche markets.

However, a large challenge to AWE technology is safety and reliability. The systems rely on automatic control to keep the devices airborne and require safety-critical decisions in an uncertain wind environment [9]. Consequently, present development platforms have typically demonstrated only a few hours of autonomous flight [7].

The working principle of AWE systems flying cross-wind can be explained by making the analogy to a conventional wind turbine. The kite of an AWE system has the same function as the outer part of the wind turbine blade: transforming the kinetic energy of the wind into usable kinetic energy of the wing (kite or blade). As a result of the work delivered by the air on the kite, the wind is slown down. If the kite flies circles, it would sweep an annulus similar to the area swept by the outer part of the blade, which is the part of the blade that generates the most power. The tether of the AWE system has the function of carrying the load similar to the inner part of the blade and the tower of the wind turbine but is less material intensive.

The cross-wind flight theory published by Loyd [18] in 1980 was arguably the first to demonstrate the potential of AWE systems. Nonetheless, it took until the decade 2000–2010 for the research and development of the technology to intensify.

Loyd derived an expression for the maximum instantaneous power that a kite can generate in cross-wind flight, in which the kite flies approximately perpendicular to the wind at a much higher speed than the wind speed. The derivation covers two types of systems: fly-gen and ground-gen systems. The system depicted in the central panel of Figure 2.1 is a fly-gen system. The small wind turbines mounted on the wing generate electricity which is transmitted to the ground with a conductive tether. The high cross-wind speed of the kite is thus directly utilised for power generation. The right-hand image of Figure 2.1 depicts a ground-gen system, or pumping system, which operates in pumping cycles with two main phases. In the traction phase, the kite pulls the tether from the winch on the ground and, thereby, drives the generator. In the retraction phase, the kite is depowered, and the generator acts as a motor to lower the kite. Since the retraction phase only consumes a fraction of the energy produced during the traction phase, the pumping cycle yields a net positive power output.

Not all AWE systems that are being developed use cross-wind flight, as shown in the classification scheme for AWE systems in Figure 2.2. One exception is the lighter-than-air concept, which is used in the Magnus-effect system of Omnidea and the aerostatic wind turbine of Altaeros. A rotational ground-gen system such as that of Windswept also uses cross-wind flight but not to keep the system airborne. Today, most AWE development is targeted at the ground-gen concept with a fixed ground station but also the development of the other concepts continues.

AWE systems can also be categorised into flexible and rigid kite systems. Flexible kites are made of fabric material and typically have a bridle system to fix the shape of the kite. Examples are the leading-edge inflatable kites of the company Kitepower and the ram-air kite of Skysails. Rigid kites look more similar to sailplanes like the kites of Ampyx, Kitemill, and TwingTec. Enerkite uses a swept semi-rigid wing also supported by a bridle system.

The flight trajectories of the kites of fly-gen and ground-gen AWE systems have different characteristics. Fly-gen systems typically fly a circular path [19]. The area swept by the kite can be approximated by an annulus in a plane which is slightly tilted forward with respect to the cross-wind plane. The major difference with respect to the swept area of a conventional wind turbine is that the radius of the annulus can be controlled. In addition to the cross-wind motion, ground-gen systems also may sweep a large vertical range in which they cover a large part of the atmospheric boundary layer or even rise above it.

The aim of this chapter is two-fold: evaluating what operational aspects and atmo-

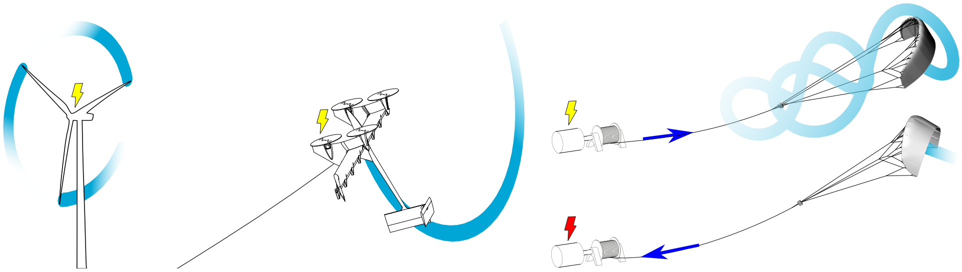

Figure 2.1: An AWE system replaces the tips of a conventional wind turbine (left) with a tethered flying device. Electricity is generated either continuously onboard and transmitted to the ground by a conducting tether (centre), or on the ground, with the tether transmitting the alternating mechanical power of a pumping cycle (right).

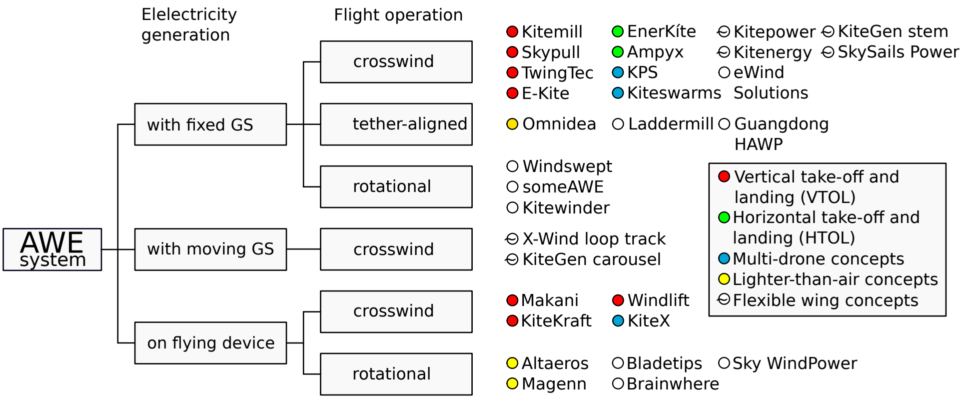

Figure 2.2: Classification of AWE systems with a list of the companies active in 2019 for every class [9].

spheric structures are relevant for the performance modelling of AWE systems and reviewing the state-of-the-art of performance models and related data available. In line with the research project REACH [14, 15], the main focus of this literature review is put on pumping AWE systems. Where relevant, also performance models developed for fly-gen systems are included in the review.

This chapter is organised as follows. First, the pumping cycle operation of ground-gen AWE systems is reviewed. Next, the characteristics of the wind environment are addressed. Subsequently, the existing performance models are reviewed by addressing how detailed the operation and atmospheric aspects are modelled. The chapter is concluded by covering available related data for performance model development consisting of the test flight and wind data.

2.2. Pumping cycle operation

The operation of the pumping cycle of a pumping AWE system greatly affects the energy production. The mean cycle power is an often-used indicator for the energy production of pumping AWE systems:

$$\bar{P}_{\text{cycle}} = \frac{E_{\text{cycle}}}{T_{\text{cycle}}} \quad , \tag{2.1}$$

in which E_{cycle} is the total energy produced in one cycle and T_{cycle} is the cycle duration. To maximise the mean cycle power, the operation of a pumping AWE system should be tailored to the wind conditions. The system is controlled by steering the kite and modulating the torque acting on the winch at the ground station. The winch controller gives control over the tether force which affects the flight of the kite including the reeling speed. The controller should aim at a high and low traction power during the traction and retraction phases, respectively. Simultaneously, the controller should strive to enhance the duration of the traction phase relative to the duration of the retraction phase. Additionally, the mean cycle power can be increased by planning the flight trajectory carefully such that the height range with the highest power potential is exploited by the kite.

The pumping cycle of an AWE system can be divided into multiple phases with different objectives. Figure 2.3 shows the distinct phases of a theoretical pumping cycle typical of a flexible-kite system. During the traction phase, or reel-out phase, the kite flies figure-of-eight manoeuvres in a fast cross-wind motion. After the traction phase follows one of the transition phases, in which the kite stops flying cross-wind and flies towards the side of the wind window, i.e., away from the position directly downwind of the ground station. Once the kite is slowed down, the kite is reeled back in; marking the start of the retraction phase, or reel-in phase. In the pumping cycle illustration, the kite is steered towards the zenith and flies at a high elevation towards the ground station. Another transition phase begins once the kite is pulled back in; the kite steers down and flies towards the starting position of the traction phase to start the next cycle.

Between the different AWE companies, a large variation exist in how they fly a pumping cycle, as illustrated in Figure 2.4. The flight trajectory closely depends on the depower approach of the kite during the retraction phase. The kite can be depowered by flying to the edge of the wind window or lowering the aerodynamic efficiency of the wing. All pumping AWE systems have in common that the crosswind flight is stopped during retraction to depower the kite by reducing its speed. Previously, Skysails only depowered their ram-air kite by flying to the edge of the wind window but recently also started pitching the kite to lower the aerodynamic efficiency [11].

Rigid kite systems, such as the systems of Kitemill and Ampyx, are more effective in lowering the aerodynamic efficiency of the wing than flexible-kite systems. This allows rigid kites to glide down during the retraction phase against the wind, nearly straight towards the ground station, with minimal tether tension. Flexible-kite systems such as the systems of Skysails and Kitepower, require tether tension to maintain the shape of the kite and secure a stable flight. Therefore, they rely more on depowering by flying to the edge of the wind window. This results in more drawn-out retraction phases and, consequently, drawn-out pumping cycles sweeping a large height range with pronounced transitions between the traction and retraction phases [20].

The interruption of the crosswind flight is more pronounced in some pumping cycles than in others. The pumping cycle of Ampyx is kept very compact and the traction and retraction phases are alternated quickly. The pumping cycles resembles a single figure-

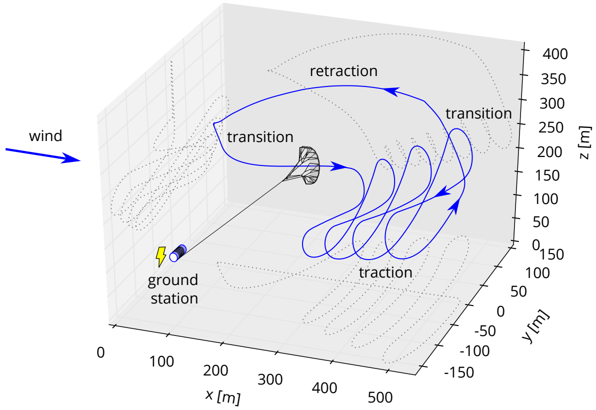

Figure 2.3: The flight path of the flexible-kite, pumping AWE system (kite & drum not to scale) adapted from [21].

of-eight manoeuvre in which the kite is reeled in at the sides of the figure of eight. In contrast, the pumping cycle of Skysails is drawn-out; the flexible kite completes approximately three multiple figure-of-eight crosswind patterns and flies further to the side of the wind window during the retraction phase. Kitepower flies at approximately four figures of eight at a nearly fixed elevation angle (angle with respect to the horizontal) during the traction phase and flies up during the retraction phase. In contrast to the previous examples, the KM1 rigid kite of Kitemill flies circular crosswind patterns. The pumping cycle includes approximately ten crosswind patterns around an elevated path with a positive z-intercept. During the retraction phase, the kite flies approximately horizontally above the crosswind patterns towards the ground station.

2

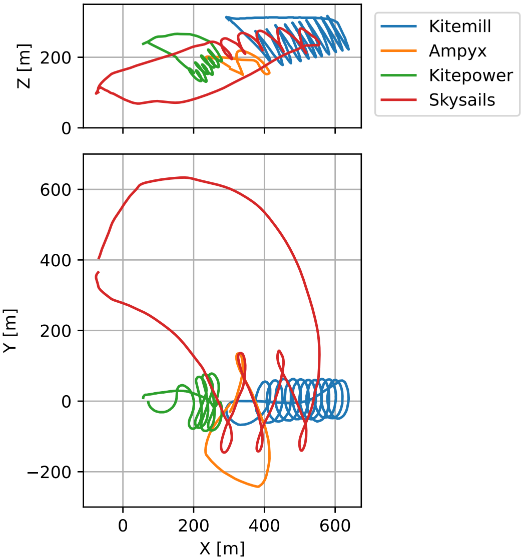

Figure 2.4: Real pumping cycle flight trajectories performed by a kite similar to the SKS PN-14 kite of Skysails [10], the V3.25B kite of Kitepower, the KM1 kite of Kitemill [10, 22, 23], and the AP2 kite of Ampyx [24].

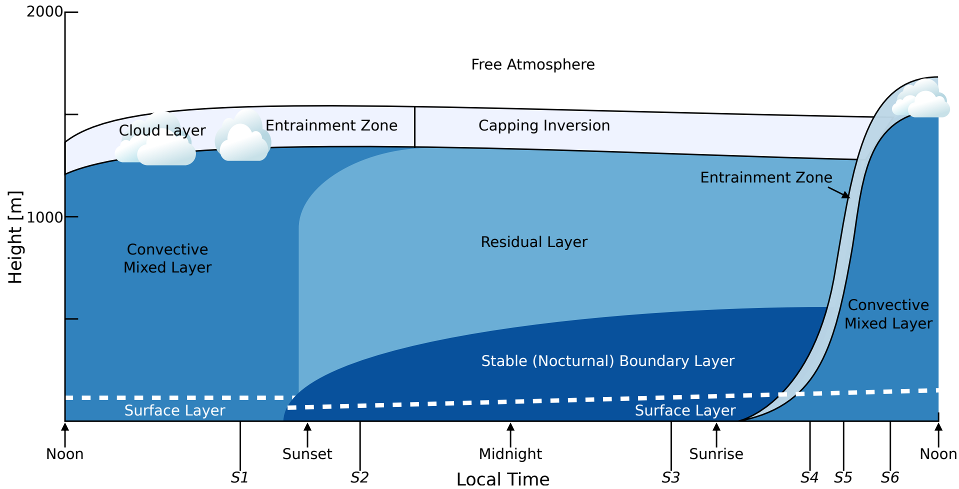

Figure 2.5: Typical extent and evolution of different sub-layers of the atmospheric boundary layer over land adapted from [26].

2.3. Wind conditions

A good understanding of the wind resource is needed to predict the performance of AWE systems. The wind resource is dictated by atmospheric conditions and orography. Therefore, the state of the atmospheric boundary layer will greatly affect the pumping cycle operation. This section addresses the relationship between the atmosphere and wind conditions. Ultimately, this knowledge is employed to determine the operational strategy of the system that maximises energy production.

2.3.1. The atmospheric boundary layer

The atmospheric boundary layer is the lower part of the troposphere that is subject to a daily cycle and in which AWE systems operate predominately. The depth of the boundary layer is typically around 1 km but varies between 200 m to 4 km depending on time and location [25]. An inversion layer with the free atmosphere above it caps the boundary layer. The geostrophic wind in the free atmosphere drives the wind in the boundary layer.

AWE systems are exposed to other sub-layers of the atmospheric boundary layer than conventional tower-based wind turbines because they typically fly well above the height of these turbines. The typical extent and evolution of these layers during the day are illustrated schematically in Figure 2.5. The boundary layer may consist of up to three different sub-layers: a very turbulent mixed layer, which transitions into a less turbulent residual layer, and a growing nocturnal boundary layer, which is only sporadically turbulent [26].

The lowest part of the boundary layer: the surface layer, or the Prandtl layer, is subject to the drag of the Earth's surface, where turbulent fluxes are constant with height. The depth of the surface layer typically equals 10% of the boundary layer depth but varies between 20 m to 200 m depending on time and location [25]. Frequently, a significant

2

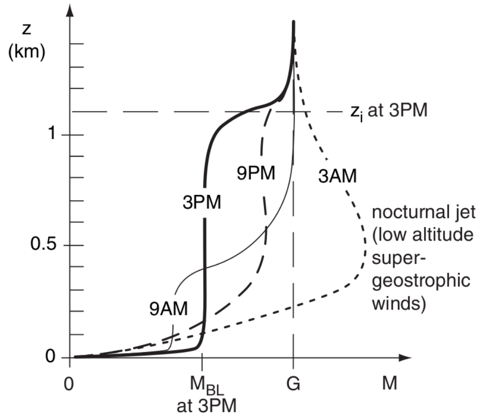

Figure 2.6: Typical wind profile evolution over land. M denotes wind speed and G the geostrophic wind speed. Depicted wind profiles: 3 PM - deep mixed layer, 9 PM - turbulence diminishes, 3 AM - reduced drag accelerates air above surface layer causing a nocturnal low-level jet, and 9 AM - shallow mixed layer after sunrise [25].

portion of the rotor of tower-based wind turbines operates in the surface layer, while AWE systems almost exclusively operate above it. In the Ekman layer, which extends from above the surface layer to the top of the atmospheric boundary layer, the surface wind speed and direction adjust to those of the geostrophic wind.

2.3.2. The vertical wind profile

The vertical wind profile describes how the wind speed changes with height at a certain location. Depending on the application, this information may be supplemented by the change in wind direction with height. Typically, the vertical wind profile describes the mean wind conditions averaged over a period of multiple minutes and varies slowly over the course of hours. The actual instantaneous wind profile will exhibit fluctuations due to turbulence and gusts.

During its daily cycle, the vertical structure of the boundary layer affects the vertical wind profile. Consequently, the wind profile will exhibit a daily cycle. Figure 2.6 depicts how the boundary layer evolution shown in Figure 2.5 relates to the evolution of the vertical wind profile. The four depicted instances are driven by the same geostrophic wind speed but exhibit different profiles in the boundary layer.

In the surface layer, the turbulent fluxes are approximately constant with height and the wind profile (below ~200 m) is found to be dependent on a small set of flow parameters. Consequently, the vertical wind profile can be conveniently parameterised in the surface layer and can be captured using empirical relationships with relatively good accuracy. This is the basis of Monin-Obukhov similarity theory [27] described later in this section. As the flow above the surface layer is affected by more parameters, such a convenient description of the wind profile is not available for the Ekman layer.

The power law is an often-used empirical relationship to estimate the vertical wind profile in the surface layer. It relates the wind speed v_w at one height z_1 to that at a

Table 2.1: Stability classes in terms of the Obukhov length [31].

Class name	Class boundaries [m]	Representative L [m]
Very unstable (VU)	$-200 \leq L < 0$	-100
Unstable (U)	$-500 \leq L < -200$	-350
Neutral (N)	$\lvert L \rvert > 500$	10^{10}
Stable (S)	$200 < L \leq 500$	350
Very stable (VS)	$0 < L \leq 200$	100

different height z_2 and has the form:

$$v_w(z_2) = v_w(z_1)\left(\frac{z_2}{z_1}\right)^{\alpha} \quad , \tag{2.2}$$

where α is an empirical shear exponent factor related to the surface properties. The power law is normally applied up to around 100–200 m [28]. The power law offers insufficient flexibility to describe the variety of measured wind profiles [29].

Another often-used wind profile relationship in the surface layer is the logarithmic wind profile. In its more general form, this relationship differentiates between the effect of mechanical and convective turbulence on wind shear. The basic relationship can be derived based on mixing length theory valid for neutral conditions over a uniform surface. The well-established Monin-Obukhov similarity theory expands its application to non-neutral conditions [27]. The theory can be applied to any stability condition, though it becomes inaccurate under very stable stratification where the surface layer becomes shallow, particularly when attempting to extrapolate above the surface layer [30]. In its non-adiabatic form, the logarithmic profile to estimate the wind speed v_w at height z is given by:

$$v_w(z) = \frac{v_*}{\kappa}\left[\ln\left(\frac{z}{z_0}\right) - \Psi\left(\frac{z}{L}\right)\right] \quad , \tag{2.3}$$

in which v_* is the friction velocity, κ is the von Karman constant, z_0 is the roughness length, Ψ is a stability dependent function, and L is the Obukhov length.

The Obukhov length can be used to evaluate atmospheric stability and is positive for stable stratification, negative for unstable stratification, and infinite for neutral stratification. A classification of stability conditions is presented in Table 2.1. Multiple stability-dependent functions are proposed in the literature. Holtslag et al. [31] propose the following functions:

$$\Psi(L \leq 0) = 2\ln\left(\frac{1+x}{2}\right) + \ln\left(\frac{1+x^2}{2}\right) - 2\arctan(x) + \frac{\pi}{2}, \qquad x = \left(1 - 19.3\frac{z}{L}\right)^{\frac{1}{4}} \quad , \tag{2.4}$$

$$\Psi(L \geq 0) = -6.0\frac{z}{L} \quad . \tag{2.5}$$

Figure 2.7a shows the logarithmic wind profiles for each stability class and a roughness length of 0.1 m representative for farmland with low crops. Moreover, the neutral logarithmic wind profile is shown for a roughness length of 0.0002 m representative for the sea. For the same friction velocity, this profile is shifted horizontally with respect to the neutral logarithmic wind profile with 0.1 m roughness length.

2

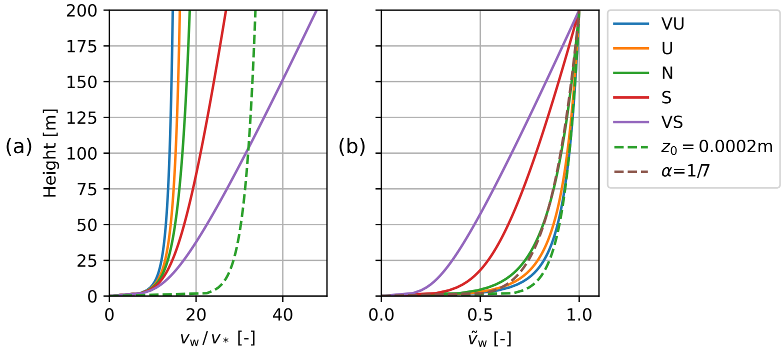

Figure 2.7: (**a**) Logarithmic wind profiles for different atmospheric stability classes (listed in Table 2.1) and roughness length, normalised with the friction velocity. The solid lines depict profiles with roughness length z_0=0.1 m and the dashed line depicts the neutral logarithmic wind profile with roughness length z_0=0.0002 m. (**b**) The same logarithmic wind profiles, normalised with the wind speed at 200 m height, together with the wind profile of the power law ($\alpha = 1/7$).

Figure 2.7b shows the same profiles but with varying friction velocities such that the profile equals one at 200 m height. The logarithmic profile is most often used to extrapolate the wind speed v_w from a height where it is known, z_1, to another height of interest, z_2, in which case the value of the friction velocity is not required:

$$v_w(z_2) = v_w(z_1) \frac{\ln\left(\frac{z_2}{z_0}\right) - \Psi\left(\frac{z_2}{L}\right)}{\ln\left(\frac{z_1}{z_0}\right) - \Psi\left(\frac{z_1}{L}\right)} \quad . \tag{2.6}$$

The logarithmic wind profile in this form allows a more direct comparison with the wind profile power law, which is plotted for a shear exponent factor $\alpha = 1/7$ in Figure 2.7b.

The wind profile relationships illustrated in Figure 2.7 cannot capture the wind profile variation in the full atmospheric boundary layer illustrated in Figure 2.6. To extend the applicability of wind profile relationships to higher altitudes, several modifications have been proposed [32, 33]. However, these theoretical formulations are only validated up to heights relevant to conventional, tower-based wind turbines.

2.3.3. Modelled wind data

Weather models are widely used for forecasting. Multiple times a day, organisations all around the world run weather models and publish the forecast data. Except for the near-ground region, where a surface layer scheme may be used, these models directly calculate the wind profile. The IFS model of the European Centre for Medium-Range Weather Forecasts (ECMWF) applies such a scheme up to approximately 10 m above the surface [34]. The wind profile evolution at the location of an operational AWE system may be extracted from these data and used to decide on the best operational approach.

Weather models can also be used for reanalysis, sometimes also referred to as hindcasting. Reanalysis combines modern weather models with historical observations to obtain the best estimate of the state of the atmosphere during a period in the past covering a large area. These reanalysis data have been used in literature to analyse how much power an AWE system potentially would have produced if it had been operating at a certain location during a certain period, e.g., Olauson[35].

Both global and regional reanalysis data exist. The former typically have a lower spatial and temporal resolution but may cover periods of multiple decades. Such data are a great resource for characterising the long-term wind resource as averaging over such a long period cancels out the year-to-year variability. In a regional reanalysis, global reanalysis data is downscaled with mesoscale models giving finer spatial resolution.

The resolution of reanalysis data is normally insufficient for conducting a siting study. Wind atlas correction methods can be employed to account for the effect of the local terrain which is subscale to mesoscale models. These corrections can be obtained from short-term on-site measurements or microscale models like WAsP [36].

An example of a global reanalysis is ERA5 from the European Centre for Medium-Range Weather Forecasts [37]. ERA5 has an unprecedented resolution and temporal coverage compared to other global reanalysis data. Data are available from 1950 onwards on a 30 km grid and at 137 vertical variable-height levels up to approximately 80 km above ground. A large set of ground-, air- and satellite-based measurements are incorporated in the analysis using data assimilation.

An example of regional reanalysis data is the Dutch Offshore Wind Atlas (DOWA) produced by the Royal Netherlands Meteorological Institute (KNMI) [38]. DOWA is produced by downscaling ERA5 data to a finer-resolution surface grid using their mesoscale weather model HARMONIE-AROME [39]. The downscaled reanalysis is performed for 10 years from 2008 until 2017. Hourly values for temperature, wind speed and direction, pressure and relative humidity are made available on a 217 x 234 grid with 2.5 km spacing and 17 heights between 10 and 600 m.

The higher resolution and non-hydrostatic nature of HARMONIE-AROME lead to an improved representation of the coastline, land surface heterogeneity, and mesoscale circulations, such as a sea breeze. Furthermore, additional observations from the KNMI's network of automated weather stations, satellite retrievals (ASCAT), and aircraft sensors (MODE-S EHS) have been assimilated by the HARMONIE-AROME model.

Kalverla [40] compares both ERA5 and DOWA against wind profile measurements up to 315 m at the offshore location of the met mast IJmuiden in the North Sea. DOWA improves on ERA5 in terms of wind speed, wind shear, and directional accuracy, as well as the representation of anomalous events such as low-level jets. ERA5 underestimates the average wind speed profile and has a roughly constant wind speed bias of -0.5 m/s with height, whereas DOWA is nearly unbiased. As DOWA does not always align well with the measurements, it exhibits a root mean square error of roughly 1.5 m/s.

2.3.4. Wind resource assessment

Wind resource assessments analyse the current wind climate at a certain height over a larger area. The practical use of these assessments is identifying suitable sites for deploying wind energy converters. Alternatively, these assessments may aim to calculate the technical potential, which quantifies how much wind energy can potentially be harvested on a yearly basis over a certain area. This requires making assumptions on the deployment capacity density, high-level performance characteristics, and the availability of land and sea. More fundamental research exists that employs climate models to account for the effect of the energy extraction on the climate [41–44].

Conventional assessments study the wind resource at the hub height of a typical tower-based wind turbine. Such an approach is not suitable for airborne wind energy converters as they can change their harvesting height depending on at what height the wind conditions are favourable. Consequently, wind resource assessments for AWE systems demand a different approach.

A first global assessment of wind power at high altitudes has been performed by Archer and Caldeira [45] based on 28 years of NCEP/DOE reanalysis data. The work analyses the optimal harvesting height below 12,000 m above ground to include the jet streams located well above the atmospheric boundary layer. The main result is the associated wind power density distribution exhibiting substantial increases compared to the distributions at fixed heights. The work does not consider the envisaged operating height limit for AWE systems.

2.4. Performance models

A long-term performance model estimates the power production, e.g., as part of a siting study. This model includes the short-term performance characteristics of a system and the wind resource at the site. A short-term performance model is often applied in optimisations to provide these short-term performance characteristics. This thesis focuses on the pumping flight operation model. The short-term performance model only covers periods up to a few minutes, whereas the long-term performance model typically covers the full lifetime of a system.

2.4.1. Pumping flight operation models

A wide variety of flight operation models for AWE systems exist in the literature. Most models have been developed with a specific system architecture and application in mind. A classification of flight operation models and their application is depicted in Figure 2.8. This review of flight operation models is limited to models for pumping systems.

Only low-fidelity flight operation models are typically labelled as performance models. These models are typically employed in optimisations as part of long-term performance assessments. With an appropriate optimisation approach, the long-term performance assessment can be conducted within a reasonable time. High-fidelity models that include a controller and aerodynamic and structural solvers for the kite are excluded from the classification due to their high computational cost.

Quasi-steady models

The quasi-steady models (QSM) are located at the lower end of the model fidelity scale, as illustrated with the top-to-bottom arrangement representing increased fidelity in Figure 2.8. Whereas dynamic simulations solve the flight path, QSMs rely on a flight path that is for a large part prescribed with a flight path parameterisation. The kite is assumed to transition through (quasi-)steady flight states along the flight path. Consequently, the forces acting on the kite are assumed to be in equilibrium and the speed of the kite can be determined at any point along the path. The flight path is discretised by moving the

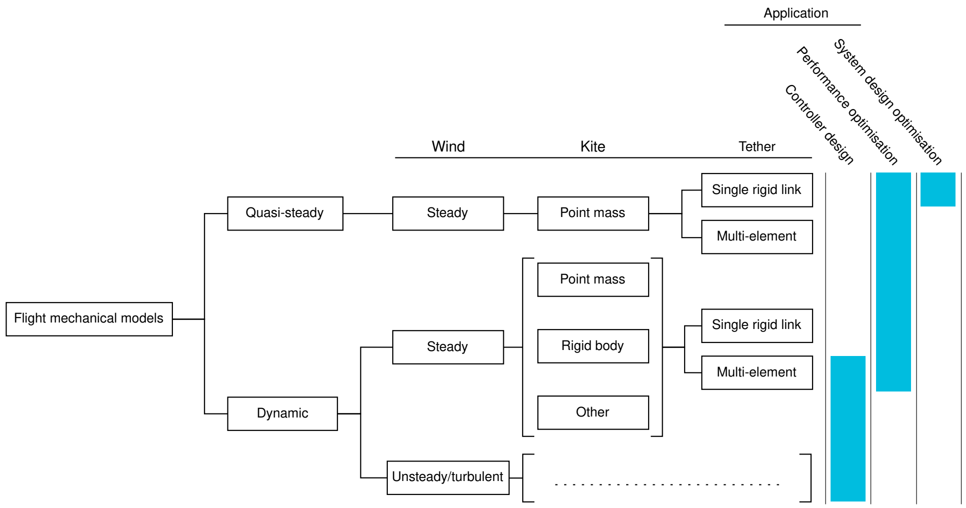

Figure 2.8: Classification of flight mechanical models. The blue bars on the right give a rough indication of which models are employed within three common AWE research topics. The dashed line indicates that the kite and tether models employed in a dynamic framework with steady wind conditions can also be employed with unsteady wind conditions. The top-to-bottom arrangement approximately represents increased model fidelity.

kite along the flight path with the resulting speed for a fixed time step. Alternatively, the flight path can be discretised a priori.

Two QSMs for pumping AWE systems with detailed representations of the pumping cycle flight path are those developed by Van der Vlugt et al. [46] and Ranneberg et al. [47]. The former is developed for flexible kite systems, whereas the latter is developed for (semi-)rigid kite systems.

An important difference between the two QSMs is how they handle the centrifugal force in the force equilibrium during the traction phase. Rigid kites typically fly at a higher speed with associated higher centrifugal force. Therefore, Ranneberg et al. consider the centrifugal force in the force balance. Rolling of the kite is included to tilt the lift force and counteract the centrifugal force. Van der Vlugt et al. neglect the centrifugal force.

The flight path discretisation of the traction phase also differs between the two QSMs. Van der Vlugt et al. do not explicitly consider the cross-wind motion of the kite and solve the radial displacement during the traction phase with a fixed time step. The cross-wind flight is accounted for by a surrogate flight state at representative elevation and azimuth angles that coincides with the middle of a figure-of-eight pattern. Ranneberg et al. collocate points lying on the surface described by a prescribed figure of eight swept along an elevated axis.

Lastly, the retraction phase is solved differently by the two QSMs. Both assume that the kite moves towards the ground station in the downwind-vertical plane. Van der Vlugt et al. solve the radial and elevation displacement of the kite to account for the kite shooting up after the traction phase resulting in an arced flight path typical of flexible kites.

Ranneberg et al. assume that the kite can glide down quickly in a straighter path and only uses a single point to model the retraction phase.

Both system models assume a straight rigid tether and no other QSM exists in the literature that models the system with a multi-element tether. However, the models could in principle be extended with the tether model proposed by Williams [48]. This model is based on a similar quasi-steady assumption and, thereby, is compatible with the system models. Given the position and velocity of the kite, the model determines the tether length based on the tether force or vice versa.

Dynamic models

The dynamic models are found at the higher end of the fidelity scale. Figure 2.8 shows that some overlap exists between dynamic models used for performance analysis and controller design application. Vermillion et al. [9] present a comprehensive overview of models developed for control purposes. The overview also covers dynamic performance models that are employed in flight trajectory optimisation using optimal control techniques to optimise performance. These models typically include a straight rigid tether.

In its simplest form, the kite model of a pumping system includes four degrees of freedom (DOF). Not all of its equations of motion are necessarily dynamic but can be purely kinematic. An often employed kinematic relation is that the kite is aligned with the apparent wind. Fechner et al. [49] introduce an aerodynamic side force as a linear function of the steering input for steering. The side force coefficient is derived from yaw rate measurements from experiments with the turning law, which states that the yaw rate is roughly proportional to the steering input. The angle of attack is inferred from the attitude of the upper tether end. A more physical way of including steering is by controlling the roll of the kite, e.g., by Houska and Diehl [50]. As the kite rolls, the lift vector tilts in the turn and thereby provides the centripetal force. The model assumes that the lift coefficient of the kite can be controlled directly and, therefore, its aerodynamic model does not require calculating the angle of attack.

Rigid kites are often modelled using a rigid body model with six DOF. A computationally efficient formulation of such a model is proposed by Gros and Diehl [51]. The control input typically includes deflections of ailerons, elevator, and rudder; requiring a detailed characterisation of the kite [52]. Malz et al. [24] use the aforementioned model formulation to compare against test flight data of the AP2 kite of Ampyx. Zanon et al. [53] present a generic procedure for obtaining a model for more complex system configurations such as a system with multiple kites or a multi-element tether.

Fechner et al. [49] have proposed a high-DOF model for a flexible kite; three point masses are used to represent the wing and one for the control unit suspended below. The point masses are connected with very stiff spring-damper elements and, thereby, the particle system behaves much like a rigid body. The tether is modelled with lumped masses connected with spring-damper elements.

Turbulence has an adverse effect on the power production [21] which is not captured by performance models. Including this performance loss requires the addition of a controller. The power loss associated with turbulence greatly depends on the performance of the controller.

Flight path optimisation

The full potential of a system can be explored by applying a performance model in optimisation to find the flight path that maximises the power output. The previously presented low-fidelity model can be used for optimisation at a limited computational cost. Applying a QSM for optimisation yields a minimally constrained problem that can be solved with generic optimisation algorithms. On the other hand, applying a dynamic model in optimisation requires a special optimisation technique to effectively explore the objective function while ensuring that the final solution satisfies the equations of motion.

The input to QSMs comprises of high-level variables that characterise the trajectory of the pumping cycle (e.g. length of the reel-out phase) and time-varying control variables (e.g. tether force and roll) that should be aligned with this trajectory. The QSM proposed by Ranneberg et al. [47] is developed particularly for optimisation purposes. The flight path is sparsely discretised to limit the number of optimisation variables while covering the most important flight aspects. Finding the force equilibrium at each flight path point is delegated to the optimiser and ensures that the control variables agree with the trajectory. Another optimisation layer can be wrapped around this core optimisation to optimise the pumping cycle trajectory.

The input to dynamic models consists only of time-varying control variables. These variables can effectively be optimised using an optimisation technique called optimal control. It can be used online to control a real-world system and offline, e.g., for flight path optimisations. One optimisation requires many evaluations of the model to find the optimal control input. The control input does not necessarily contain realistic control actions such as from a flight control computer. Instead, the controlled variables may consist of high-level kinematic quantities such as the kite roll. With the exception of the model of Fechner et al., the aforementioned dynamic models have been used in optimal control problems for optimising the pumping flight trajectory.

2.4.2. System performance characterisation

Characterising the performance of a wind energy converter is a stepping stone towards estimating long-term performance. The power curve expresses the power output against the wind speed at hub height and is the conventional way of characterising the power output of tower-based wind turbines. A typical power curve of an AWE system is depicted in Figure 2.9. Since the power curve expresses the power output as a function of a single variable, it has limited capability of covering the variation in wind conditions.

Ideally, the power output characterisation is site-independent and can thus be used for estimating long-term performance at any location. In practice, the wind profile variability at a site will slightly affect the measured power curve. Especially for large wind turbines, the variations of the wind profile introduce uncertainty to the conventional power curve [54].

No consensus has been reached on what is the best power output characterisation for an AWE system. Often, a simple power curve expressed against the wind speed at a fixed reference height is used. Alternatively, Airborne Wind Europe proposes in its glossary to express the power curve against the average pattern trajectory height [55], as shown in Figure 2.9. Without information about the (varying) operating height range, this type

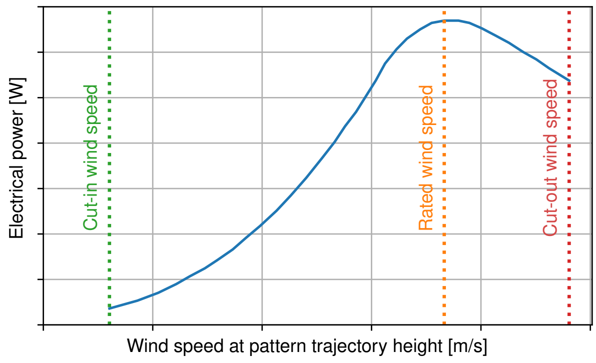

Figure 2.9: AWE system power curve and associated characteristics according to the definitions in the Airborne Wind Europe Glossary [55].

of power curve is of limited use to long-term performance assessments as it cannot be connected straightforwardly to the wind resource.

More than for conventional wind turbines, the power curve of AWE systems will suffer from uncertainty due to wind profile variations. This can be attributed to the larger height range swept by AWE systems. Consequently, AWE systems will encounter more sub-layers and a greater range of turbulent structures in the atmospheric boundary layer associated with its diurnal cycle. Moreover, the kite may even cross the capping inversion and operate in the free atmosphere in case of a shallow boundary layer. These aspects lead to non-uniformities in the wind profile and, thereby, a higher uncertainty of the power curve.

This uncertainty in the characterisation can be mitigated by parameterising the power output with respect to wind profile properties other than the wind speed at a reference height. Despite the uncertainty, a more detailed characterisation is not likely to replace the power curve altogether, e.g., a simple curve is a very powerful tool for wind resource assessments.

In a preliminary step towards assessing the long-term performance of an AWE system at a medium and a high-wind speed location, Licitra et al. [56] generate a power curve expressed against the wind speed at 10 m height using a model-based approach. This requires making some site-dependent assumptions about the wind environment. Licitra et al. assume a logarithmic wind profile with a roughness length of 0.1 m irrespective of the location. The power curve is obtained by solving an optimal control problem for a range of wind profile magnitudes. No cut-in wind speed is used. Instead, the system consumes power below a certain wind speed to fly holding patterns to keep the system airborne.

Ranneberg et al. [47] employ a more generic approach by generating a family of power curves for logarithmic wind profiles with different roughness lengths. The study shows that the power curve is not very sensitive to the roughness length when expressed against the wind speed at the average operating height. Note that this observation is specific to the choice of expressing the logarithmic wind profile as function of the wind speed at the average operating height instead of, e.g., against the friction velocity or the wind speed near the ground. Any other effects to the wind profile shape are not considered, such as the effect of the atmospheric stability. Apart from the sensitivity study, a family of power curves is generated for different average operating heights, producing a look-up table that maps wind speed and height to power output. To calculate the power output, the form of the logarithmic wind profile function is adjusted for each evaluated height, expressing it as a function of the wind speed at that particular height.

Instead of assuming an analytical wind profile, Sommerfeld et al. [57] infer power curves based on a small sample of one year of wind profiles simulated with the Weather Research and Forecasting model with a 10-minute resolution. A targeted sampling approach is used based on the clustering of the simulated wind profiles. The operation of a pumping AWE system is optimised for each of the profiles in the sample. The optimal power outputs are mapped to the mean wind speed within the operating height range and curve fitting is used to infer a power curve. Despite this fitting approach, the power curve shows significant fluctuation, suggesting an insufficient sample size or the requirement for a more detailed power output parameterisation.

2.4.3. Long-term performance estimation

A long-term performance estimation connects the short-term performance characteristics to the wind resource to calculate performance metrics such as the annual energy production (AEP) and capacity factor. When the two are closely coupled, conducting calculations with a flight operation model is typically an integral aspect of the estimation. Otherwise, the short-term performance characteristics may be calculated beforehand and provided with an existing performance characterisation, such as a power curve.

Licitra et al. [56] assess the AEP for two fictitious locations: a medium and a high-wind speed location using a statistical summary of the wind resource as described by the IEC standards. The frequency of occurrence of the wind profile magnitudes at the two locations is assumed to be described by the Rayleigh distribution. The AEP is calculated with the integral of the product of the wind speed distributions at 10 m height and the power curve. Sommerfeld et al. [57] use a similar approach to calculate the AEP at an onshore and offshore location. Instead of using a standardised wind distribution, Sommerfeld et al. infer the statistical wind resource from the simulation results of a mesoscale atmospheric model.

Instead of summarising the wind resource statistically, Ranneberg et al. [47] connect forecasted hourly wind profile data with power curves for five operating heights. For every hourly wind profile, the hypothetical power outputs at the five heights are calculated. Of those five power outputs, only the maximum power output is considered when calculating the AEP. The calculation assumes that the power output is independent of the exact wind profile shape.

Besides avoiding summarising the wind resource, Malz et al. [58] do not characterise the power output of the fly-gen AWE system in a separate step of the assessment. Instead, the maximum power output is determined by solving an optimal control problem for each 3-hourly wind profile in 3 months of global reanalysis data. A site-specific power output characterisation is obtained by arranging the optimal power outputs in descending order. In later work [59], these power duration profiles are obtained for one year of hourly wind profile data. Managing the computations of this analysis can be challenging as they are time-consuming. Conducting this analysis for a pumping AWE system is even more computationally costly as the optimal control problem is larger for these systems.

2.4.4. Validation

The extent to which the previously presented models are validated is very limited. This can partly be attributed to the limited availability of good-quality experimental data. Working with experimental data from a real-scale system is challenging as the system is affected by the continuously varying wind field, which is difficult to measure. Moreover, flight data is often proprietary. An alternative source of experimental data could be scale model experiments. However, experiments with a full-system scale model in a controlled environment, such as a wind tunnel, are sparse due to the technical challenge of controlling the kite in an air flow with kinematic similarity.

Measuring the wind velocity at the ground is of limited use for estimating the wind velocity at the kite. The wind velocity is subject to turbulence and changes with height and can thus cause large temporal and spatial variations. A good solution to measuring the wind profile is using a lidar system next to a system in operation. Such a measurement campaign has been done for the first time by Kitemill and it has just been completed at the time of writing [60]. Alternatively, the wind velocity at the kite can be inferred from the flow and kite velocity measurements. A larger uncertainty is associated with this indirect way of measuring the wind. Moreover, only the wind velocity at the kite can be inferred and not the full wind profile.

To avoid the accumulation of modelling errors over time, Malz et al. [24] applies small changes to the measured kite states and wind velocity using a fitting problem such that the states satisfy the dynamic equations of the model and the differences between a selection of modelled and measured states are minimised. The resulting differences are small and imply a small power output error. Hence, it is concluded that the model does not contain large deficiencies and that the straight tether assumption is justified. However, since minimising the same differences is part of the fitting objective, this may lead to an optimistic accuracy judgement. It is unclear how effective the methodology is at identifying subtle model errors.

The simplified flight mechanics considered in the QSMs increase the gap between the measurements during an actual flight and the model output. Consequently, Van der Vlugt et al. [46] only perform a visual comparison with flight data. More often only computed and measured time-averaged metrics such as the mean cycle power are compared for model validation. Ranneberg et al. [47] compare the power curve calculated with a QSM against measurements and results of a dynamic model. Similarly, Van der Vlugt et al. [61] compare the power curve produced with the QSM of Fechner and Schmehl [62] with measurements. Both studies show large fluctuations of the measured power around the modelled power curves.

Long-term power production measurements would be the ultimate reference for validating AEP estimations. The availability of long-term power production measurements is limited as the existing development systems have only now started to operate more continuously. Moreover, AWE companies are hesitant to publish such data as they contain sensitive information. Also, power curve measurements of the latest development platforms are typically not made public. Power curve validation is essential for proving the long-term performance and economic potential of AWE systems. This type of validation would help to push the technology forward and should get a high priority [63].

2.5. Kitepower B.V.

This research was carried out as part of the REACH consortium together with Kitepower B.V. with the final goal: a commercial 100 kW system [16]. This collaboration between industry and academia aimed at promoting knowledge transfer and intensifying the use of flight data in research.

The development strategy of Kitepower is gaining a lot of practical experience with a relatively low complexity system to rapidly reach commercialisation. Although commercialisation is not yet reached, major steps have been taken towards this goal with an important milestone being the first operation on a tropical island [64], depicted in Figure 2.10. The longest flight lasted 22 hours as the low wind speed conditions did not allow continuous operation.

Figure 2.10: AWE system of Kitepower B.V. with a 60 m^2 kite and ground station integrated in a standard 20 ft container in operation on the Caribbean island Aruba in October 2021 (photo courtesy of Kitepower B.V.).

2.5.1. Test flight

At the start of this research, a test flight was selected in consultation with Kitepower which would serve as the subject of investigation of this research. The main criterion of the selection was the length of the test flight. A three-hour test fight was selected as it provides a significant dataset that is useful for gaining statistical insights into the flight behaviour of the kite. In approximately three hours, the system completed 87 pumping cycles in moderate wind conditions.

The selected test flight took place at the former naval air base Valkenburg in the Netherlands with the ground station positioned at a latitude of 52.1691°, a longitude

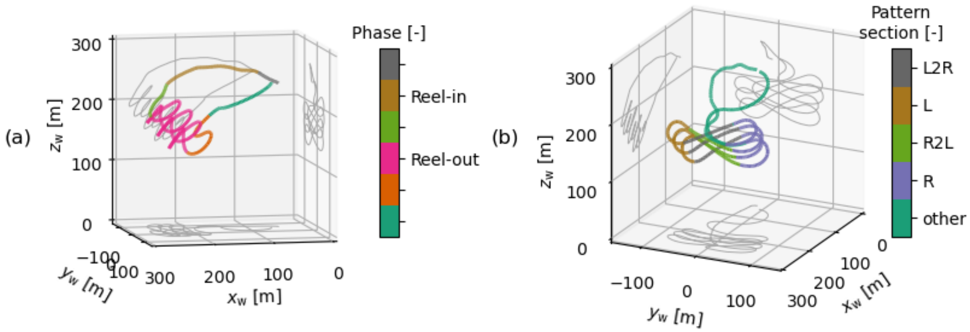

Figure 2.11: The reference pumping cycle of Kitepower's test flight illustrated with the flight path of the kite. The wind blows in the positive x-direction. (**a**) The distinction between phases within the cycle is depicted with the colour scale. The controller treats the orange segment as part of the reel-out phase. The transition phases are the segments lying between the reel-out and reel-in phases. (**b**) The reel-out phase is divided into left and right turns (from a downwind perspective) and the straight sections in between.

of 4.4310°, and an elevation of -5 m (below sea level). It lies close to Leiden and near the coast which is roughly 3.5 km away towards the west. The immediate surroundings are mostly flat open land with some more vegetation towards the dune area at the coast. The wind conditions were moderate with a westerly wind of roughly 5 m/s at the start, which gradually changed to a west-southwesterly wind of roughly 7 m/s measured 6 m above the ground.

Conservative operational settings were used for this specific flight because its purpose was to test new hardware and software components of the system and to acquire data rather than maximising energy production. The 25 m^2 operated kite was substantially smaller and less performant than the 60 m^2 kite shown in Figure 2.10 that Kitepower B.V. develops for the commercial 100 kW system [65]. Considering all this, the power output during the test was substantially lower than for the nominal operation of the commercial system.

Each pumping cycle consists of a traction phase where the kite flies figure-of-eight manoeuvres and a retraction phase with transition phases in between. All cycles have a similar altitude profile and flight path. To outline the flight path, the 65$^{\text{th}}$ pumping cycle, shown in Figure 2.11, is described next. The traction phase fully comes into effect at the lowest altitude around 130 m. In the traction phase, the kite flies three and a half figure-of-eight manoeuvres after which it exceeds 200 m altitude. In the subsequent transition phase and at the start of the reel-in phase, the kite rises up to roughly 270 m after which the kite is lowered by retracting the tether.

2.5.2. System configuration

For this specific experiment, the 100 kW ground station was employed. The rated power is the average power during a cycle. Consequently, the instantaneous power experienced by the generator may exceed this rating, i.e., the power during the reel-out phase is on average approximately 30% higher. The system was equipped with the 25 m^2 leading edge inflatable V3.25B kite. This kite is a derivative of the TU Delft LEI V3 kite designed for the 20 kW demonstrator system and was previously studied in 2017 by Oehler and Schmehl [66]. Figure 2.12 shows a picture of the kite before the flight test and Figure 2.13 shows a schematic drawing of the kite. A thick tether with a 10 mm diameter was used which is compatible with a 60 m^2 kite. A small ram-air turbine provided onboard electrical power, as shown in [67].

The wind at the ground was measured using an anemometer and wind vane mounted to the ground station at 6 m height. The air speed at the kite was measured using a Pitot tube rigidly-mounted to the front bridle lines at the connection to a power line. The apparent wind speed was measured with a Pitot tube attached to the front bridle lines at the connection to a power line. This flow sensor is visible in the foreground of Fig. 2.12, also featuring a flow vane to measure the angle of attack. The side slip angle was not measured in this setup.

Two Pixhawks® were mounted to the wing, i.e., one on each of the two struts adjacent to the symmetry plane of the kite, see Figure 2.13. The Pixhawk® has integrated IMU and GPS sensor, and barometer for recording position and attitude of the kite. The default Kalman filter implementation was used to enhance the quality of the measurements. The position data is based on measurements of sensor 0, which have been processed using the default Kalman filter implementation of the Pixhawk®. The velocity measurements used in this thesis come from the same sensor. However, the velocity is also measured with sensor 0. For an unknown reason, sensor 0 did not measure acceleration. Therefore, the acceleration measurements of sensor 1 are utilised.

The tether force was measured at the pulley at the tether outlet of the ground station. A load cell measured the force on the pulley. This force is corrected for the angle of the tether with respect to the load cell to infer the tether force. The tether reeling speed was inferred from the measured rotational speed of the drum.

Figure 2.12: Fully instrumented V3.25B kite before launch. Courtesy of Kitepower B.V.

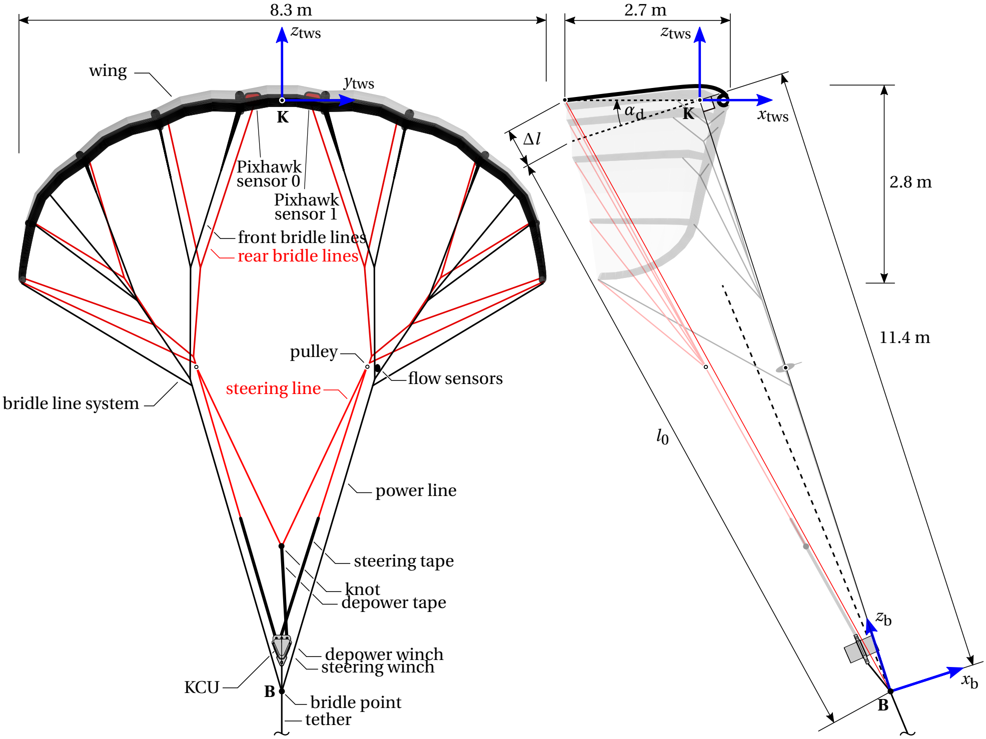

Figure 2.13: Front-view (left) and side-view (right) of the LEI V3 kite. Also depicted are the top wing surface (TWS) reference frame $x_{\text{tws}}, y_{\text{tws}}, z_{\text{tws}}$, with origin **K** at the point around which the wing pitches when changing the angle of attack, and the bridle reference frame $x_{\text{b}}, y_{\text{b}}, z_{\text{b}}$ with origin at the bridle point **B**. The two Pixhawk® sensors 0 and 1 approximately measure in the TWS reference frame while the relative flow approximately measures in the bridle reference frame. Adapted from Oehler and Schmehl [66].

2.6. Research questions

This literature review resulted in the research questions stated below. The conclusions of this chapter are summarised to support these questions.

How does flexible-height operation change the accessible wind resource? (Chapter 3)

The global wind resource accessible to wind energy converters using flexible-height operation has been assessed before. However, previous studies did not consider realistic operating height limits. Consequently, the implications in terms of the annual energy production of realistic AWE systems are unclear.

How can the wind profile variability be integrated into a statistical description of the wind climate? (Chapter 4)

Annual energy production (AEP) estimations for AWE systems are either done based on a univariate statistical wind climate summary or a historical wind profile dataset. The former is efficient but may lead to high uncertainties of the AEP, while the latter is precise but computational costly. The computational effort can be reduced by narrowing the analysed time window. Establishing a statistical description of the wind profile climate based on a historical dataset may provide the best of both worlds.

Existing wind profile relationships may not be a good starting point for expressing the wind profile climate as they are not strictly valid in the full operating height range of AWE systems. This research question has been formulated to investigate alternative approaches.

How can the pumping flight operation be modelled efficiently with due regard to the influence of the wind profile variability? (Chapter 5–7)

Quasi-steady models (QSMs) are efficient and easy to use but may not be as accurate as dynamic models. Multiple studies have shown that power output measurements substantially fluctuate around the power curve modelled with a QSM. More detailed validations of QSMs are needed to identify if such models can be confidently applied in AEP estimations and with what level of detail the flight path should be considered.

The dynamic performance models (without controllers) are relatively complex to use and require more advanced optimisation techniques to find the control output that maximises the power production, which is an integral part of the AEP estimation. The existing models are developed for rigid-kite AWE systems and have not been applied to flexible-kite AWE systems. Consequently, it is unclear how valid the models are for flexible kite systems.

The characterisation of the flight characteristics of the kite is a key aspect of both model types. The accuracy of the estimated energy production will greatly depend on how accurately the aerodynamics and turning of the kite are reflected in the model. The flight behaviour of rigid kites is well understood while the flight behaviour of flexible kites is more complex and not yet fully understood.

The research question is answered using the three sub-questions below:

- How should the aerodynamics of a flexible kite be modelled? (Chapter 5)

- How should the turning of a flexible kite be modelled? (Chapter 6)
- How should the flight path be modelled? (Chapter 7)

How sensitive is the power output to the wind profile? (Chapter 8)

The sensitivity of the power output to the wind profile has not been systematically studied before. One study evaluated the effect of subtly different logarithmic wind profiles to the mean cycle power using a fixed operating height range. Another study only covered the effect of completely different wind profiles to the mean cycle powers. A more systematic approach could provide clarity on this topic.

Understanding the sensitivity of the power output is crucial to understanding how the AEP will change from site to site. Moreover, the sensitivity is decisive for the level of detail with which the wind profile climate should be considered in the AEP estimation. A very low sensitivity would suggest that a simple statistical wind climate summary would yield a sufficiently accurate AEP estimation.

3

Exploring the potential of variable-height operation

This chapter performs a spatial analysis to identify the potential benefits of flexible-height operation in comparison to fixed-height operation across Europe. Locations exhibiting a shift of the probability towards higher wind power densities indicate favourable conditions for airborne wind energy systems. This chapter only considers the characteristics of the wind resource and does not cover the energy conversion process.

3.1. Introduction

Previous wind resource assessments have shown that the wind power available in the atmosphere could theoretically power the world [68]. The precise extent of this power potential is however still a subject of scientific debate. Uncertain is, for example, what effect large-scale energy extraction would have on the overall resource and how the vertical energy exchange between atmospheric layers would influence extraction on this scale. Miller et al. estimate a maximum of 18 to 68 TW that can be extracted over land [41]. Jacobson and Archer also include coastal ocean regions outside Antarctica and revise the saturation wind power potential to 80 TW [44]. Adams and Keith use a mesoscale model and predict that the power potential is significantly lower at 20 TW [69], which is approximately consistent with Miller et al. [43]. Emeis estimates the total extractable wind power potential to be about 61 TW [70]. Dupont et al. review these estimates and conclude that the global wind power potential is substantially lower than previously established when both physical limits and a cut-off value for the energy returned on energy invested (EROI) is greater than 10 [71].

Most of these studies account only for energy extraction close to the surface, using conventional tower-based wind turbines. A first global assessment of wind power at high altitudes has been performed by Archer and Caldeira [45]. The study, based on 28 years of NCEP/DOE reanalysis data, resulted in a global high-altitude wind atlas [72] and was

Parts of this chapter have been published in Renewable Energy **141** (2019).

one of the scientific drivers for the exploration of airborne wind energy. As part of the study, the optimal harvesting height was determined, and the effect of intermittency was investigated as well as the global climate effects of large-scale energy extraction from higher atmospheric layers. Using a climate model, Miller et al. estimated the maximum sustainable extraction from jet streams of the global atmosphere to be about 7.5 TW [73]. The work concluded that jet stream wind power does not have the potential to become a significant source of renewable energy. Using a similar approach, Marvel et al. found that tower-based wind turbines could extract at least 400 TW globally, whereas high-altitude wind power could extract more than 1800 TW [42]. They further state that uniformly distributed wind turbines generating the entire global primary power demand of 18 TW are unlikely to affect the climate substantially. Archer et al. show that also jet-like structures below 3 km above ground (below the jet streams) offer a great potential to AWE systems [74] and find that more than 7.5 TW technical wind power potential is contained in these jet-like structures.

Gambier et al. present a detailed modelling framework for AWE system designs and combine this with COSMO-EU and NCEP/DOE model data for 12 locations in and around Germany [76] as well as lidar measurements up to 1200 m at two locations in Germany [75]. The measurements reveal strong wind shear between 200 and 1000 m altitude during night time, while during the day the wind shear is small. Lunney et al. present a techno-economic study of airborne wind energy harvesting in Ireland [77]. The high-altitude wind resource was modelled on the basis of NCEP/DOE AMIP-II Reanalysis (R-2) data which provides an updated 6-hourly global analysis of atmospheric variables such as wind and temperature with 143×73 grid points in the horizontal with a spacing of 2.5° ranging from the year 1979 to the present. Yip et al. use MERRA-2 data to identify possible deployment areas for AWE in the Middle East, computing also the optimal height at which the systems would operate [78]. Emeis discusses AWE systems in his outlook chapter [70].

Olauson accurately calculates the wind power generation of several countries and regions [35] using ERA5 reanalysis data [37]. The quality of these predictions and the global availability of an unprecedented spatial and temporal resolution has motivated us to use ERA5 data to do a similar study as done by Archer and Caldeira [45].

In this study, the wind resource available to conventional wind turbines and AWE systems is compared without considering the energy conversion process. A fixed harvesting height is assumed for wind turbines, and the harvesting height for AWE systems is assumed to adjust continuously to the varying wind profile. Instead of only evaluating mean values, the distribution of the wind power density is investigated spatially. This allows an investigation of power production intermittency which is an important aspect of a power system with a high penetration of renewables. In contrast to Archer and Caldeira [45], this analysis is restricted to height ranges compatible with the technology expected for the first implementations of AWE. In contrast to Mann et al. [79], the wind resource is analysed over a large area.

3.2. Obtaining the wind statistics

This study uses 11 years of ERA5 data between 2010 and 2020 covering Western and Central Europe on a 0.25° by 0.25° grid (approximately 30 km) to evaluate the wind resource.

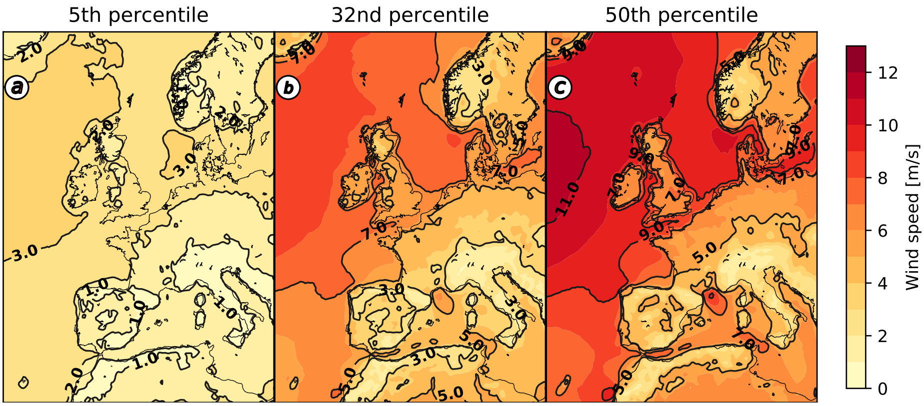

Figure 3.1: 5th, 32nd, and 50thpercentiles of the wind speed probability distribution at a fixed height of 100 m.

The atmospheric variables, including the wind velocity, are provided at model levels which vary in height and pressure. ERA5 uses in total 137 model levels to resolve the state of the atmosphere. The highest model level reaches up to approximately 80 km, covering most of the depth of the atmosphere. In this study, only the lowest 25 levels are used. On average, these levels reach up to 1.6 km above ground. However, depending on the state of the atmosphere it may be hundreds of meters more or less.

The height above the ground of the model levels is calculated for every grid point in time and space. This is inferred from the temperature, humidity, and surface pressure [37]. In an intermediate step, the pressure at each model level is calculated using the surface pressure and the model level definition. From the pressure, the air density ρ can be calculated using the ideal gas law. The power density is calculated from the air density and wind velocity v_{w}:

$$P_{\mathrm{w}} = \frac{1}{2}\rho v_{\mathrm{w}}^3. \tag{3.1}$$

To study the wind climate, the 5th, 32nd, and 50thpercentile wind speeds are evaluated in the spatial domain. Effectively, only the part of the distribution within the cut-in and cut-out limits of the converter contributes to the energy yield. Therefore, a large fraction of the area of the distribution contained within these limits is desirable. The 5th, 32nd, and 50thpercentiles are also employed by Archer and Caldeira [45].

Firstly, these percentiles are determined for the wind speed and power at a 100 m fixed-height in the ERA5 data, shown in Figures 3.1 and 3.2. This is a typical hub height of a modern tower-based wind turbine and, therefore, this height is often used for performing wind resource assessments. For this reason, the corresponding fixed-height harvesting results are used as a reference in later comparisons with variable-height harvesting results.

All three wind speed percentiles in Figure 3.1 show relatively high wind speeds over the Atlantic Ocean west of the United Kingdom and the lowest is found south of the Alps. The contour lines are roughly oriented parallel to the coastline with higher winds offshore than onshore. Moreover, the distributions show some concentrated peaks, e.g., in the Mediterranean off the coast of France.

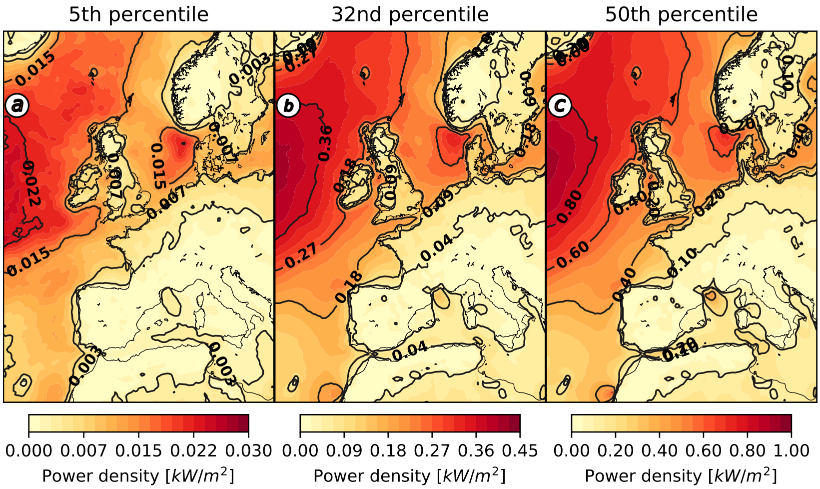

Figure 3.2: 5$^{\text{th}}$, 32$^{\text{nd}}$, and 50$^{\text{th}}$percentiles of the wind power density probability distribution at a fixed height of 100 m.

The limited height ranges covered in conventional wind resource analyses give a poor representation of the resource accessible by AWE systems. Therefore, this study does not only assess the wind resource for fixed-height harvesting but also for variable-height harvesting. The variable-height harvesting analysis accounts for two main features of AWE technology:

- the possibility of accessing higher altitudes than tower-based turbines, and
- the ability to continuously adjust the harvesting height to the varying wind conditions.

The wind resource statistics relevant to AWE are acquired by searching the maximum wind speed of the vertical wind profile for each grid point in time and space in the ERA5 data. The wind profile is examined between 50 m and 500 m height because existing development AWE systems typically adopt a maximum height of 500 m [46, 61, 80]. Although AWE systems may access even higher altitudes, this may penalise the performance and may not be permitted due to air space legislation [81]. Note that, the conversion process and thus the performance penalty is disregarded.

Resolving the flight trajectory of the kite is out of the scope of this analysis. Instead, it is assumed that the AWE system operates continuously at the maximum accessible wind speed. This assumption may be reasonable for fly-gen systems but not necessarily for ground-gen systems as they may cover a large height range depending on the operational approach. Moreover, it is assumed that the AWE system can instantly adjust its operational height. Therefore, irrespective of the step size, instant adjustments between hourly-determined optimal harvesting heights are allowed.

The software used to compile the presented results is implemented in Python and can be downloaded from the publicly accessible repository [82]. An archived version

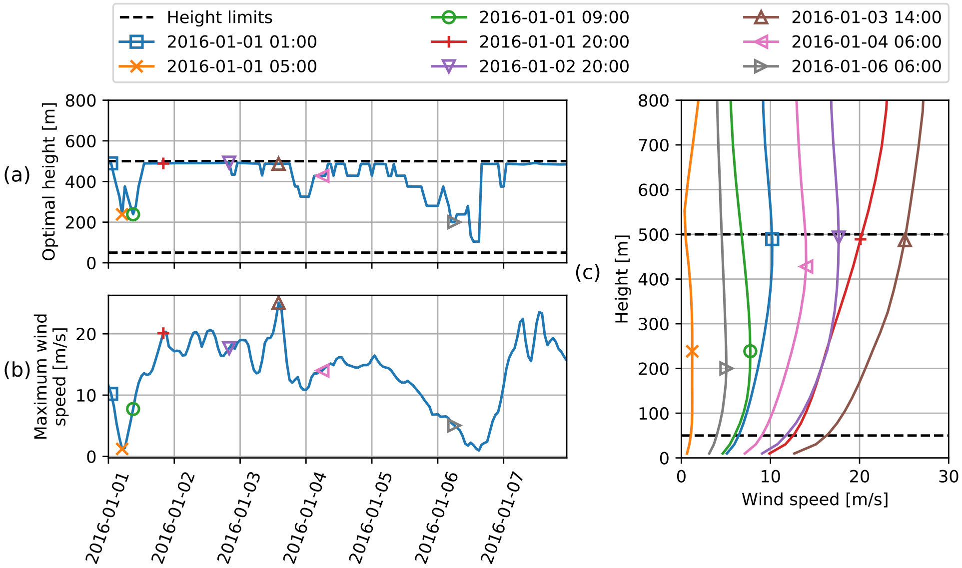

Figure 3.3: Optimal height analysis at the location 51.0°N, 1.0°E in the English Channel during the first week of 2016. (**a**) Optimal harvesting height over time, (**b**) corresponding wind speed, and (**c**) wind profiles at marked instances.

of the source code is provided that is packaged together with the original datasets used for this analysis [83]. Also, instructions are provided on how to download the required ERA5 dataset and how to run the scripts. This will allow future researchers to compile detailed wind statistics at any location in the world. These can be used for analysing the suitability of specific deployment locations.

3.3. Wind statistics over the English Channel

This section illustrates the previously introduced analysis for a location in the English Channel and introduces the metrics that are used in the next section to analyse the wind resource over a larger area. This offshore site is selected as the wind conditions are expected to be favourable near the surface. Thereby, this site is expected to yield a high energy production without variable-height harvesting and give a conservative image of the benefits of variable-height harvesting.

Figure 3.3a and b show the evolution of the maximum available wind speed and the corresponding harvesting height for the first week of 2016. The figure shows that the optimal height frequently coincides with the ceiling of the operational range. A selection of the underlying wind profiles is displayed in Figure 3.3c. The markers with the same colour connect the points in the evolution to the wind profiles. The diversity in wind profiles is considerable during this particular week, including a weak low-level jet on 2016-01-02 at 20:00.

Figure 3.4 summarises the results of the analysis with the entire 11 year of ERA5 data. Figure 3.4a also confirms that the optimal height frequently coincides with the ceiling

Table 3.1: 5th, 32nd, and 50th percentile wind speeds corresponding to the lines in Figures 3.4c and 3.4d

Fixed height	Height [m]	5th percentile [m/s]	32nd percentile [m/s]	50th percentile [m/s]
	100	2.2	5.9	7.6
	500	2.4	6.8	9.3
	1500	2.7	7.2	9.6
Variable height	Ceiling height [m]	5th percentile [m/s]	32nd percentile [m/s]	50th percentile [m/s]
	300	2.6	6.9	9.1
	500	2.8	7.1	9.5
	1000	3.5	7.7	10.1
	1500	4.1	8.4	10.7

of the operational range and rarely occurs at the lower end. More often than not the optimal harvesting height remains approximately the same from one hour to the next. The changes in harvesting height that do occur are predominately small, as shown in Figure 3.4b. This gradual variation justifies the assumption that the harvesting operation can be adjusted instantly.

Figures 3.4c–e shows the resulting wind speed distributions of multiple analyses with slightly different operational limits. To facilitate a visual comparison between the shapes of the distributions, Weibull probability density functions are fitted to the discrete wind speed distributions:

$$f_{\mathrm{Weibull}} = \frac{k}{\lambda}\left(\frac{v_{\mathrm{w}} - v_{\mathrm{w0}}}{\lambda}\right)^{k-1} e^{-[(v_{\mathrm{w}}-v_{\mathrm{w0}})/\lambda]^k}. \tag{3.2}$$

The fitted parameters are the shape parameter k, the scale parameter λ, and the minimum wind speed v_{w0}. The latter is included to emphasise distribution differences at the lowest wind speeds.

Figure 3.4c shows the distributions for different fixed-height cases and the variable-height case with a 500 m ceiling. The 100 m fixed-height distribution is an apparent outlier. Compared to the other distributions, it is clearly shifted towards lower wind speeds. This is due to the effect of wind shear which is greatest in the surface layer leading to significantly lower wind speeds compared to aloft. The effect of wind shear can also be observed in the wind profiles displayed in Figure 3.3c. The other distributions, corresponding to higher harvesting heights, are very similar. Compared to the 500 m fixed-height distribution, the 1500 m distribution is shifted only slightly towards higher speeds. The 500 m ceiling case avoids weaker winds by adjusting the height and virtually never experiences zero wind. The distribution exhibits a greater than zero minimum wind speed, and its centroid is shifted farthest towards higher wind speed, i.e., the distribution exhibits the highest mean wind speed. Both these distribution properties are expected to increase the energy yield.

Alongside the Weibull functions, the diagrams also include the 5th, 32nd, and 50th percentiles represented by square, cross, and circle markers, respectively. These percentile values are also listed in Table 3.1. All distributions have their maximum between the 32nd and 50th percentiles (markers × and ∘, respectively) which is the range of the most probable wind speeds.

To compare the 500 m ceiling and reference (100 m fixed-height) cases, the ratio f_{p_n}

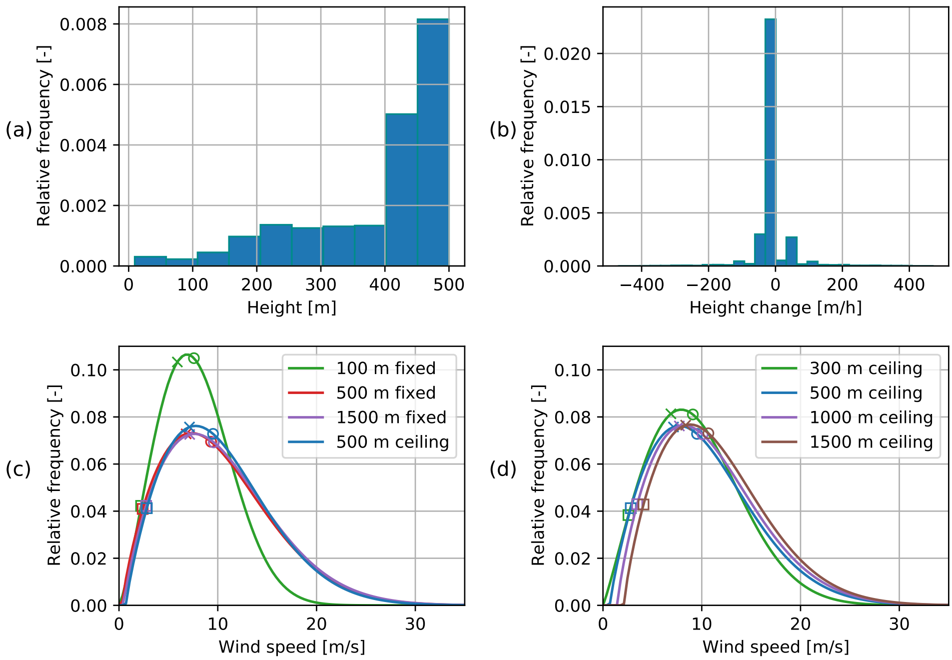

Figure 3.4: Resulting statistics of the optimal height at the location 51.0°N, 1.0°E in the English Channel for the full data set. (**a**) Probability distribution of optimal harvesting height. (**b**) Probability distribution of hourly change in optimal harvesting height. (**c**, **d**) Weibull distribution fits for the fixed-height and variable-height harvesting cases. □, ×, and ∘ refer to the 5$^{\text{th}}$, 32$^{\text{nd}}$, and 50$^{\text{th}}$percentile wind speeds listed in Table 3.1.

between the n^{th}-percentiles p_n of the wind speed (or wind power density) of the two cases is introduced:

$$\mathrm{f}_{p_n} = \frac{p_n(v_{\mathrm{w,500m\ ceiling}})}{p_n(v_{\mathrm{w,100m\ fixed}})}, \tag{3.3}$$

which is referred to as the increase factor. The increase factor is included in the comparison in Table 3.2. Even though this factor is highest for the 5$^{\text{th}}$percentile, the corresponding absolute increase is lowest. The increase factor is lowest for the 32$^{\text{nd}}$percentile, and the absolute increase is highest for the 50$^{\text{th}}$percentile.

The centroid of the probability distribution is shifted towards substantially higher wind speeds when harvesting above 100 m. The shift is particularly large when increasing the fixed height from 100 to 500 m. The shift is smaller when switching from the 500 m fixed-height case to the 500 m ceiling case. In conclusion, allowing access to heights greater than 100 m yields a relatively high increase in the mean wind speed. This increase is expected to be even higher for onshore sites which typically have a higher wind shear. This observation emphasises the potential of harvesting energy beyond the reach of conventional tower-based wind turbines.

The 5$^{\text{th}}$percentile of the wind speed distribution gives an indication of the consistency of the wind, i.e., a high wind speed suggests a high consistency. Table 3.1 shows that the 5$^{\text{th}}$percentile wind speed of the 500 m fixed-height case is only 0.2 m/s higher

Table 3.2: Comparison of the percentiles of the 100 m fixed-height case and 500 m ceiling case.

Percentile	Fixed-height case [m/s]	Variable-height case [m/s]	Absolute increase [m/s]	Relative increase factor [-]
5th	2.2	2.8	0.6	1.26
32nd	5.9	7.1	1.2	1.21
50th	7.6	9.5	1.9	1.26

than the 100 m fixed-height case. The difference is 0.4 m/s between the 500 m fixed-height case and the 500 m ceiling case. This emphasises that the ability to continuously adjust the harvesting height to the varying wind conditions is paramount to access more consistent winds.

Figure 3.4d shows the differences between the distributions of the variable-height harvesting cases with different ceiling heights. The most pronounced effect of increasing the ceiling is the increase in minimum wind speed leading to a higher 5thpercentile wind speed and thus higher consistency of the wind.

3.4. Wind statistics over Europe

This section analyses the wind resource over a spatial grid covering Western and Central Europe. In Section 3.4.1, the 500 m ceiling case is compared to the reference case using contour plots of the 5th, 32nd, and 50thpercentiles. Section 3.4.2 investigates how the accessible wind power density changes with alternative ceiling heights up to 1250 m.

3.4.1. Variable-height operation up to 500 m

The effect of variable-height harvesting on the wind speed distribution is shown in Figure 3.5. The contour plots of the wind speed percentiles (Figure 3.5a–c) show similar trends as the reference case, i.e., the contour lines are roughly oriented parallel to the coastline. However, the values have changed as can be observed in the contour plots of the increase factors (Figure 3.5d–f).

The increase factors are the smallest for the 32ndpercentile wind speed, as was previously observed for the location in the English Channel. Above the continent, the coastal areas, and the Mediterranean Sea, the increase is more than 10%, whereas the increase above the Atlantic Ocean is less significant. The 5thpercentile wind speed shows the highest increase. Above coastal areas a 30% increase is common. Over predominately mountainous areas the increase exceeds 70%. Peaks in the increase factor of more than 2.5 are found, e.g., south of the Alps. Note that the high relative increase at these sites is mostly explained by low reference values (see Figure 3.1).

The wind power density stated in Equation 3.1 is a measure of the wind power that is locally available for conversion. Figure 3.6 shows the effect of variable-height harvesting on the distribution of the wind power density with the contour plots of the 5th, 32nd, and 50thwind power density percentiles and the corresponding increase factors relative to the reference case. Note that the colour scales differ per panel. Since the wind power density is a cubic function of the wind speed, the plotted percentile scores and increase factors are approximately the cubes of those for the wind speed. Therefore, peaks and

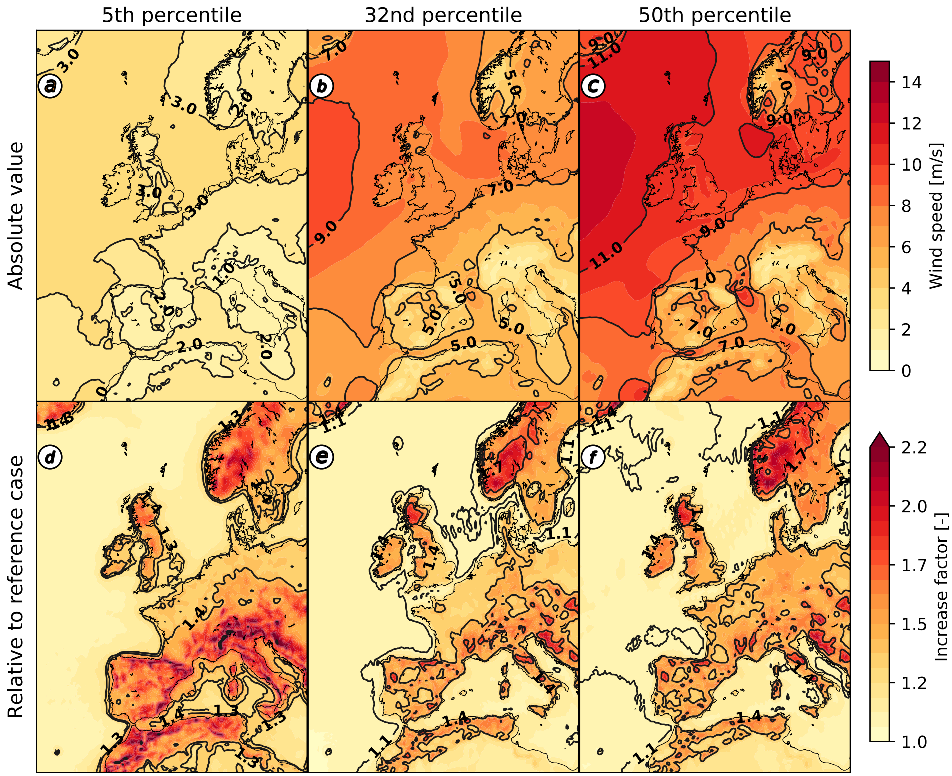

Figure 3.5: 5^{th}, 32^{nd}, and 50^{th}percentiles of the wind speed probability distribution for the 500 m ceiling case (**a**–**c**) and the relative increase with respect to the 100 m fixed-height case (**d**–**f**).

valleys of the power density percentiles coincide with those of the wind speed. Again, the 5^{th}wind power density percentile shows a higher increase factor than the $32^{nd}$and 50^{th}percentile.

As an alternative to assessing the wind power availability with the 5^{th}wind power density percentile, as done by Archer [84], the availability A is assessed using the percentile ranks PR of the wind power densities of 40, 300, and 1600 W/m^2 (equivalent to approximately 4, 8, and 14 m/s wind speed):

$$A = 100\% - PR. \tag{3.4}$$

The 40 and 1600 W/m^2 power densities correspond to typical cut-in and rated wind speeds of conventional wind turbines, respectively. The variable-height harvesting causes an increase in the availability of the cut-in wind speed and, thereby may increase the operational time of the system. Also, the availability of the rated wind speed is increased, leading to an increase in energy yield. Additionally, the operating height could be adjusted such that not only low wind speeds are avoided but also high wind speeds. Although out of scope, such an approach enables further tailoring of the wind speed probability distribution to optimise the energy yield of an AWE system.

Figure 3.7a–c shows the contour plots of the availabilities corresponding to the three power densities. The peaks and valleys are found at similar locations as for the wind power density percentiles. For 40 W/m^2 wind power density, the 90% availability contour line approximately follows the coastline of Northern Europe. Apparent is the low availability, below 50%, in the south of the Alps. At the coastline of Northern Europe, the availability of 1600 W/m^2 wind power density is approximately 22%.

Figure 3.7d–f depict how the availability changes with respect to the fixed-height reference case:

$$\Delta A = A_{\text{500 m ceiling}} - A_{\text{100 m fixed}}\,. \tag{3.5}$$

Around the North Sea, the increase in availability for 40 W/m^2 wind power density is small due to the high reference availability at a fixed height of 100 m. In contrast, most of the Mediterranean coastal areas exhibit a high availability increase. The availability increase for 1600 W/m^2 shows an opposite trend. The increase is highest for the United Kingdom, Denmark, and southern Sweden. These countries also show a fair score on the availability increase for 40 W/m^2.

The availability increase for 40 W/m^2 wind power density associated with variable-height harvesting indicates that the fraction of the time without power production decreases. This suggests that AWE systems have improved base load capabilities compared to conventional wind turbines. Not only in Mediterranean coastal areas, where the highest availability increase is observed, but over most of Europe the increase is substantial. For instance, over the coastline of Northern Europe the availability is increased by 5%, which is significant considering the already high reference availability of approximately 80% for conventional wind technology.

3.4.2. Variable-height operation with alternative ceilings

So far, only the variable-height case with a harvesting ceiling of 500 m has been discussed. In the next step, the effect of the ceiling height on the availability is investigated by repeating the 40 W/m^2 wind power density availability calculation for harvesting ceilings of 300, 1000, and 1250 m. In this section, the availability increase is defined relative to the 500 m ceiling case instead of the 100 m fixed-height case:

$$\Delta A = A_{\text{alternative ceiling}} - A_{\text{500 m ceiling}}\,. \tag{3.6}$$

Figure 3.8a–c show the resulting 40 W/m^2 wind power density availability for the three alternative ceilings. For a 300 m ceiling, the 90% availability contour line is roughly parallel to the coastline of Northern Europe. The area for which 40 W/m^2 has an availability of at least 90% expands by increasing the ceiling. Again, the lowest availability is observed in the south of the Alps, approximately 20%, 40%, and 50% for ceilings of 300, 1000, and 1250 m, respectively (below the colour scale range). The highest availability can be found west of the United Kingdom peaking at 94%, 96%, and 97% for ceilings of 300, 1000, and 1250 m, respectively.

Figure 3.8d–f show the availability increase for the three alternative ceilings with respect to the 500 m ceiling case. The highest increase in wind power availability can be found south of the Alps and in the Mediterranean and Norwegian coastal areas. Note that these areas exhibit a low baseline availability and thus allow more room for improving the availability. The availability increase of 40 W/m^2 wind power density associated with switching from fixed-height to variable-height harvesting (Figure 3.7d) is similar to the availability increase due to increasing the ceiling (Figure 3.8d–f). This similarity suggests that both changes to the operational approach have a similar effect on the availability. This does not necessarily hold for the availability of 300 and 1600 W/m^2 wind power density.

The ceiling in the variable-height analysis significantly affects the availability of 40 W/m^2 wind power density, e.g., at the centre of France and Germany the 300 m ceiling decreases the availability with 2.3% compared to the 500 m ceiling case and the 1000 m and 1250 m ceilings increase the availability with 4.3%, and 6.4%, respectively. Assuming that harvesting low winds at greater heights is feasible, these increases in the order of magnitude of percentages potentially could cause a significant reduction of power production intermittency. However, to enable drawing a more definite conclusion, the energy conversion process needs to be considered in the analysis.

3.5. Conclusion

The available wind resources have been assessed over a large part of Europe using the recently released ERA5 reanalysis data for a period of 11 years at a spatial resolution of 0.25° using the lower 25 model levels, which on average reach up to 1.6 km above ground. The analysis is focused on the paradigm of airborne wind energy (AWE): harvesting at higher altitudes where winds are generally stronger and continuously adjusting the operating height below a predefined ceiling. For the first envisaged AWE systems, a ceiling of 500 m is assumed. The operational details, conversion efficiency, and economic boundary conditions vary strongly between different AWE concepts because they are optimised for different conditions and applications. Therefore, in this study, only the accessible wind power density is assessed and no account is taken of the specifics of the energy conversion.

The effect of variable-height harvesting on the wind speed probability distribution is demonstrated for a location in the English Channel. The analysis illustrates the potential for obtaining access to stronger winds by harvesting energy beyond the reach of conventional tower-based wind turbines. The 5$^{\text{th}}$percentile of the wind speed increases by 0.6 m/s for a 500 m ceiling suggesting a significant increase in the consistency of the wind. The contribution to the wind consistency increase of enabling access to higher altitudes is less significant than that of continuously adjusting the harvesting operation to the varying wind conditions.

The effect of variable-height harvesting with a 500 m ceiling is analysed in the spatial domain by comparing it to the wind resource at a fixed height of 100 m representative of a typical hub height of a wind turbine. First, the increases in the 5$^{\text{th}}$, 32$^{\text{nd}}$, and 50$^{\text{th}}$percentile wind speeds are investigated. The increase of the 5$^{\text{th}}$percentile exceeds 30% over most of the continent, and the highest increases are found over mountainous areas. The increase is less substantial for the 32$^{\text{nd}}$and 50$^{\text{th}}$percentile wind speeds.

The spatial analysis is repeated for the wind power density as it can be more directly related to the power production. The 5$^{\text{th}}$percentile wind power density increases by more than 100% over most of Europe compared to the 100 m fixed-height case.

Additionally, the availabilities of the wind power densities of 40, 300, and 1600 W/m^2 are analysed. The availability of 40 W/m^2 reflects the percentage of time for which a typical cut-in wind speed of a wind turbine is exceeded and provides an intuitive metric for assessing energy production intermittency. Around the North Sea, the increase in this availability due to variable-height harvesting is small due to the high reference availability at a fixed height of 100 m. Nonetheless, it may still give AWE systems better base load capabilities than conventional tower-based wind turbines. The Mediterranean coastal areas exhibit a higher potential for improving the base load capabilities using variable-height harvesting.

Finally, the effect of the ceiling heights of 300, 1000, and 1250 m above ground on the availability is studied. Again, around the North Sea, the gain in availability by increasing the ceiling height is small. On the other hand, the availability increase over Mediterranean coastal areas by increasing the ceiling height from 500 to 1250 m commonly exceeds 10 %.

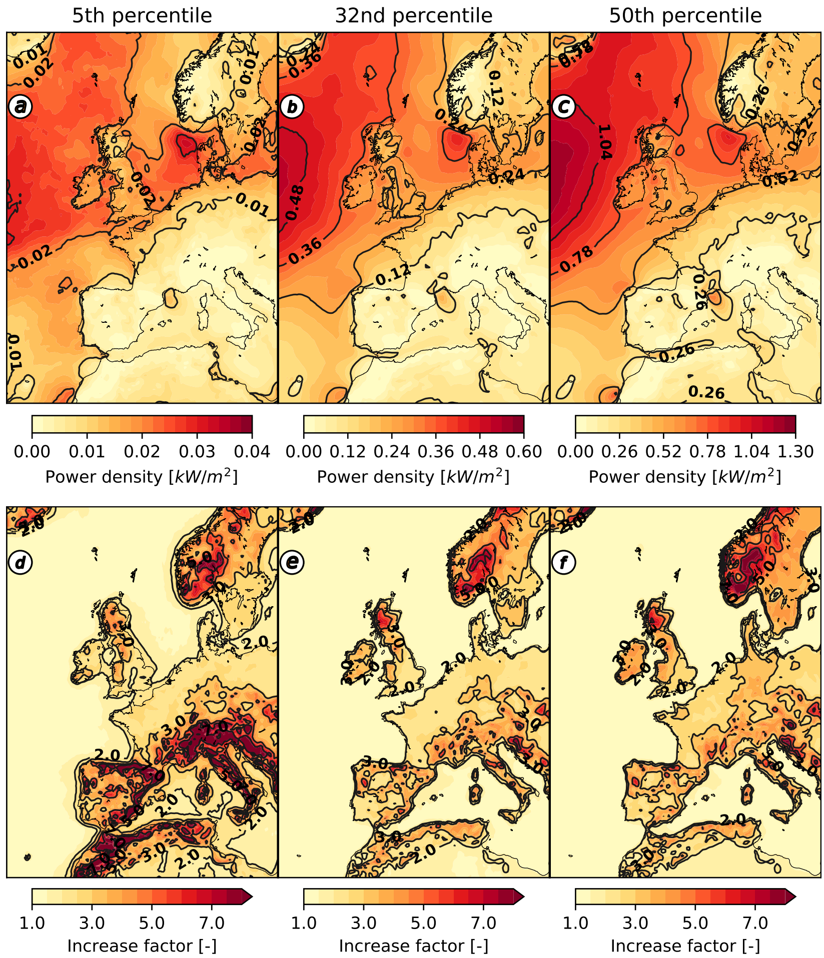

Figure 3.6: 5$^{\text{th}}$, 32$^{\text{nd}}$, and 50$^{\text{th}}$percentiles of the wind power density probability distribution for the 500 m ceiling case (**a**–**c**) and the relative increase with respect to the 100 m fixed-height case (**d**–**f**).

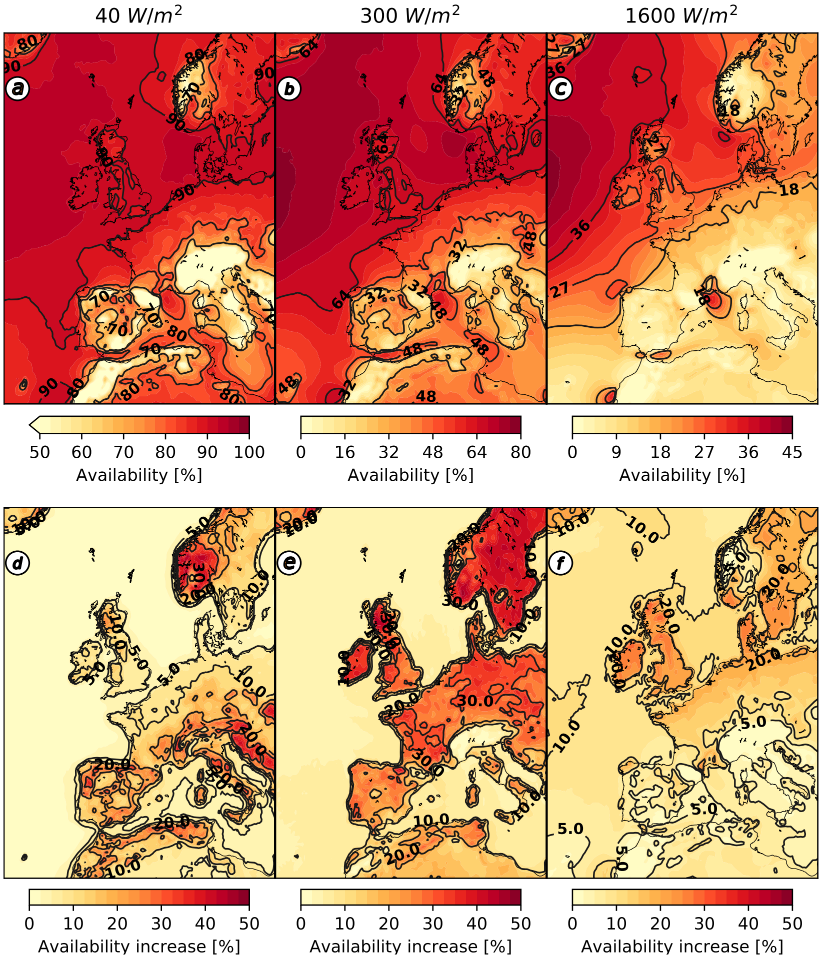

Figure 3.7: Availability of 40, 300, and 1600 W/m^2 wind power density for the 500 m ceiling case (**a**–**c**) and relative increase with respect to the 100 m fixed-height case (**d**–**f**).

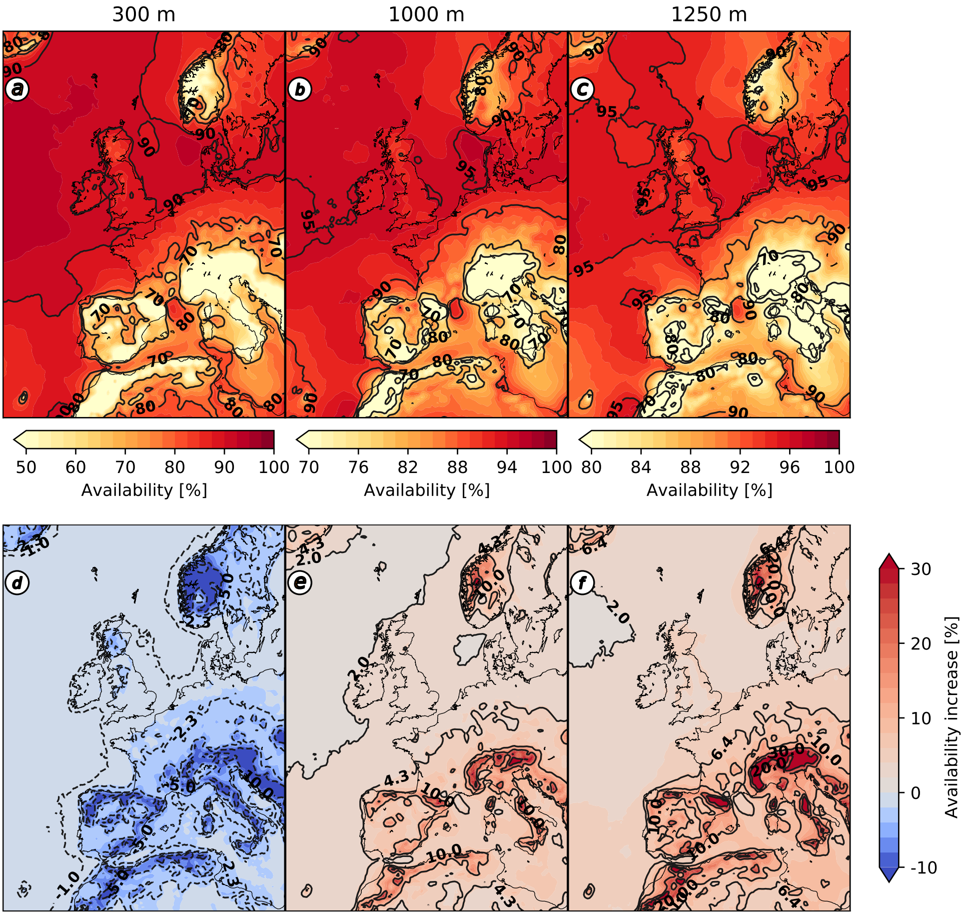

Figure 3.8: Availability of 40 W/m^2 wind power density for variable-height cases with 300, 1000, and 1250 m ceiling (**a**–**c**) and the corresponding increase relative to the 500 m ceiling case (**d**–**f**).

4

Establishing the wind profile climatology

The previous chapter incorporated the paradigm of variable-height harvesting in the assessment of the available wind power. This chapter summarises the wind climate in such a way that it can adequately be coupled to the flight operation model. The prevalent vertical wind profiles identified in this chapter are input for characterising the power production.

4.1. Introduction

To estimate the annual energy production (AEP), measured or modelled wind speed statistics close to the ground are commonly extrapolated to higher altitudes to obtain the wind speed statistics in the full operational height range of the AWE system using either the wind profile power law or the logarithmic profile. This way of representing the wind resource introduces substantial uncertainties since the aforementioned wind profile relationships are not strictly valid beyond the surface layer. Moreover, within this layer, not all wind profiles can be described well with these relationships.

The wind direction can vary substantially with height in the lower atmosphere [85, 86]. A limitation of both the power law and logarithmic profile is that they provide no information about any wind direction dependence with height. In addition, the relationships assume that wind speed increases monotonically with height. In practice, low-level maxima in wind speed, with decreasing wind speed above (low-level jets), are likely to occur, which is also observed in reanalysis data [87, 88].

Computationally expensive brute-force energy production calculations avoid the assumption of any kind of wind profile relationship by using historical wind data over the full operational height range. Ranneberg et al. combine COSMO-DE forecast data with power curves for multiple heights (independent of the wind profile shape) to estimate the AEP [47]. This is a valid approach if the system is operating at a nearly constant

Parts of this chapter have been published in Wind Energy Science **5**, 3 (2020).

height. However, the wind profile shape has to be considered if the system operates in a larger height range, as is the case for a flexible-kite AWE system [46]. AEP calculations become more computationally expensive if the wind profile shape is considered, especially when identifying the optimal cycle settings for all time points. Malz et al. use three months of three-hourly MERRA-2 reanalysis data and speed up the computation by a factor of 20 by using the solution of the previous optimisation to initialise the next [58]. In a follow-up study, Malz et al. use this approach to determine the AEP of an AWE system for 20 locations in Europe [59]. The current state-of-the-art lacks a methodology that can be confidently used to make efficient AEP calculations for a pumping AWE system that sweeps a non-negligible height range.

4

Atmospheric stability strongly influences wind shear and is often used to characterise the wind profile shape. The Obukhov length L is used as a measure of atmospheric stability. This length is not easily measured or derived from model data and is generally inferred indirectly. One way to do this is to fit a functional form of the logarithmic wind profile with stability correction to the wind velocity magnitude profile. Basu uses this approach based on the wind speeds at three levels [89]. Another common way of estimating L, is by inferring it from the gradient Richardson number, $\mathrm{Ri_G}$. By assuming a functional form of the stability correction, L can be derived from $\mathrm{Ri_G}$ [31]:

$$\frac{\bar{z}}{L} = \begin{cases} \frac{\mathrm{Ri_G}}{1-5\mathrm{Ri_G}}, & \text{if } \mathrm{Ri_G} \geq 0 \\ \mathrm{Ri_G}, & \text{otherwise} \end{cases} \quad , \tag{4.1}$$

in which $\bar{z}$ is a reference height which is commonly taken as either the arithmetic or geometric mean of the heights used to determine the temperature and wind speed differences. To approximate $\mathrm{Ri_G}$, a finite difference can be used which is equivalent to the bulk Richardson number, $\mathrm{Ri_B}$. This property expresses the ratio between the temperature stratification and the wind shear:

$$\mathrm{Ri_B} = \frac{g}{\bar{\theta}_v} \frac{\Delta\theta_v \Delta z}{\Delta v^2} \quad , \tag{4.2}$$

in which g is the gravitational acceleration, $\bar{\theta}_v$ is the mean virtual potential temperature, and $\Delta\theta_v$ and Δv are the virtual potential temperature difference and the horizontal wind speed difference, respectively, determined over the height difference Δz. Positive (negative) $\mathrm{Ri_B}$ values indicate stable (unstable) stratification and values close to zero indicate neutral stratification. To approximate L, $\mathrm{Ri_G}$ is substituted with the calculated value for $\mathrm{Ri_B}$ in Equation 4.1.

As an alternative to categorising wind profiles based on atmospheric stability, data-driven techniques have been used to identify wind profile patterns based on features of the wind profile itself. Sommerfeld et al. apply k-means clustering to subdivide wind profile datasets from lidar observations into two clusters for a location in a mostly flat area in northern Germany [90]. Durán et al. use self-organising maps to characterise wind profile data for two locations (Cabauw in the centre of the Netherlands and the FINO-1 platform in the North Sea, 45 km north of the German/Dutch coast) from Weather Research and Forecasting (WRF) modelled data using 2300 clusters [91]. The clusters are used for forecast verification and to investigate diurnal and seasonal cycles.

This study proposes a clustering procedure for obtaining representative wind profile shapes from measured or modelled data that include the vertical variation of the wind speed and direction. The data is partitioned into a small number of clusters, and the corresponding cluster-mean wind profile shapes are determined. A data-driven approach is chosen for identifying these shapes and, thereby, the use of a wind profile relationship with limited validity is avoided. Nevertheless, the observed features are interpreted physically. In contrast to earlier studies that cluster wind profiles, the wind profiles are normalised before clustering. Non-dimensionalising the wind profile is often done in wind profile relationships and can yield a more compact wind resource representation. Additionally, the variation of the wind direction with height is included in the wind resource representation as it affects the operation of an AWE system.

This chapter outlines the development of the wind profile climate description based on 10 years of historical wind data. Section 4.2 introduces the onshore and offshore reference locations that are used to demonstrate the climate description. Section 4.3 discusses the data processing and clustering techniques, complemented by interim results. Section 4.4 first addresses the clustering of DOWA data and presents the results for the reference locations. Subsequently, the DOWA data of 45 other locations are clustered altogether to generate a generalised set of wind profile shapes that is applicable to an area which includes a wide range of location types. Finally, Section 3.5 summarises the conclusions of this study.

4.2. Reference locations

In principle, any dataset containing time series of wind speeds and directions for multiple altitudes can be used as input for the proposed methodology. An hourly temporal resolution of the datasets suffices for capturing the diurnal cycle of the wind profile. Atmospheric phenomena with shorter time scales are considered subscale to the wind profile and only impose brief, small disturbances. For the relevance to AWE, it is desirable to have wind data at least up to 500 m height. The vertical resolution should be adequate to assess the shape of the wind profile with sufficient detail for the performance calculations. Both long-term lidar observations and modelled data qualify as input. This study focuses on using modelled data, i.e. the DOWA dataset, which provides good spatial and temporal coverage of 10 years, as discussed in Section 2.3.3. As supplementary information, ERA5 is only used to determine the atmospheric stability at the time and location of the analysed wind profiles.

An on- and offshore location in the Netherlands and the North Sea, respectively, are selected to demonstrate the methodology. The offshore location, that of the met mast IJmuiden, is located 85 km off the Dutch coast in the North Sea. The onshore location, namely, the met mast Cabauw, is located in the centre of the Netherlands. The area directly surrounding the mast is flat open grassland for at least 400 m in all directions and up to 2 km in the dominant wind direction, i.e. west-south-west. Furthermore, within a radius of 20 km, the terrain is predominantly grassland and virtually flat. The met mast sites, shown in Figure 4.1, are selected because they are well-known in the literature. Note that the anemometer or lidar measurements of the met masts are not used in this study. The other 45 depicted locations are used to evaluate the full DOWA domain and are selected such that onshore, coastal, and offshore locations are equally represented.

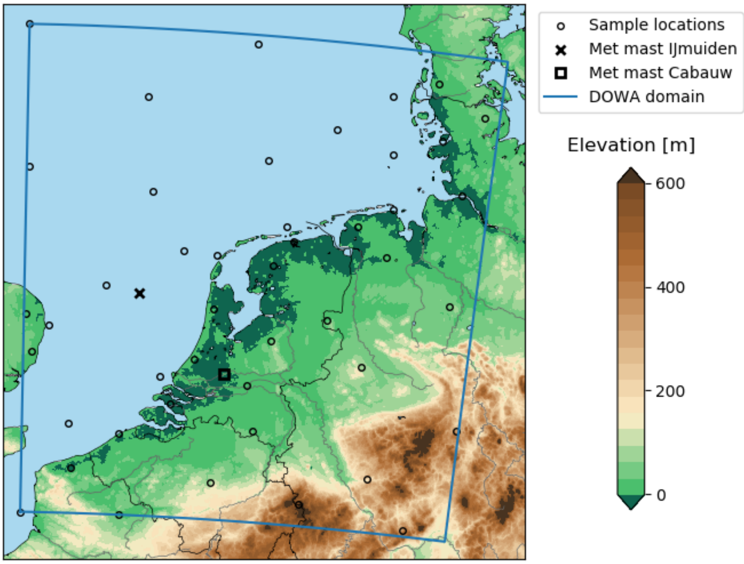

Figure 4.1: The DOWA domain, framed by the blue line, covers the Netherlands, a substantial part of the North Sea, and adjoining coastal areas. The ×, □ and ∘ markers depict the locations analysed in Sect. 4.4.1, 4.4.3, and 4.4.4, respectively. The sea is depicted in light blue, and the colour scale shows the elevation of the land surface [92].

The datasets for the met masts Cabauw and IJmuiden and the 45 locations are referred to as the onshore, offshore, and multi-location datasets, respectively.

4.3. Clustering procedure

This section illustrates the clustering procedure for the offshore location. The data is filtered and normalised, and its dimensions are reduced using a principal component (PC) analysis. Next, the clustering performance is analysed, and the number of clusters is chosen for the wind resource representations analysed in Section 4.4.

4.3.1. Preprocessing of the wind data

The operation of an AWE system is affected by the variation of wind speed and direction with height. Therefore, wind profile shapes are studied with both these features included. Each wind profile sample consists of westerly and southerly wind velocity components for multiple heights (vertical grid points) at a given time and location and is processed in two steps to obtain its shape. Firstly, similar to Kalverla et al. [93] and Malz et al. [58], the wind velocity components are expressed as parallel and perpendicular components relative to the wind velocity at a reference height of 100 m. As a result, the value for the perpendicular wind velocity at 100 m is zero, and the reformatted wind profile is independent of the wind direction at 100 m. Secondly, the wind velocity components are normalised using the 90th percentile of the sample's wind velocity magnitudes. Using the percentile makes the normalisation less sensitive to outliers than using the maximum value. The normalised parallel and perpendicular wind speeds together form the *wind profile shape* of a sample. The normalisation yields a more compact wind resource representation, however, it is prone to producing irregular wind profile shapes for low winds. Therefore, the wind profiles that have a mean wind speed below 5 m s^{-1} are filtered out to obtain the representative wind profile shapes. Note that low wind conditions only contribute a small part of the AEP of a wind energy system. Therefore, they do not necessarily need to be well-represented in the wind profile shape set considered in the AEP calculation. Nevertheless, they cannot be disregarded altogether in the AEP calculation and need to be reintroduced at a later stage such that their frequency is properly considered.

4.3.2. Principal component analysis of the wind profile shape dataset

The mean wind profile shape for the offshore location is illustrated in Figure 4.2a by plotting the normalised wind speed $\tilde{v}$ against height using profiles for the parallel and perpendicular velocity components. As expected for an offshore location, the mean shape exhibits low wind shear. The hodograph in Figure 4.2e shows how the normalised wind velocity changes with height by plotting the parallel and perpendicular normalised wind speed ($\tilde{v}_{\parallel}$ and $\tilde{v}_{\perp}$) for every height. In accordance with Ekman theory, the mean shape shows wind veer (wind direction turns clockwise with height).

A logarithmic profile with roughness length z_0=0.0002 m, a representative value for open water [94], is fitted to the lower 200 m of the mean shape. 200 m is used as a proxy for the depth of the surface layer for which the logarithmic profile is valid. Though, in very stable and unstable conditions the surface layer could be considerably smaller. Following the approach recommended by [95], the profile is fitted by varying the friction velocity v_* and the stability function Ψ, which are constrained to the functional forms given in Equations 2.4 and 2.5. From this, a mean value of the Obukhov length L can be inferred. The best-fitting profile corresponds to a value L=-3391 m, implying a neutral logarithmic profile in the surface layer (assuming neutral conditions if $|L| > 500$). Above 200 m, the fit slightly deviates from the mean profile.

Prior to clustering, a PC analysis is used to reduce the dimensionality of the dataset, while preserving most of the variance. This reduces the computational effort and thus speeds up the clustering. The PC analysis specifies a transformation from the original to the PC coordinate system with its origin coinciding with the mean of the dataset. The first axis is oriented such that it accounts for most of the variance in the data. Subsequent axes are perpendicular to their predecessors and oriented such that they account for as much of the variance as possible. As a result, the last axis accounts for the least of the variance. The PCs are unit vectors in the direction of the positive PC axes.

The compositions of the first two PCs of the offshore dataset are illustrated in the second column of Figure 4.2. The coefficients of each PC describe the relationship between the PC and the parallel and perpendicular normalised wind velocity components at the 17 heights. The absolute values of the PC coefficients quantify the contribution of the respective normalised wind velocity components to the PC. The contributions of the perpendicular components account for most of PC1, indicating that PC1 mostly characterises wind veer. In contrast, the contributions of the parallel components account for most of PC2, indicating that PC2 mostly characterises wind shear. Both PCs show large

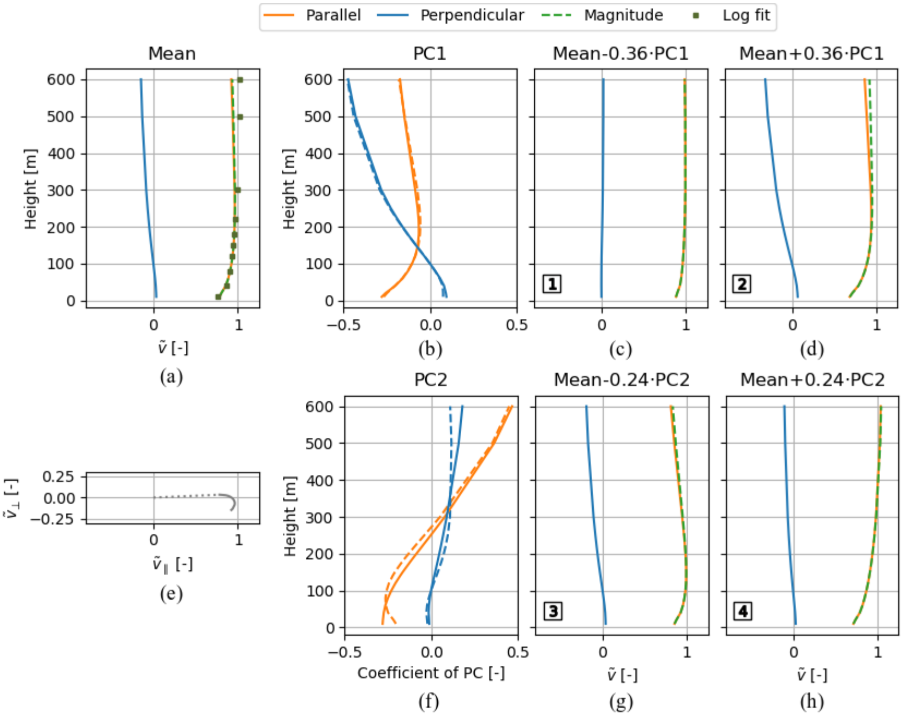

Figure 4.2: Mean wind profile shape and corresponding non-adiabatic logarithmic profile fit (a) and corresponding hodograph (e) for the filtered offshore dataset. Composition of the first and second PCs (b and f). The average of the PC1 (PC2) profiles from the two reference locations is plotted alongside the offshore PC1 (PC2) profile using the dashed line. PC multiplicands superimposed on the mean wind profile shape using minus (c and g) and plus (d and h) one standard deviation as multipliers. The wind profile shape numbers 1–4 refer to the markers in Figure 4.4a.

contributions at both ends of the height range, which indicates that most variance in the dataset is found at these heights. In the PC-space, the data is expressed by multiplicands of the PCs superimposed on the mean wind profile shape. The third and fourth columns of Figure 4.2 show the wind profile shapes that correspond to the points on the PC1 and PC2 axes at -1 and +1 standard deviation and illustrate how the shape varies along both PCs. 68 % of the PC1 (PC2) values lie between the values used for generating wind profile shapes 1 and 2 (3 and 4). Indeed, the wind veer differs substantially between wind profile shapes 1 and 2 and the wind shear between 3 and 4.

The percentage of variance retained after dimensionality reduction depends on how many PCs are used to express the data. The relation between the percentage of variance retained and the number of PCs follows from the PC analysis and is shown in Figure 4.3. The first four PCs already account for more than 90 % of the variance in the offshore dataset. Since the wind velocities of neighbouring vertical grid points are highly correlated, most of the variance in the data is retained using a limited number of PCs. Retaining 90 % of the variance or more is considered acceptable for this application. Although

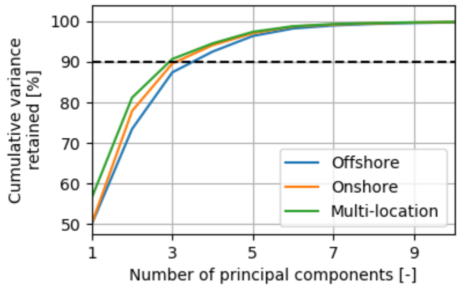

Figure 4.3: Relationship between the percentage of variance retained and the number of PCs for the filtered offshore, onshore, and multi-location datasets analysed in Sect. 4.4.1, 4.4.3, and 4.4.4, respectively.

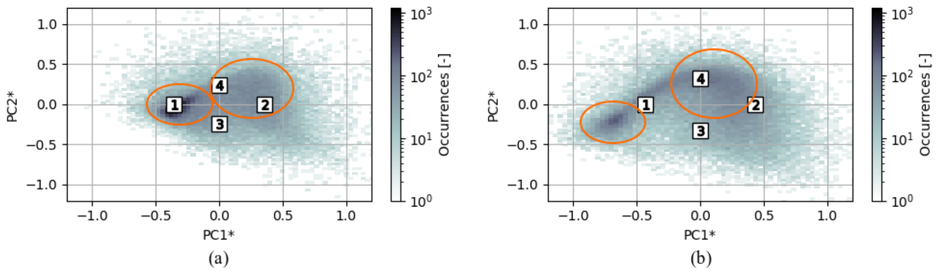

Figure 4.4: Sample frequency distributions in the PC1, PC2-space for the offshore (a) and onshore (b) locations. The origins coincide with the mean wind profile shapes, and the markers with the wind profile shapes numbered 1–4 in Figures 4.2 and 4.11. The orange ellipses indicate the visually identified clusters. The plots share the same coordinate system with the x-axis (y-axis) aligned with the average of the PC1 (PC2) unit vectors from the two reference locations. The averaged PCs are denoted by an asterisk and their profiles shown in Figures 4.2 and 4.11 (b and f).

the use of four PCs already meets this requirement, five PC are used since the variance retained still increases a few per cent between four and five PCs. The preprocessed data is mapped onto the PC1–5-space and used as input for the clustering.

Figure 4.4a shows the frequency distribution of the wind profile shapes in the PC1, PC2-space. The PC1, PC2-projection of the wind profile shapes in the third and fourth columns of Figure 4.2 are indicated with the markers. By visual inspection, two relatively dense groups of data points are identified: a confined group and a less confined group which resembles a tail extending from the first group, marked with the left and right ellipses, respectively. Figure 4.4b shows results for the onshore location and will be discussed in Sect. 4.4.3.

4.3.3. Choosing the number of clusters

K-means clustering [96] is applied to identify the set of wind profile shapes that are used for representing the wind resource. Each cluster is represented by its centroid, and each sample is assigned to the cluster with the nearest centroid. The clustering algorithm iteratively searches for the positions of the centroids that minimise the sum of the squared Euclidean distances between the centroids and their associated samples. This cost function is also referred to as the within-cluster sum of squares (WCSS). The resulting centroids reflect the *cluster-mean wind profile shapes* in the dataset, which follow from back-transforming the cluster-centroids from the PC to physical space.

K-means clustering is always able to produce a result, which makes it very powerful but also potentially deceptive. The algorithm tends to produce spherical clusters with equal radius and sample size and works best on data with such a structure. The previous visual analysis of Figure 4.4 revealed a different structure type for the wind profile shape datasets with two unevenly sized groups of data points. The number of clusters k generated by the algorithm needs to be specified by the user, and it is often not evident how many clusters to choose. The elbow and silhouette methods are used for finding an appropriate number for k. Moreover, the choice for k is evaluated in the context of applying the cluster-mean wind profile shapes to represent the wind resource.

The elbow method investigates the trend of WCSS against k. Increasing the number of clusters is equivalent to reducing the WCSS. Kinks in the trend indicate appropriate choices for k. The elbow plot in Figure 4.5a shows no distinct kinks for more than three clusters.

The silhouette score expresses the similarity of a sample to the other samples in its cluster relative to its similarity to the nearest neighbouring cluster's samples. The dimensionless score ranges from -1 to 1: a negative value suggests that the sample is assigned to the wrong cluster, a value around zero indicates that the sample lies between two clusters, and a high value indicates that the sample is assigned to a distinct cluster. Figure 4.5b shows the mean silhouette score is highest for two clusters. The division of the dataset into two clusters thus yields the most cohesive clusters, which is in agreement with the visual inspection of Figure 4.4a. The decreasing trend of silhouette score with k implies that, in general, a small number of clusters should be used to maintain cluster cohesiveness.

After obtaining the cluster-mean wind profile shapes, they are used for constructing the *cluster representation* of the wind resource. Each sample's absolute vertical wind speed profile is approximated by scaling the associated cluster-mean wind profile shape using the normalisation wind speed used in the pre-processing. The accuracy of this cluster representation is assessed using the mean fit error over all filtered samples. The fit error of the j^{th} sample is calculated by the root mean square of the errors at each vertical grid point. Two different expressions are used to evaluate the error at the i^{th} vertical grid point: the wind velocity magnitude error $\varepsilon_{i,j}$ and that which includes both the parallel and perpendicular wind speed errors $\varepsilon_{\parallel,i,j}$ and $\varepsilon_{\perp,i,j}$. The resulting magnitude and two-component forms of the mean fit error, E_{mag} and E_{2c}, are given by:

$$E_{\text{mag}} = \frac{1}{n_{\text{s}}} \sum_{j=1}^{n_{\text{s}}} \left(\sqrt{\frac{1}{n_{\text{h}}} \sum_{i=1}^{n_{\text{h}}} \varepsilon_{i,j}^2} \right) \tag{4.3}$$

and

$$E_{\text{2c}} = \frac{1}{n_{\text{s}}} \sum_{j=1}^{n_{\text{s}}} \left(\sqrt{\frac{1}{2\, n_{\text{h}}} \sum_{i=1}^{n_{\text{h}}} \left(\varepsilon_{\parallel,i,j}{}^2 + \varepsilon_{\perp,i,j}{}^2 \right)} \right) \quad , \tag{4.4}$$

in which n_h is the number of heights, n_s is the number of samples. The relation between both mean fit errors and the number of clusters is shown in Figure 4.5c.

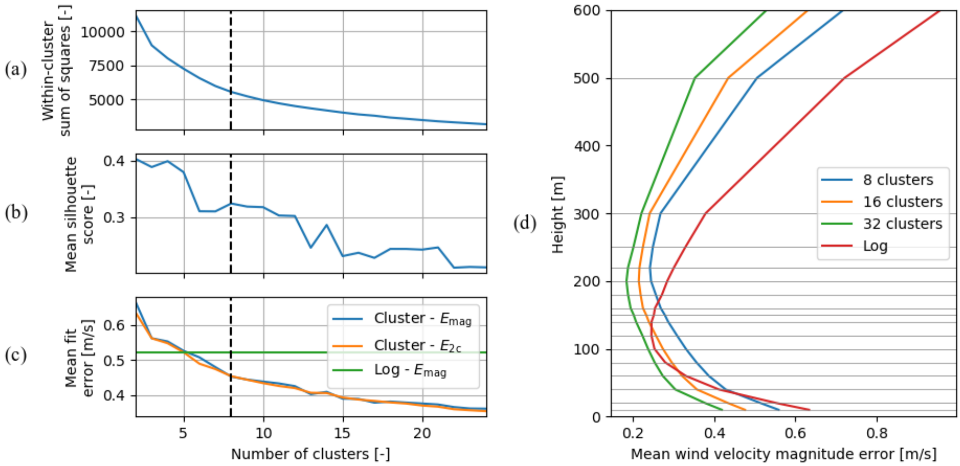

Figure 4.5: Sensitivity of the k-means clustering performance to the number of clusters over the full vertical grid (a–c) and for each height (d) for the filtered offshore dataset. Cost function of the clustering algorithm (a), cluster cohesiveness metric (b), and the mean wind speed fit error (c) against the number of clusters. The dashed vertical lines depict the final choice for eight clusters.

The use of the cluster representation is considered valid when it yields a higher accuracy than a representation that uses logarithmic profiles to approximate the vertical variation of the horizontal wind speed. The logarithmic wind resource representation is obtained by fitting logarithmic profiles with roughness length z_0=0.0002 m to each sample. Here, the Obukhov length L passed to the Ψ stability function is restricted to the representative values of the five stability classes, listed in the third column of Table 2.1. Moreover, the fit is performed to the full height range, i.e. 10–600 m, to minimise the fit error of the wind resource representation and, therefore, allow the application of the logarithmic profile relationship beyond the surface layer. As the logarithmic representation does not include information about the wind direction variation with height, its accuracy is only assessed using E_{mag}. The fit error of the cluster representation is evaluated in relation to the number of clusters and compared to the fit error of the logarithmic representation. Figure 4.5c shows that whether the fit error of the cluster representation is evaluated using E_{mag} or E_{2c} makes little difference. The cluster representation is more accurate than the logarithmic representation when using five clusters or more.

The wind resource representations do not yield the same accuracy for each vertical grid point. To investigate the height dependency, the mean wind velocity magnitude error over all filtered samples is calculated for each vertical grid point. The results are shown in Figure 4.5d, in which the horizontal lines depict the 17 heights of the vertical grid points of DOWA. The fits have a relatively low error around 150 m height and a higher error at the top and bottom of the vertical grid. Around 150 m height, the grid is relatively fine, which is equivalent to allocating more weight to the 100–200 m interval for the logarithmic profile fitting procedure. As a result, the fitting favours minimising the errors in this interval over those at both ends of the height range. Note that the sensitivity of the cluster representation to the grid spacing is limited by the PC analysis prior

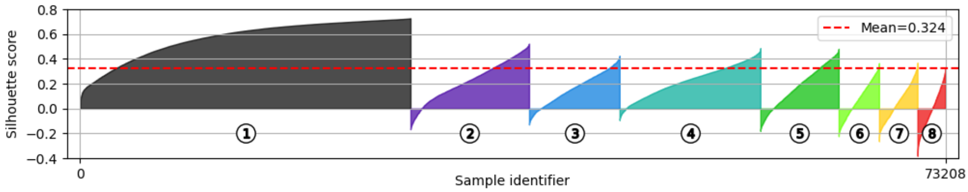

Cluster label	Mean silhouette score
MMIJ-1	0.567
MMIJ-2	0.226
MMIJ-3	0.164
MMIJ-4	0.226
MMIJ-5	0.184
MMIJ-6	0.080
MMIJ-7	0.056
MMIJ-8	-0.009
Overall	0.324

Figure 4.6: The silhouette scores of the individual samples grouped by cluster and in ascending order for the filtered offshore dataset. The numbered markers and filled area colours indicate to which cluster a sample belongs. The overall mean score is indicated by the dashed line, and the table below the figure states the mean score for each cluster.

to the fitting. As stated before, the PC1 and PC2 profiles show that most variance in the dataset is found at both ends of the height range. Due to the relatively high variance and fit model deficiencies, the fit error is also expected to be the largest at these heights. Although the error of the cluster representations at 100 m is higher than that of the logarithmic representation, on average they perform substantially better.

The choice of the number of clusters used to represent the wind resource depends on the type and application of the analysis. Here, eight clusters are chosen for investigating their characteristics in Sect. 4.4. This choice follows from a trade-off between the mean wind profile fit error, the silhouette score, representation validity, and the aim to present a concise analysis and meaningful interpretation of the resulting clusters. To get more insight into the structures of the eight offshore clusters (MMIJ-1–8), the mean silhouette score is calculated for each cluster. The higher the mean silhouette score, the more likely that a cluster represents a natural structure in the data. Figure 4.6 shows that a large fraction of the samples have high silhouette scores for MMIJ-1–4, indicating that MMIJ-1–4 are relatively cohesive clusters. The silhouette score distributions of MMIJ-5–8 indicate less uniform sets of samples, especially that of MMIJ-8. Note that the MMIJ-1 score is roughly a factor of 2.5 larger than the second-biggest cluster despite the tendency of k-means clustering to produce equally sized clusters.

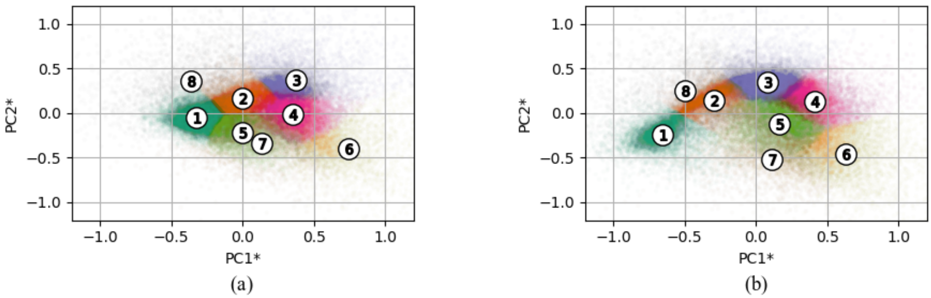

Figure 4.7: Projection of the samples onto the PC1, PC2-space for the offshore (a) and onshore (b) locations. The colour indicates to which cluster a sample belongs and the markers represent the cluster centroids. The plots share the same coordinate system with the x-axis (y-axis) aligned with the average of the PC1 (PC2) unit vectors from the two reference locations. The averaged PCs are denoted by an asterisk and their profiles shown in Figures 4.2 and 4.11 (b and f).

4

4.4. Cluster wind resource representation

In this section, the physical interpretation of the cluster representations is discussed. Firstly, the clusters and their cluster-mean wind profile shapes that result from the offshore dataset are presented. For each cluster, patterns in the times of occurrence are studied together with their association to wind properties at 100 m and atmospheric stability. The analysis is then repeated for the onshore location. The cluster sets for both reference locations are compared to shed some light on the similarities and differences between them. Finally, data from 45 locations are combined to obtain a single set of clusters that is applicable to the entire DOWA domain. For each of the resulting clusters, a map is generated depicting the cluster frequency distribution over the DOWA domain.

4.4.1. Cluster representation for the offshore location

The clustering of the dataset for the offshore location at the met mast IJmuiden yields eight clusters (MMIJ-1–8), which are represented by their centroids shown in Figure 4.7a. The clusters are well spread over the PC1, PC2-space, with the exception of MMIJ-5 and 7, which are relatively close to each other. Note that only two axes of the five-dimensional PC-space are shown. Table 4.1 lists all five PC-coordinates of the cluster-centroids. It shows that MMIJ-5 and 7 are furthest apart in the PC3-direction. The PC4 and PC5 values have a substantially smaller range than that for PC1–3 and are superfluous for distinguishing between the eight clusters.

The cluster-mean wind profile shapes of the offshore clusters are shown in Figure 4.8. Logarithmic profiles with roughness length z_0=0.0002 m are fitted to the magnitude profiles and shown for comparison. Here, the Obukhov length used in the stability function is varied freely, and the fit is restricted to the lower 200 m. The values for the Obukhov lengths inferred from the fits are plotted as 500 m/L in Figure 4.9 and categorised using the stability classes in Table 2.1. The comparison serves to show to what extent the cluster shapes deviate from non-adiabatic logarithmic profiles, particularly above the surface layer.

To investigate the characteristics of each cluster, Figure 4.10a–c show how the clus-

Table 4.1: Principal component coordinates of the cluster-centroids for the filtered offshore dataset. The centroid positions in the PC1, PC2-space are depicted in Figure 4.7a with the numbered markers.

Cluster label	PC1	PC2	PC3	PC4	PC5
MMIJ-1	-0.33	-0.05	-0.04	-0.02	0.01
MMIJ-2	0	0.17	-0.08	0.05	0.02
MMIJ-3	0.38	0.38	-0.01	0.05	0
MMIJ-4	0.35	-0.02	0.04	-0.09	-0.06
MMIJ-5	0	-0.22	-0.16	0.07	-0.04
MMIJ-6	0.74	-0.4	0.02	0	0.12
MMIJ-7	0.14	-0.33	0.44	0.09	0.03
MMIJ-8	-0.36	0.36	0.45	0.04	0.02

ters are distributed over the years, months, and hours of the day. Figure 4.10a shows that the inter-annual variability is limited, which asserts that the results can safely be generalised to the lifetime of a wind energy system (~20 years). The absolute frequency on the y-axis serves to show the cluster sizes. Figure 4.10d–f show the relative frequency of each cluster for different conditions in terms of wind speed, wind direction, and atmospheric stability. As for the logarithmic profile fits, the stability of each sample is classified using Table 2.1.

For generating Figure 4.10f, the stability class distributions are derived using the bulk Richardson number $\mathrm{Ri_B}$ converted to the Obukhov length L using Equations 4.2 and 4.1. The arithmetic mean of the model level heights is used for $\bar{z}$ in order to convert $\mathrm{Ri_B}$ to L. The data from either ERA5 or DOWA could be used to derive $\mathrm{Ri_B}$, however, a comparison of the corresponding Obukhov length ranges shows that the data from the two lowest ERA5 model levels, i.e., ~10–31 m yield the most realistic values.

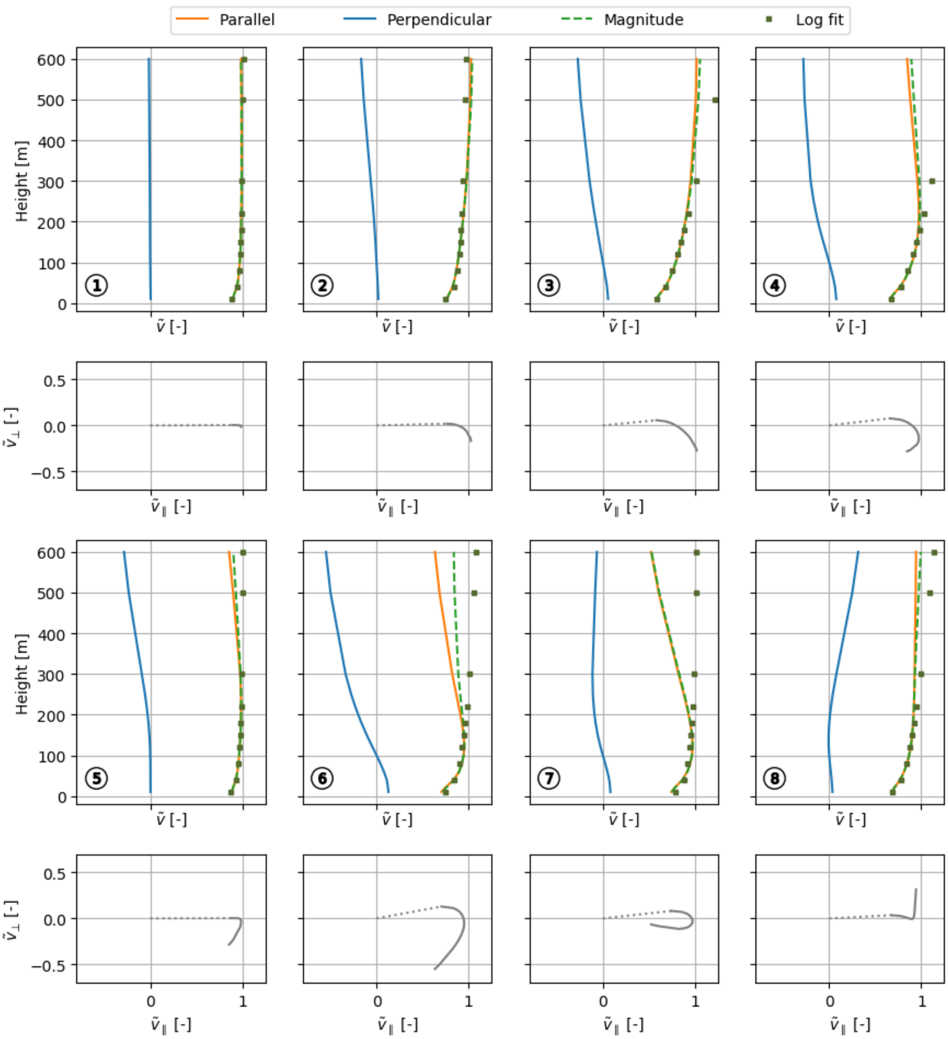

Figure 4.8: The eight cluster-mean wind profile shapes of the offshore clusters (MMIJ-1–8). Each shape is depicted by the normalised wind velocity components with height (first and third rows) with the corresponding hodograph below (second and fourth rows). Non-adiabatic logarithmic profile fits are plotted alongside the shapes. In each hodograph, the lower end of the profile is indicated by the dotted line connecting the lowest height point to the origin. All plots share the same x-axis.

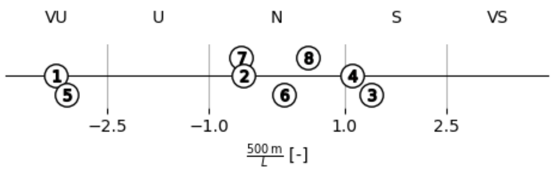

Figure 4.9: Obukhov lengths (plotted as 500 m/L) found by fitting logarithmic profiles to the offshore cluster-mean wind profile shapes in Figure 4.8. The stability classes are adopted from Table 2.1.

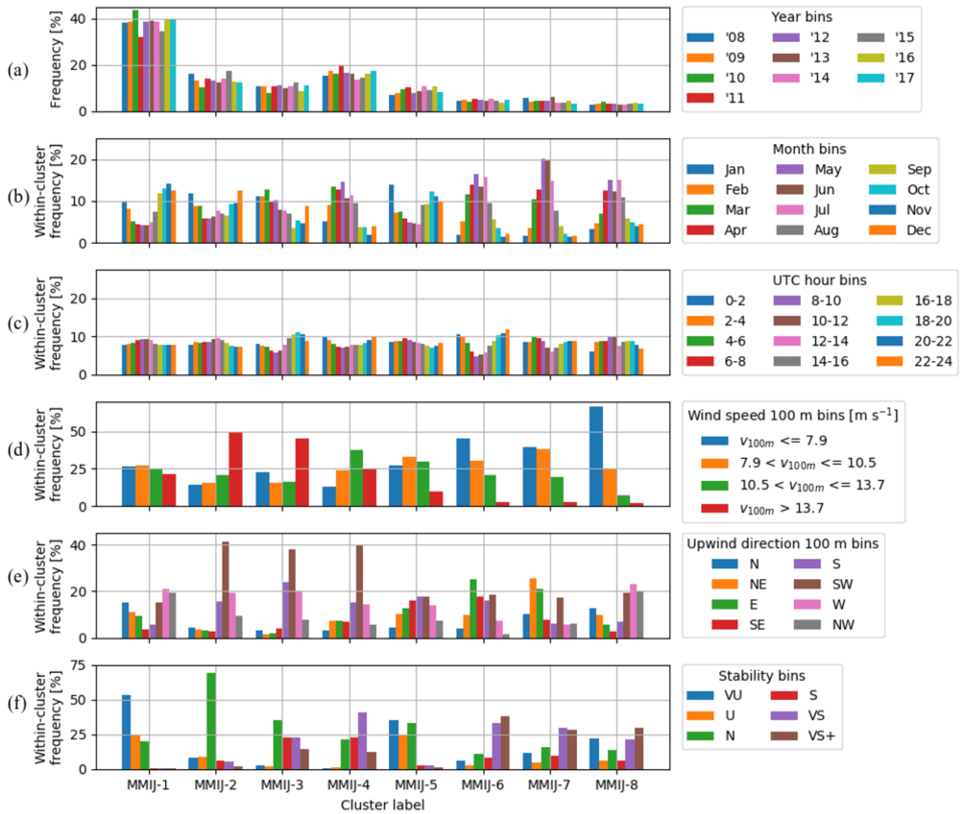

Figure 4.10: Frequency distributions broken down into bins by time of occurrence (a, b, and c), wind speed and direction at 100 m (d and e), and atmospheric stability (f) for the filtered offshore dataset. The wind speed bin limits are chosen such that the frequency over all clusters for each bin is roughly the same. The stability bins correspond to the classes in Table 2.1 together with the VS+ bin ($Ri_B \geq 0.2$). The other distributions have equal bin widths.

4.4.2. Interpretation of the offshore cluster representation

Examining Figures 4.8 and 4.9 shows that the cluster-mean wind profile shapes differ from standard logarithmic profiles, particularly above the surface layer. Moreover, by referring to Figure 4.10, it is possible to investigate the conditions under which the clusters occur and to gain insight into the weather systems that are causing them.

Figure 4.8 shows that the MMIJ-1 and 2 magnitude profiles are well-described with logarithmic profiles. The MMIJ-1 profile shape suggests a well-mixed convective profile with little wind shear and veer. MMIJ-1 occurs predominantly in autumn and is slightly more frequent in the morning hours. The wind is more frequently weak or moderate than strong and mostly comes from the westerly, north-westerly or northerly directions. Furthermore, this cluster occurs predominantly during unstable conditions. These observations make sense as in autumn, the relatively warm sea water favours neutral to unstable stratification; the dominant wind directions have long fetches over sea which allows the boundary layer to reach an equilibrium state due to the relatively constant surface forcing.

The MMIJ-2 and 3 profile shapes show an increase in wind shear relative to MMIJ-1. The MMIJ-2 profile shape closely resembles a neutral logarithmic profile up to 600 m, whereas that of MMIJ-3 only shows a good fit with a stable logarithmic profile in the surface layer. These clusters occur typically during strong winds, predominantly from the southwest. Strong south-westerly winds are characteristic of the wind climate at this mid-latitude location, which is dominated by the frequent passage of low-pressure systems. The relatively strong winds coincide with the highest occurrence of near-neutral conditions, in particular for MMIJ-2. For MMIJ-3, also stable conditions are frequent. MMIJ-2 occurs more often in the late autumn and MMIJ-3 in winter and the start of spring. The colder sea water in spring favours the formation of stable stratification, which explains the difference in stability distribution between the two clusters. Stable stratification suppresses turbulent mixing, which helps to sustain a strong wind shear, consistent with the increasing wind shear and veer seen in Figures 4.7a and 4.8.

The MMIJ-4–7 profile shapes are all jet-like. Because wind speed increases monotonically with height in the logarithmic wind profile relationship, it cannot describe these types of profile shapes. The wind direction and stability distributions associated with the MMIJ-4 cluster are correlated with south-westerly winds and stable stratification. The seasonal cycle is very pronounced and peaks in spring when stable stratification is frequent. The winds recorded for MMIJ-4 are mostly moderate to strong. The distributions associated with the MMIJ-5 cluster are very similar to those of MMIJ-1, with the exception of the wind direction distribution, which shows an opposite trend. The winds with a southerly component are dominant for MMIJ-5 and typically have shorter fetches over sea than the north and westerly winds seen for MMIJ-1. The hodograph of the MMIJ-5 profile shape indicates a rather abrupt kink around 140 m, suggesting a discontinuity such as a boundary-layer top. The MMIJ-6 profile shape shows a maximum at 120 m. Although the magnitude profiles of MMIJ-5 and 6 look similar, the MMIJ-6 profile shape veers more. The MMIJ-7 profile shape shows the most pronounced jet-like shape, also peaking around 120 m. MMIJ-6 and 7 occur almost exclusively for very stable conditions in spring and for weak wind situations. Both clusters occur predominantly for winds with an easterly component and show a diurnal cycle with fewer occurrences around noon. Such a diurnal cycle is in agreement with various studies that have linked low-level jets and the diurnal variation of both the land-sea temperature difference and the intensity of turbulent mixing [97–100].

The hodographs of MMIJ-5 and 8 both show a sharp bend around 140 m. However, the wind direction turns anticlockwise with height above the bend for MMIJ-8, which is opposite to the veering of the other profile shapes. Despite the peculiarities of the wind direction profiles, the magnitude profiles of MMIJ-5 and 8 are described reasonably well below 200 m with very unstable and neutral logarithmic profiles, respectively. MMIJ-8 occurs mostly in spring, under stable conditions, and more often for winds with a westerly rather than a southerly component. Note that this shape belongs to an incohesive cluster and, therefore, gives a relatively poor representation of the cluster samples.

4.4.3. Comparing the on- and offshore cluster representations

The dataset for the onshore location at the met mast Cabauw is clustered using the same approach. The eight resulting clusters are referred to as MMC-1–8. The results of the PC analysis of the onshore dataset are shown in Figure 4.11 and compared to the offshore results, shown in Figure 4.2. A logarithmic profile with roughness length z_0=0.1 m is fitted to the mean wind profile shape as before. This value falls in the higher end of the measurement-inferred roughness lengths for the area surrounding the mast Cabauw [101]. With a stability function value corresponding to L=476 m, the mean profile shape below 200 m is in accordance with a stable logarithmic profile. Above that, the fitted logarithmic profile rapidly diverges from the mean shape. A higher wind shear is observed than for the offshore location due to the higher surface roughness. The hodograph in Figure 4.11e shows that also the wind veer is substantially increased. Despite the apparent differences in mean shape, the PC1 and PC2 profiles are very similar for both reference locations. The average of the PC1 (PC2) profiles from the two reference locations is plotted alongside the onshore PC1 (PC2) profile using the dashed line. To enable a direct comparison between results, the same coordinate system is used for Figures 4.4a and b. The x-axis (y-axis) is aligned with the average of the PC1 (PC2) unit vectors from the two reference locations.

The distribution in Figure 4.4b shows a similar pattern to Figure 4.4a: a dense, confined group of samples, marked with the left ellipse, with a tail of samples extending from this group at around 45° towards the right ellipse. In general, the samples of the onshore dataset are more spread out than the offshore samples, particularly along the PC1 axis. Also, the confined group is less dense for the onshore location. Figure 4.7 shows that, for both locations, the samples of these confined groups belong to the on- and offshore clusters with number 1. The remaining onshore clusters with monotonic wind speed and veering profiles, MMC-2–4, account for most of the tail, see Figure 4.12. Note that the clustering algorithm labels the clusters arbitrarily with numbers. These labels are manually adjusted such that the onshore cluster numbers align with the offshore cluster numbers. This allows us to draw parallels between them and show that the resulting profiles are very similar between both locations, e.g., the first offshore clusters (MMIJ-1–3) also have monotonic profiles.

Logarithmic profiles with roughness length z_0=0.1 m are fitted to the cluster-mean wind profile shapes and plotted alongside them in Figure 4.12. The values for the Obukhov lengths inferred from the fits are shown in Figure 4.13 and categorised by stability class. For the offshore location, Figure 4.9 shows that six out of eight logarithmic profiles found by fitting are neutral or stable and those for MMIJ-1 and MMIJ-5 are more unstable. Figure 4.13 shows that only one unstable logarithmic profile is found for the onshore location, next to six stable and one neutral logarithmic profiles. Since there is little diversity in the shape of the unstable profiles, all the associated samples are grouped together by the clustering. The fact that this type of profile is well-mixed with little shear and a relatively high boundary layer height explains why the diversity is small. By contrast, the neutral and stable profiles can have a wide range of shear, and in addition, particularly in stable conditions, the boundary layer height can be quite low which will have a strong influence on wind shear. This means that a greater diversity of profile shapes is to be expected under neutral or stable conditions.

The profile shapes for MMC-1 to MMC-3 show an increase in wind shear. Between MMC-3 and MMC-4, the wind veer increases and the wind shear reduces. The profile shapes for MMC-5–7 are jet-like, as is the case for the offshore clusters MMIJ-4–7. MMC-5 and MMC-6 have similar wind velocity magnitude profiles with a relatively weak wind speed maximum around 200 m, but MMC-6 shows a much stronger wind veer. MMC-7 shows the strongest fall-off above 200 m. Like its offshore counterpart, MMC-8 is characterised by an anticlockwise-turning profile with a sharp bend, which is most clearly visible in the hodograph. Recall that the offshore wind profile shape for MMIJ-5 also showed a sharp bend, albeit, in combination with clockwise turning. These features are not observed for any of the MMC profile shapes.

Figure 4.14 shows that MMC-3 is the most frequent cluster in the filtered onshore dataset, with a frequency of 20.6 %. The first five onshore clusters have similar total frequencies, whereas MMIJ-1 dominates for the offshore location. As for the offshore clusters MMIJ-6–8, the onshore clusters MMC-6–8 are less frequent.

Figure 4.14 shows clusters that typically occur during spring/summer (MMC-1, MMC-7 and MMC-8) or autumn/winter (MMC-2–6). The diurnal cycles of the onshore location are highly pronounced in contrast to those of the offshore location. This effect is caused by the lower heat capacity of the land surface which promotes a more immediate heat transfer to or from the atmosphere. Convection created by solar irradiation leads to more turbulent mixing during the day than at night. Indeed, MMC-1 and MMC-2 are mixed profiles and predominantly occur during the day, whereas MMC-3–8 show profiles with less mixing and predominantly occur during the night. Note that low-level jets, and stable conditions in general, occur almost exclusively at night. Figure 4.14c indicates a pronounced diurnal cycle in atmospheric stability for the onshore location, whereas for the offshore location the seasonal cycle, shown in Figure 4.10b, is more pronounced. Figures 4.10d and 4.14d display almost identical frequency distributions over the bins. However, the actual wind speed distributions differ due to the different bin limit values. Note that the chosen limits give the same total frequency for each bin. Thereby, the distributions of the individual clusters are easily related to the uniform general distribution and compared with one another. Also, the wind direction distributions show similar patterns for both locations. In the case of the stability distributions, the onshore location shows a tendency to more stable conditions for all clusters.

In conclusion, very similar cluster-mean wind profile shapes have been identified for

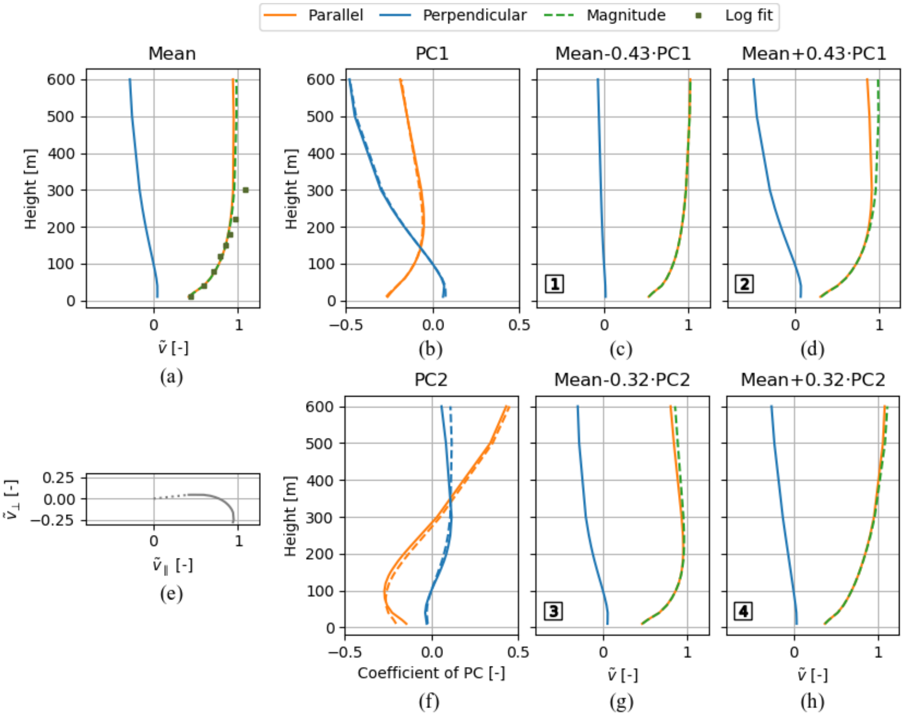

Figure 4.11: Mean wind profile shape and corresponding non-adiabatic logarithmic profile fit (a) and corresponding hodograph (e) for the filtered onshore dataset. Composition of the first and second PCs (b and f). The average of the PC1 (PC2) profiles from the two reference locations is plotted alongside the onshore PC1 (PC2) profile using the dashed line. PC multiplicands superimposed on the mean wind profile shape using minus (c and g) and plus (d and h) one standard deviation as multipliers. The wind profile shape numbers 1–4 refer to the markers in Figure 4.4b.

the on- and offshore reference locations. Moreover, similar profiles seem to be related to similar conditions in terms of wind speed, wind direction, and atmospheric stability. The strongest winds typically act to neutralise the stratification, leading to monotonic profiles with relatively little veer. These profiles are relatively well captured by logarithmic wind profiles. For weaker winds, atmospheric stability acts to enhance wind shear and veer, up to the point where low-level jets are observed. However, whereas stability at the offshore location is governed by a clear seasonal cycle in the underlying sea surface, stability over land is regulated by the relatively rapid diurnal heating cycle of the land surface. Over sea, the wind direction also seems to play a more pronounced role, since it controls the characteristics of the prevailing fetch.

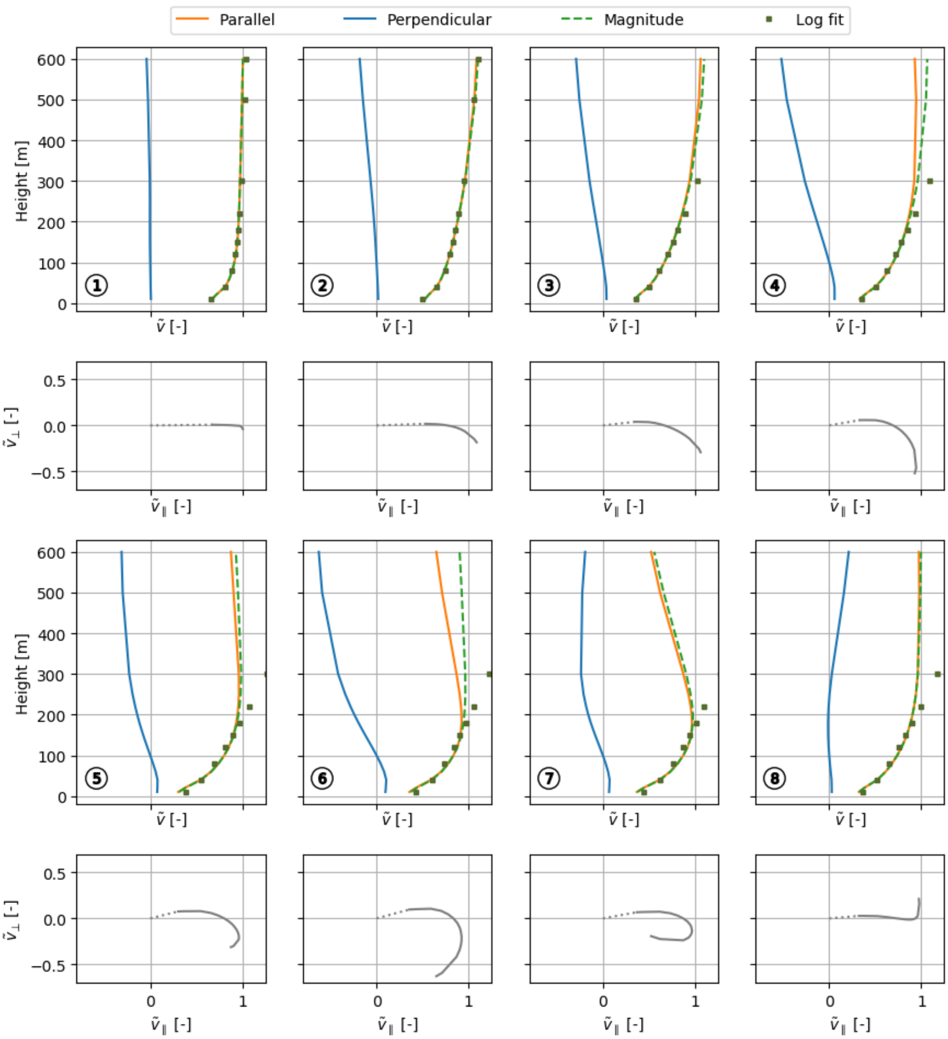

Figure 4.12: The eight cluster-mean wind profile shapes of the onshore clusters (MMC-1–8). Each shape is depicted by the normalised wind velocity components with height (first and third rows) with the corresponding hodograph below (second and fourth rows). Non-adiabatic logarithmic profile fits are plotted alongside the shapes. In each hodograph, the lower end of the profile is indicated by the dotted line connecting the lowest height point to the origin. All plots share the same x-axis.

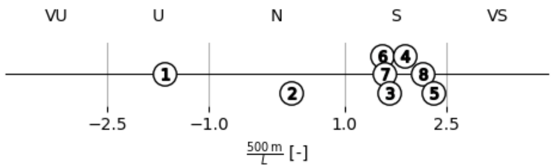

Figure 4.13: Obukhov lengths (plotted as 500 m/L) found by fitting logarithmic profiles to the onshore cluster-mean wind profile shapes in Figure 4.12. The stability classes are adopted from Table 2.1.

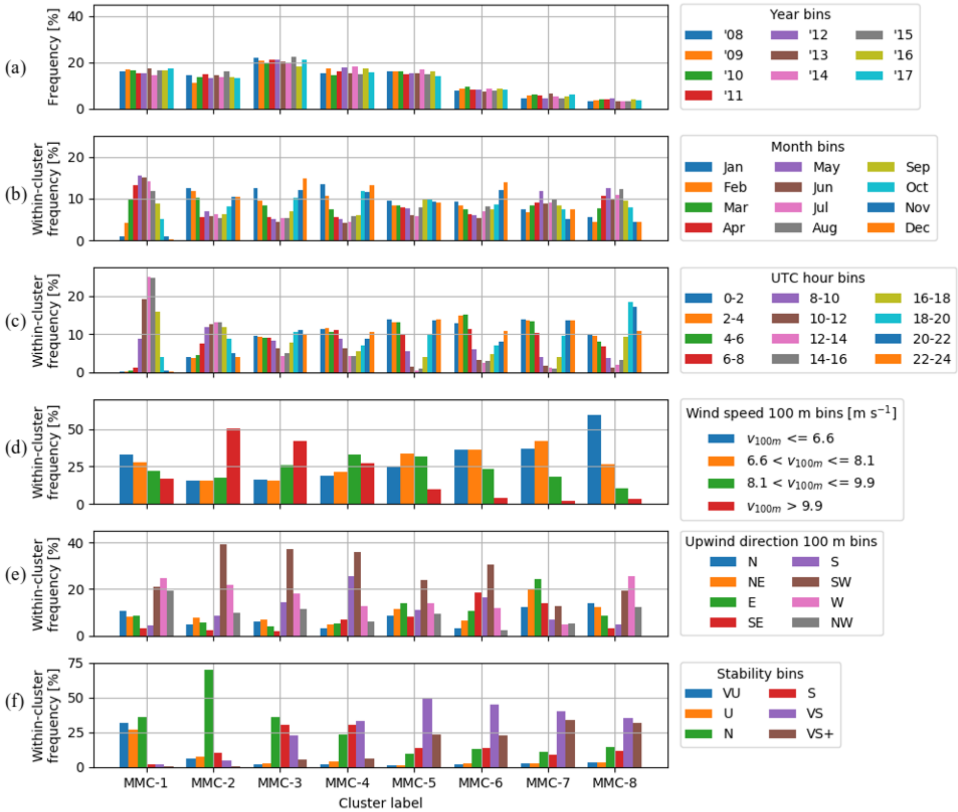

Figure 4.14: Frequency distributions broken down into bins by time of occurrence (a, b, and c), wind speed and direction at 100 m (d and e), and atmospheric stability (f) for the filtered onshore dataset. The wind speed bin limits are chosen such that the frequency over all clusters for each bin is roughly the same. The stability bins correspond to the classes in Table 2.1 together with the VS+ bin ($Ri_B \geq 0.2$). The other distributions have equal bin widths.

4.4.4. Spatial frequency distribution of wind profile shape clusters

The clustering algorithm is applied to a dataset that includes wind data from a variety of locations to investigate how the frequency of the clusters varies spatially. The multi-location dataset (filtered to exclude low wind samples) includes wind data from 45 DOWA grid points that are selected such that onshore, coastal, and offshore locations are equally represented. For each location type, 15 grid points are chosen pseudo-randomly to yield a good coverage of the full DOWA domain (50778 grid points in total). The sampled grid points are marked on the map in Figure 4.1. The aim is to give some insight into the spatial variability of wind profile characteristics, in particular, to see how the clustering approach highlights profile characteristics of the on- and offshore environments. In principle, the multi-location approach gives a set of profile shapes that could be used for an AEP assessment, though a site-specific set would be better suited if a more accurate assessment is required. Whilst the applicability of the cluster representation is increased to a larger area by keeping the number of clusters the same, the accuracy is compromised. For the purpose of evaluating the spatial variability of the wind profile shapes, the accuracy of the representation is not critical and does not require increasing the number of clusters. The eight resulting multi-location clusters are referred to as ML-1–8.

Figure 4.15 shows the cluster-mean wind profile shapes for each of the multi-location clusters. Each sample of every grid point in the DOWA domain is assigned to the cluster with the closest centroid. For each cluster, a map is generated showing the spatial distribution of its frequency of occurrence, see Figure 4.16. Note that the colour scale is different for each map such that spatial patterns are easier to observe. Table 4.2 lists the frequency of each cluster at the on- and offshore reference locations. It is interesting to compare the multi-location clusters with the site-specific clusters identified earlier. With a frequency of 48.5 %, ML-1 is dominant at the met mast IJmuiden and, therefore, is expected to be similar to MMIJ-1, which was identified as the dominant cluster in the offshore analysis. Comparing Figures 4.8 and 4.15 indeed shows that the cluster-mean wind profile shapes of ML-1 and MMIJ-1 look alike. Similarly, ML-7 has the highest frequency at the met mast Cabauw, i.e. 21.7 %, and has a profile shape somewhere in between those of MMC-3 and 4, the most frequent clusters resulting from the onshore analysis. Every multi-location cluster is manually linked to the single-location clusters based on the resemblance of their cluster-mean wind profile shapes, see Table 4.2.

The maps in Figure 4.16 show a distinct division between clusters that mostly occur over sea (ML-1–3) and over land (ML-4–8). The latter group is sub-divided into coastal and onshore clusters, see Table 4.2. The sharply defined patterns in the frequency maps of ML-5–8 coincide with orographic features and thus suggest a strong relationship between the clusters and orography. Other site characteristics such as recurring weather systems and land cover also affect the clusters and thus the frequency maps. Over land, the frequency maps of ML-5 and ML-6 suggest an inverse relationship: the frequency of ML-5 peaks at high elevations, whereas that of ML-6 is highest in the river valley in the lower right corner of the DOWA domain. A similar inverse relationship is observed between ML-7 and ML-8. Also, the frequency maps of ML-7 and ML-8 show contours coinciding with the elevation map, though the relationship between the frequency and elevation is not as direct as for ML-5 and ML-6.

Table 4.2: Classification of the multiple-location clusters and frequencies of occurrence of the clusters at the on- and offshore reference locations (met masts Cabauw and IJmuiden).

Cluster label	Class	Similar single location cluster(s)	Frequency at offshore location	Frequency at onshore location
ML-1	offshore	MMIJ-1	48.5 %	5.8 %
ML-2	offshore	MMIJ-2, 3	22.0 %	4.2 %
ML-3	offshore	MMIJ-4, 6	14.1 %	4.2 %
ML-4	coastal	MMC-1	8.6 %	16.1 %
ML-5	onshore/ coastal	MMC-2	3.0 %	17.4 %
ML-6	onshore	MMC-6	2.0 %	13.3 %
ML-7	onshore	MMC-3, 4	1.2 %	21.7 %
ML-8	onshore	MMC-5	0.6 %	17.3 %

4.5. Conclusion

This chapter presented a methodology for including multiple wind profile shapes in a wind resource description. A data-driven approach is used to identify a set of wind profile shapes that characterise the wind resource. These shapes go beyond the height range for which conventional wind profile relationships are developed, such as the logarithmic profile. Moreover, they include non-monotonic wind profile shapes such as low-level jets. The methodology is demonstrated for an on- and offshore reference location using DOWA data.

To obtain the wind profile shapes of the DOWA samples, the wind profile of each sample is expressed relative to its wind velocity at the 100 m reference height and normalised. A PC analysis shows that three PCs already account for about 90 % of the variance in the dataset. The first and second PCs are very similar for the datasets of the onshore and offshore locations. The first PC mostly characterises wind veer, whereas the second PC mostly characterises wind shear. Moreover, the analysis reveals a natural structure of the data in the principal component space with two relatively dense groups of data points. The data points for the onshore location are more spread out, indicating a larger variety of wind profile shapes.

The dataset is partitioned using k-means clustering. The resulting cluster-mean wind profile shapes are used to approximate the vertical variation of the wind, yielding the cluster wind resource representation. This representation reduces the wide variety of wind conditions in the DOWA dataset to a reasonable number of wind profile shapes. The accuracy of the representation using three or more clusters is already higher than that of a representation using logarithmic wind profiles. The eight cluster-mean wind profile shapes of the offshore representation include three monotonic profiles, four jet-like profiles, and an anticlockwise-turning, sharply-bent profile. Very similar cluster-mean wind profile shapes have been identified for the onshore location occurring under similar conditions. A single set of clusters is generated that is representative of the entire DOWA domain and used to analyse the spatial variability of the frequency of occurrence of the clusters. The cluster frequency maps indicate a clear distinction between onshore and offshore clusters. The sharply defined patterns in the frequency maps of the onshore clusters coincide with orographic features and thus suggest a strong relationship between the wind profile shape and orography.

The presented methodology has the capability to produce a single set of wind profile shapes that is valid for a large area. Such a set can facilitate the standardisation of wind conditions for which AWE systems are benchmarked in terms of power production. Moreover, the multi-location cluster representation enables a spatial analysis to assess which installation site is best for an AWE system in terms of its AEP, which makes this methodology a very powerful tool for project developers.

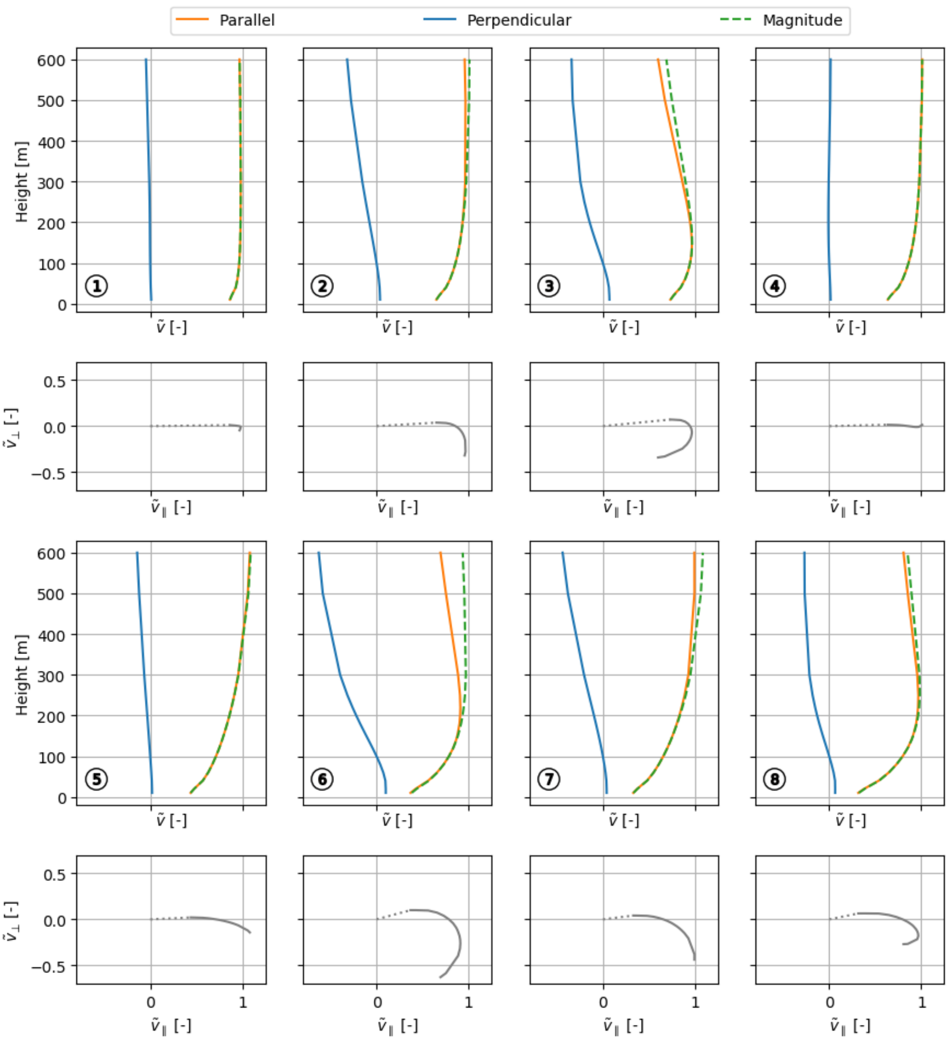

Figure 4.15: The eight cluster-mean wind profile shapes of the multi-location clusters (ML-1–8). Each shape is depicted by the normalised wind velocity components with height (first and third rows) with the corresponding hodograph below (second and fourth rows). In each hodograph, the lower end of the profile is indicated by the dotted line connecting the lowest height point to the origin. All plots share the same x-axis.

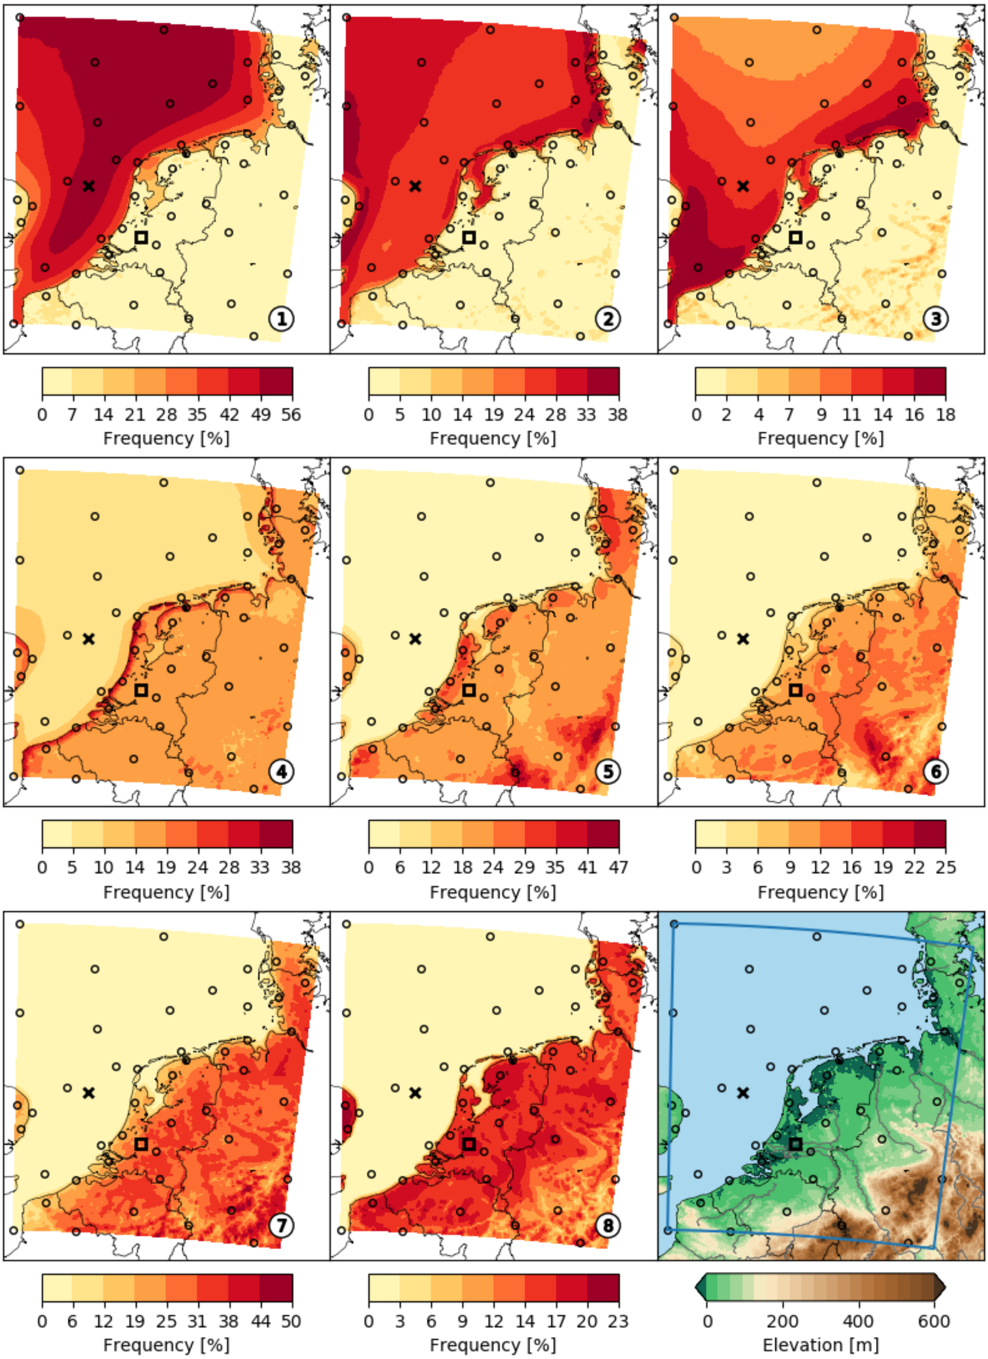

Figure 4.16: Frequency of occurrence of each multiple-location cluster (ML-1–8) mapped over the DOWA domain. The × and □ markers depict the reference locations of the met masts IJmuiden and Cabauw. The ∘ markers show the sampled grid points. The lower right plot is a repetition of Figure 4.1.

5

Reconstructing the wind profile and characterising the kite aerodynamics

The previous chapter established a wind profile climate description that can be coupled with a flight operation model. Selecting and configuring an adequate flight operation model requires a good understanding of the flight behaviour of real-world systems. This chapter investigates such flight behaviour using flight data from an actual system and estimates the instantaneous vertical wind profile using the relative wind measured at the kite. The estimated wind profiles are input to the model validation in Chapter 7.

5.1. Introduction

Loyd [18] was arguably the first to compute the potential of cross-wind flight and derived the expression for optimal instantaneous power generated by a kite in steady state at the direct downwind position:

$$P = \frac{2}{27}\rho v_\mathrm{w}^3 S C_L \left(\frac{C_L}{C_D}\right)^2 \quad , \tag{5.1}$$

in which S is the area of the kite and C_L and C_D are the lift and drag coefficients, respectively. After Loyd, many other researchers built on this theory. Schmehl et al. [102] further developed the steady flight state approximation in spherical coordinates to include the effect of not flying at the direct downwind position and accounting for the weight of the system. The extended expression for the instantaneous power of a massless kite is:

$$P = \frac{1}{2}\rho v_\mathrm{w}^3 S C_R \left[1 + \left(\frac{C_L}{C_D}\right)^2\right] f\left(\cos\beta\cos\phi - f\right)^2 \quad , \tag{5.2}$$

Parts of this chapter have been published in Journal of Physics: Conference Series **1618**, 3 (2020).

in which $C_R = \sqrt{C_L^2 + C_D^2}$, f is the reeling factor, β is the elevation angle, and ϕ is the azimuth angle of the kite. An iterative procedure is proposed to calculate the power when accounting for weight.

Modelling the flight of an AWE system requires knowledge about its flight characteristics. Flight characteristics of aircraft are typically obtained with experiments on a scale model in the controlled environment of a wind tunnel. Experiments on flexible kites should be conducted at the original scale since the similarity parameters cannot be respected at a reduced scale. Consequently, the flight characteristics of a scale model will not be representative of the real kite. Due to the scaling issue, only a few wind tunnel experiments have been conducted on kites. Alternatively, outdoor towing tests have been used to identify the flight characteristics of the kite.

Another source of data for aerodynamic identification is test flights in the field. The challenge with using test flight data is that the degree of control over the wind environment and system states in an experiment is limited. Also, these states are hard to measure and often incomplete, e.g., typically the wind is only measured at the ground, the accuracy of the kite velocity measurement is limited, and measurements at the tether are sparse. The uncertainty about the states is passed on to the identified model parameters.

The validation of system models requires knowledge about the true wind velocity at the kite and not just the apparent wind velocity. Measuring the wind velocity at the ground is of limited use for estimating the wind velocity at the kite because the wind velocity is subject to turbulence and changes with height and can thus cause large temporal and spatial variations. To enable an effective validation, it is best to measure the wind profile directly to reduce uncertainty, e.g., using a lidar system next to the operating system. Alternatively, the wind velocity at the kite can be inferred from the flow and kite velocity measurements. However, the associated uncertainty may nullify the validation.

Oehler and Schmehl [66] designed a flow measurement setup with a self-aligning Pitot tube and employed it on the test flight on 24 March 2017. Flow measurements conducted at the kite yield the apparent wind velocity. The apparent wind velocity, tether force measurements, and estimated bridle attitude were used to characterise the V3-kite of Kitepower B.V.

More recently, Kitepower has equipped its kite with a more robust flow measurement setup. The new setup includes a Pitot tube rigidly mounted to the front bridle of the kite. The three-hour test flight of Kitepower on 8 October 2019 produced a dataset of improved quality and quantity and motivated conducting a more comprehensive aerodynamic identification in which the bridle attitude is inferred from the measured wing attitude. Moreover, the inferred attitude is used to reconstruct the wind velocity at the kite and estimate the wind profile.

First, this chapter provides an overview of the investigated test flight and the measurements that were conducted. Next, the prevailing wind conditions are reconstructed to provide input to the aerodynamic characterisation. Subsequently, the steady flight state approximation is introduced and used to find the lift and drag coefficients. The relationships between these coefficients and the angle of attack of the flow experienced by the kite are investigated. Finally, the identified aerodynamic coefficients are applied in power estimations with the steady flight state approximation and compared with measurements.

5.2. Test flight overview

As a first evaluation of the test flight, the following time-averaged performance indicators of the 87 pumping cycles are determined: the mean cycle power, the cycle duration, and the duty cycle. The duty cycle is defined as the reel-out and cycle duration ratio. Figure 5.1 shows that the indicators fluctuate between cycles and that the mean cycle power and duration exhibit clear trends. The mean cycle power shows an increasing trend with time, while the duration decreases. In the three-hour flight, also the wind conditions are expected to change. This is confirmed by the measured wind speed at the ground, which shows an increasing trend from approximately 5 to 7 m/s. The increasing wind speed explains the increase in mean cycle power.

The measured wind speed at the ground station is compared to the wind speed from ERA5 at three heights: 10, 150, and 250 m. The first height is the lowest height available in ERA5 and lies closest to the height at which the wind speed is measured at the ground station. At the start of the test flight, the ERA5 wind speed at 10 m is higher than the measured speed, and the wind speed is similar towards the end. The other two heights are more representative of the operational height of the kite at which energy is harvested from the wind. Between 10 and 150 m the ERA5 wind speed increases by more than 2 m/s, while the increase is approximately 0.5 m/s between 150 and 250 m.

The remainder of this chapter focuses on instantaneous quantities to evaluate the flight behaviour of the system. The 65$^{\text{th}}$ pumping cycle of the test flight depicted in Figure 2.11 is used to illustrate the presented analyses. This particular pumping cycle is selected as it does not show unexpected flight behaviour or large variations in wind conditions. After interpreting the results for the pumping cycle, the results of the full test flight are presented.

Figure 5.2a-c shows the measured airspeed, force, and inflow angle during the reference pumping cycle. In general, the airspeed increases in the turns (indicated by the shaded areas) as the speed of the kite increases due to the downward flight. This goes hand in hand with an increase in the tether force measured at the ground station. Also, the inflow angle seems to follow a periodic pattern synchronous to the figure-of-eight cross-wind manoeuvres. Dips in the inflow angle occur during the right turns, but not during the left turns.

5.3. Wind reconstruction

The wind conditions dictate how much energy the system can potentially harvest and thus should be known to assess the system performance. Figure 5.1d shows that the mean cycle wind speed (at the ground) can vary substantially between cycles. This quantity does not change gradually as the ~2 minute cycle duration is too short to average out the turbulence. The mean power output of two cycles with the same mean cycle wind is not necessarily the same. In fact, they can differ substantially due to turbulence, e.g., one cycle might experience a near-steady wind field, while during the other cycle, a large eddy (turbulent flow structure) passes through and causes large wind speed and direction variations. This will greatly affect the flight behaviour of the system and, consequently the power output.

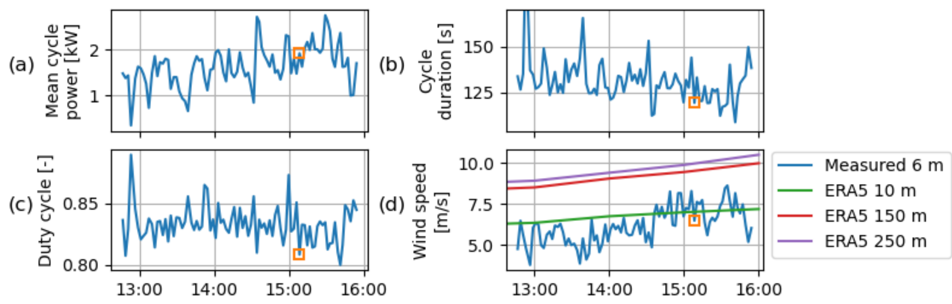

Figure 5.1: Evolution of the key performance indicators and the mean cycle wind speed from anemometer measurements at the ground station over the 87 pumping cycles. The latter plot is complemented with ERA5 (reanalysis) wind speeds at three heights. The reference pumping cycle is indicated with □.

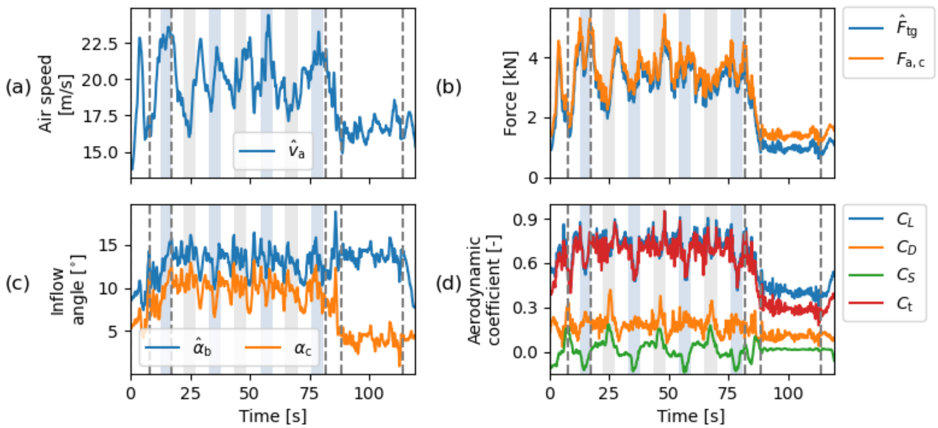

Figure 5.2: Measured flow, measured tether force, and calculated aerodynamic coefficients during the reference pumping cycle. The reel-out and reel-in phases occur between 17–82 s and 89–114 s, respectively. The grey and blue shades indicate left and right turns, respectively.

5.3.1. Instantaneous wind at operational height

Knowledge of the wind field swept by the kite and tether is required to fully explain the motion of the system. Estimating this wind field based on measurements close to the ground is inaccurate. The distance between the ground station and the kite reaches up to 350 m for the reference pumping cycle. In between, myriad turbulent flow structures of varying sizes will exist. As such, the wind speed at the ground is not closely correlated to that at the kite, and they cannot be inferred from one another. Therefore, wind speed measurements close to the ground are of limited use in explaining the flight behaviour of the kite.

More helpful than the wind velocity at the ground for explaining the motion of the system is the wind velocity at the kite. The wind velocity at the kite can be reconstructed from the measurements of the kite velocity and apparent wind velocity. When assuming that the kite does not induce a velocity variation to the free stream wind velocity, the wind velocity $\mathbf{v}_\mathrm{w}$ at the kite can be determined from the apparent wind $\mathbf{v}_\mathrm{a}$ and kite velocity $\mathbf{v}_\mathrm{k}$:

$$\mathbf{v}_\mathrm{w} = \mathbf{v}_\mathrm{a} + \mathbf{v}_\mathrm{k}\,. \tag{5.3}$$

The apparent wind velocity $\mathbf{v}_\mathrm{a}$ expressed in the bridle reference frame, shown in Figure 2.13, can be reconstructed from the airflow measurements:

$$\mathbf{v}_\mathrm{a}^\mathrm{b} = \hat{v}_\mathrm{a} \begin{bmatrix} -\cos\hat{\alpha}_\mathrm{b} \\ 0 \\ \sin\hat{\alpha}_\mathrm{b} \end{bmatrix} \quad , \tag{5.4}$$

in which $\hat{v}_\mathrm{a}$ is the air speed measured by the Pitot tube, which is aligned with the x_b-axis, and α_b is the inflow angle measured by the wind vane. The Pitot tube measurements are assumed to be insensitive to misalignment with the inflowing air. The side slip angle is not measured and is assumed to be zero. Consequently, the lateral apparent wind velocity component is also zero. This assumption implies that the kite is always heading into the apparent wind. In the experiment of Oehler and Schmehl [66], side slip recordings up to ten degrees were common, implying that the zero side slip assumption is not strictly valid.

5

The bridle attitude is required to express the apparent wind velocity in the earth reference frame. The bridle attitude is inferred from the measured wing attitude. Appendix A describes how the bridle reference frame is inferred from the measured wing attitude. Alternatively, Oehler and Schmehl approximate the bridle attitude by assuming that the tether is aligned with the kite position vector and finding the tilt angle of the bridle with respect to the tether using a force balance.

The resulting wind velocity is expressed in the pseudo wind reference frame where the x-axis is aligned with the moving average of the wind direction measured at the ground station. Note that the reference frame is virtually fixed for short periods and thus not continuously turned with the instantaneous wind direction. Thus, a non-zero wind direction indicates that the instantaneous wind is not aligned with the average wind at the ground. Initially, a westerly wind is blowing, gradually becoming west-southwesterly. Thereby, the reference frame rotates approximately 20° counter-clockwise during the test flight.

Figure 5.3a shows that the horizontal wind speed varies substantially. Not all of the depicted variation can be attributed to the actual wind speed. Especially during the turns, the reconstruction is imprecise due to the inaccuracy of the kite velocity and the zero side slip approximation while some side slip is expected during the turns. Consequently, the reconstruction during the straight paths and the reel-in are considered more accurate and plotted with coloured line segments. Based on these flight sections, the wind speed varies approximately between 8 and 11 m/s during the reference cycle.

Figure 5.3b shows a cyclic pattern in the wind direction synchronous to the figure-of-eight cross-wind manoeuvres: the right-to-left straight sections systematically exhibit a larger wind direction than the left-to-right straight sections. This difference cannot be attributed to the wind and results from a systematic error in the calculation caused by the erroneous zero-side-slip assumption.

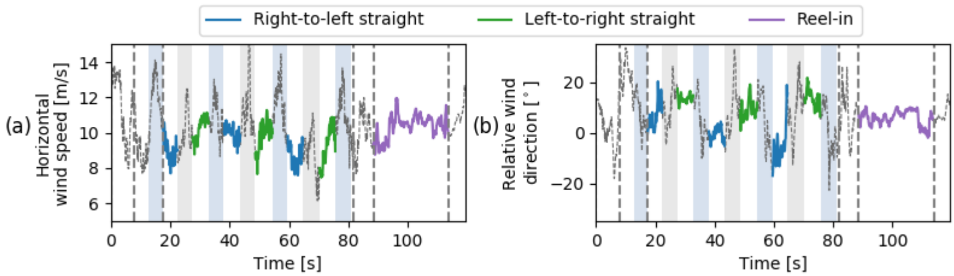

Figure 5.3: Reconstructed horizontal wind speed (**a**) and relative wind direction (**b**) at the varying operational heights of the kite (approximately between 130–270 m). The reconstruction is considered more reliable for the flight sections depicted with the coloured segments. The wind direction is expressed with respect to the average wind direction at the ground.

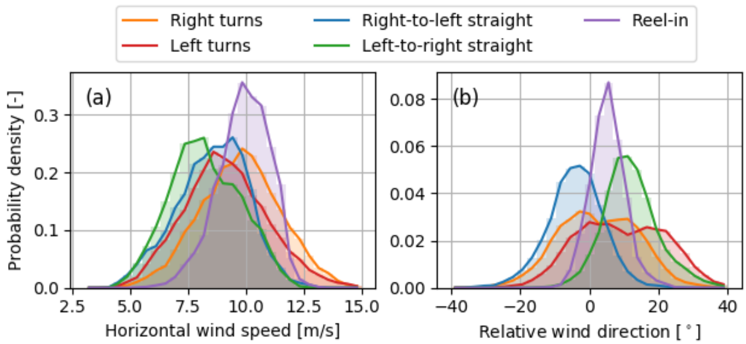

Figure 5.4: Probability density distributions of the reconstructed horizontal wind speed (**a**) and relative wind direction (**b**) at the varying operational height of the kite per flight section for the full test flight.

5

One cycle is too short to get significant wind statistics; therefore, the reconstructed wind is evaluated for the full flight data. Figure 5.4a shows a mean wind speed of approximately 10 m/s during reel-in, which is substantially higher than the mean wind speed of approximately 8.5 m/s during reel-out. This suggests that wind speed increases with height as the kite flies higher up during reel-in.

Figure 5.4b, which shows the equivalent relative wind direction distributions, confirms the systematic error for the straight flight sections during reel-out. The corresponding distributions (shown in blue and green) are approximately mirrored versions of one another about a wind direction of 3°. The relative wind direction distribution for the reel-in flight peaks at approximately 5°. This suggests that the wind is slightly backing, i.e., turning counter-clockwise with height. The collective non-zero mean relative wind direction of the straight flight sections during reel-out suggests that the figure of eight is not aligned with the wind and may cause asymmetric flight behaviour of the kite.

The aligning and turning conditions for the figure-of-eight flight control are based on the wind direction measured at the ground station. In the absence of wind veer, this would not lead to the observed mismatch in airspeed between the turns, given that they

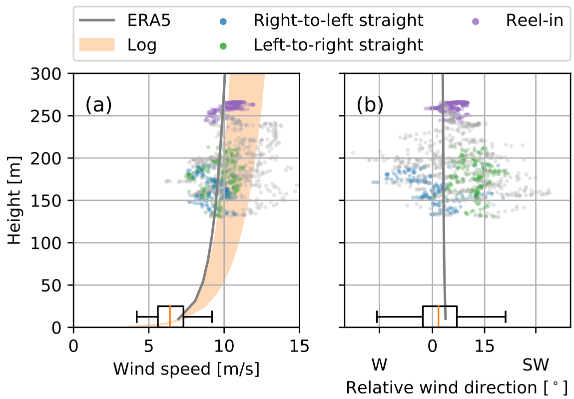

Figure 5.5: Reconstructed horizontal wind speed (**a**) and relative wind direction (**b**) as a function of the height of the kite. The grey points correspond to the transition phases of the pumping cycle and the turns during reel-out, whereas the coloured points correspond to the flight sections with a more reliable reconstruction. The boxplots depict the statistics of the wind measured at the ground station at 6 m height. The grey lines show the ERA5 profiles at 15:00, and the orange area covers logarithmic profiles with roughness lengths between 0.01–0.1 m obtained with the mean wind speed at the ground.

5

are controlled identically and no asymmetries exist in the system.

5.3.2. Reconstructing the vertical wind profile

The performance of an AWE system strongly depends on the vertical wind profile. Therefore, measuring the vertical wind profile is important for assessing the system performance. The wind profile cannot be deduced accurately from measurements at a single height and thus requires wind measurements at multiple heights. As the kite sweeps a substantial height range during a pumping cycle, it can be used as a measuring device to reconstruct the wind profile.

The wind profile reconstruction results presented hereafter are compared to ERA5 data to act as a tentative validation. A definite validation would require wind profile measurements, which are not available. Only large-scale weather phenomena are resolved in ERA5; therefore, it is more suitable as a reference for the wind climate but not for instantaneous wind conditions.

To evaluate the prevailing wind speed profile during the reference cycle, the reconstructed wind speeds are plotted against height in Figure 5.5a. The wind speeds exhibit an increasing trend with height that agrees well with the wind profile obtained from ERA5. The boxplot depicts the wind speed statistics measured at the ground station at 6 m height. Also plotted is the wind speed range covered by the neutral logarithmic profile corresponding to the mean value of the ground wind speed and roughness lengths ranging from z_0=0.01 m to z_0=0.1 m.

The logarithmic profile with the lowest roughness length matches the trend exhibited by the reconstructed wind speeds. Extrapolating the wind speed measured at the ground using the logarithmic wind profile to find the wind speed higher up introduces much uncertainty. At the ground station, the wind is more prone to deviate from the logarithmic wind profile as it is located closer to the roughness elements and local obstacles and thus is more directly affected by them. The logarithmic wind profile does not account for these effects and only includes the collective effect of the roughness elements, which is reflected in the roughness length. The extrapolation amplifies a deviation in wind speed at the ground and can introduce substantial errors, i.e., the presented logarithmic profiles may double the error between the ground station at 6 m height and 250 m height.

The wind direction profile is also reconstructed and plotted in Figure 5.5b. A clear separation between the data points of the two straight flight sections can be observed, with the ERA5 profile lying roughly in between. Recall that this is believed to be an artefact of the zero side slip assumption. This suggests that the true wind direction lies in the middle and agrees well with the ERA5 profile. Note that the ERA5 profile is rather straight but slightly veers, i.e., turns clockwise with height.

Hourly subsets of the flight data are analysed to identify more significant trends in the wind profiles. Figure 5.6 shows how the reconstructed mean, lower quartile, and upper quartile of the wind speed and direction change with height. The reconstructed wind speed for the straight sections during reel-out gradually converges towards the ERA5 profile with height (Figure 5.6a–c). During reel-in, the reconstructed wind speed matches the ERA5 profile well.

The larger differences with respect to the ERA5 profile at lower heights might be explained by the relatively narrow figures of eight at the start of the reel-out. The width of the manoeuvres is likely to affect the side slip and, thereby the reconstruction. Disregarding induction may also contribute to the mismatch. The kite-measured wind speed is not the free stream wind speed, as the presence of the kite affects the airflow in its vicinity. The magnitude of the induced wind speed is, however expected to be too small to explain the observed differences with respect to the ERA5 profile.

Similar to the reference cycle, the logarithmic profiles (obtained from the mean wind speed measured at the ground) tend to overestimate the wind speed higher up for each hourly subset. This difference between the reconstructed and logarithmic profiles grows over time. For the first hour, the lowest-roughness logarithmic profile still lies on top of the reconstructed reel-out profiles, whereas for the last hour, the lowest-roughness logarithmic profile exceeds the reconstructed reel-out profiles. The increased difference indicates that the wind speed measured at the ground exhibits a larger increase in wind speed over time than observed for the lowest point in the ERA5 data. This can also be observed in Figure 5.1d, which shows that the wind speed measurements more frequently exceed the ERA5 data towards the end of the test flight.

The reconstructed wind direction for the straight sections during reel-out converges with height towards the reconstructed wind direction during reel-in, see Figure 5.6d–f. This suggests that the wind velocity reconstruction becomes more reliable with height. The reconstruction seems to agree well with the ERA5 profile for the first hour. However, the reconstruction diverges from the ERA5 wind direction for the subsequent hours. A larger hourly change in mean wind direction is observed in the ground measurements than in the ERA5 data. At first, the wind direction for the lowest point in the ERA5 data is more southerly than the measured wind direction but slightly westerly at the end of the test flight.

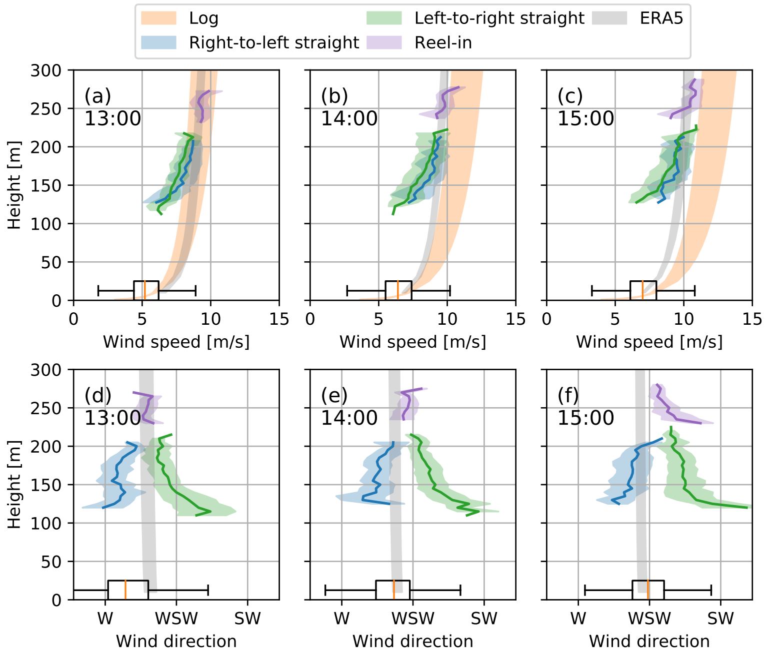

Figure 5.6: Hourly wind profile reconstruction statistics, together with hourly wind measurement statistics at the ground station, the corresponding range of logarithmic wind profiles with roughness lengths between 0.01–0.1 m, and the range bounded by the ERA5 profiles at the beginning and end of the hour.

In conclusion, the wind velocity reconstruction is believed to be too uncertain for identifying the wind profile, e.g., the reconstructed wind speed with height exhibits more shear than is deemed realistic. Side slip measurements should be incorporated as the first step toward a more reliable reconstruction. Moreover, wind profile measurements are needed to provide a more reliable reference than ERA5 data and enable a more definite validation. Lastly, more test flights should be analysed to yield more significant error statistics.

5.4. Steady flight state approximation

The steady-state approximation is used to analyse separate instances recorded in the test flight data. For each instance, all forces acting on the kite are assumed to be in equilibrium. The evolution of the kite position does not need to be solved, as it is recorded in the flight data.

In the simplest form, the system is modelled with a single point mass at the position of the kite. This does not mean that the tether can be disregarded altogether. The weight and drag of the tether substantially affect the flight and thus need to be considered. This is done by lumping the forces to the point mass. Effectively, only the fractions of the weight and drag carried by the kite are lumped to the point mass.

The force balance comprises forces acting on the wing, the KCU, and the tether (subscripts $\frown$, kcu, and t; respectively):

$$\mathbf{L}_{\frown} + \mathbf{D}_{\frown} + \mathbf{F}_{\mathrm{s},\frown} + \mathbf{W}_{\frown} + \mathbf{D}_{\mathrm{kcu}} + \mathbf{W}_{\mathrm{kcu}} + \mathbf{F}_{\mathrm{t}} = 0\,, \tag{5.5}$$

in which $\mathbf{L}_{\frown}$ is the lift force generated by the wing, $\mathbf{D}$ is a drag force, $\mathbf{F}_{\mathrm{s},\frown}$ is the side force generated by the wing, $\mathbf{W}$ is a weight, and $\mathbf{F}_{\mathrm{t}}$ is the tensile tether force acting on the kite.

The aerodynamic forces of the kite are given by:

$$\mathbf{L}_{\frown} = \frac{1}{2}\rho\,\|\mathbf{v}_{\mathrm{a}}\|\,\mathbf{v}_{\mathrm{a}} \times \mathbf{e}_{y_{\mathrm{f}}}\,C_{\mathrm{L}}\left(\alpha_{\frown}\right)\,S\,, \tag{5.6}$$

$$\mathbf{D}_{\frown} = \frac{1}{2}\rho\,\|\mathbf{v}_{\mathrm{a}}\|\,\mathbf{v}_{\mathrm{a}}\,C_{\mathrm{D}}\left(\alpha_{\frown}\right)\,S\,, \tag{5.7}$$

$$\mathbf{F}_{\mathrm{s}} = \frac{1}{2}\rho\,\|\mathbf{v}_{\mathrm{a}}\|^{2}\,\mathbf{e}_{y_{\mathrm{f}}}\,C_{\mathrm{S}}\left(\beta_{\frown}\right)\,S\,, \tag{5.8}$$

$$\mathbf{D}_{\mathrm{kcu}} = \frac{1}{2}\rho\,\|\mathbf{v}_{\mathrm{a}}\|\,\mathbf{v}_{\mathrm{a}}\,C_{\mathrm{D,kcu}}\,A_{\mathrm{kcu}}\,, \tag{5.9}$$

in which ρ is the air density; $\mathbf{v}_{\mathrm{a}}$ is the apparent wind velocity; $\mathbf{e}_{y_{\mathrm{f}}}$ is the unit vector of the y-axis of the flow reference frame; C_{L}, C_{D}, and C_{S} are lift, drag, and side force coefficients; S is the projected wing area, and A_{kcu} is the frontal area of the KCU.

Following the lumping approach of van der Vlugt et al. [46], the tether force can be approximated using:

$$\mathbf{F}_{\mathrm{t}} = \mathbf{D}_{\mathrm{t}*} + \mathbf{W}_{\mathrm{t}*} + \mathbf{F}_{\mathrm{t}*}\,, \tag{5.10}$$

in which the asterisk denotes that the forces are manipulated due to the lumping. In reality, the latter drag and weight do not act directly on the tether but are reflected in the tensile tether force acting on the kite. Between the ground station and the kite, the tether force tilts due to tether sag and increases as the tether needs to carry its own weight and drag.

This lumping approach considers the effect of the cross-radial component of the tether weight on the tether shape. The sag due to tether drag is not considered. $\mathbf{F}_{\mathrm{t}*}$ equals the radial component of tether force at the ground:

$$\mathbf{F}_{\mathrm{t}*}^{\tau} = \begin{bmatrix} 0 \\ 0 \\ -F_{\mathrm{tg,r}} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ -\sqrt{F_{\mathrm{tg}}^{2} - \left(\frac{1}{2}\cos\beta\, m_{\mathrm{t}}\, g\right)^{2}} \end{bmatrix}, \tag{5.11}$$

in which F_{tg} and $F_{\mathrm{tg,r}}$ are the tether force at the ground and its radial component, β is the elevation angle, m_{t} is the tether mass, and g is the gravitational constant. $\mathbf{W}_{\mathrm{t}*}$ has a non-vertical component induced by the lumping approach:

$$\mathbf{W}_{\mathrm{t}*}^{\tau} = \begin{bmatrix} \frac{1}{2}\cos\beta \\ 0 \\ -\sin\beta \end{bmatrix} m_{t}\, g\,. \tag{5.12}$$

For the lumped tether drag, the well-known definition is used that preserves the moment of the tether drag around the ground station [103]:

$$\mathbf{D}_{\mathrm{t}*} = \frac{1}{8}\rho\,\|\mathbf{v}_{\mathrm{a}}\|\,\mathbf{v}_{\mathrm{a}}\,C_{\mathrm{D,t}}\,r_{\mathrm{k}}\,d_{\mathrm{t}}\,, \tag{5.13}$$

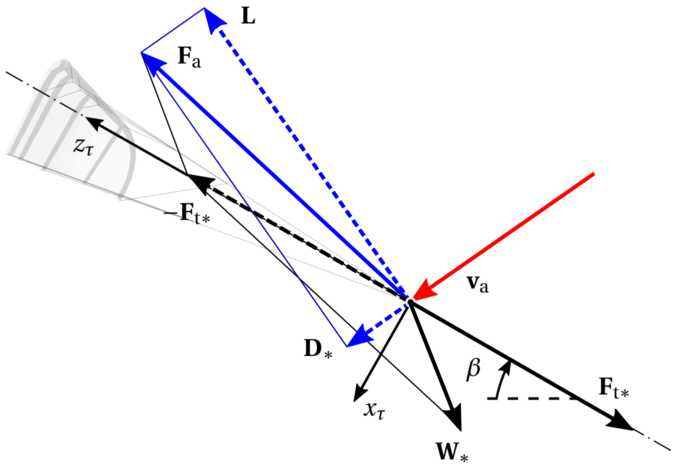

Figure 5.7: Force balance for the special case where the kite is downwind from the ground station and flies up. Note that the depicted drag and weight are the combined forces of the wing, the KCU, and the tether. Adapted from [66].

5

in which $C_{D,t}$ is the tether drag coefficient, r_k is the radial position of the kite, approximated with the tether length, and d_t is the tether diameter.

The force balance is depicted in Figure 5.7. For illustrative purposes, the drag and weight forces of the different components are combined. When flying up or down, the lateral axis of the kite $\mathbf{e}_y$ is aligned with the y-axis of the tangential reference frame. However, when the kite is flying cross-wind, the lift and thus the kite need to be rolled to counterbalance the weight.

Reasonable estimates of the tether and KCU drag coefficients are available, which are assumed to be constant. Table 5.1 lists the coefficients together with other system properties. On the other hand, the aerodynamic coefficients of the wing can vary substantially. The aerodynamic characteristics of flexible kites are less well-established than those of conventional wings. The relationship between the aerodynamic coefficients of conventional airfoils and the angle of attack α is well understood and typically explains much of the variation in the coefficients. Although the coefficients of LEI kites are expected to vary with the angle of attack, other factors, such as the kite deformation, might cause substantial variation.

The kite position and tether force can be directly taken from the flight data to solve the aerodynamic force of the wing. Figure 5.2b shows the evolution of the force during the reference cycle. The approximated aerodynamic force and the tether force measured at the ground $\hat{F}_{tg}$ have approximately the same magnitude. This confirms that the aerodynamic force mostly counterbalances the tether traction force.

5.5. Characterising the aerodynamics of the kite

The steady flight state approximation is used to identify the aerodynamic coefficients of the kite for the reference cycle. Instead of finding the steady state with predefined

Kite		Tether	
Projected wing area, S	19.75 m^2	Density	724 $\mathrm{kg/m}^3$
Wing mass, $m_{\frown}$	14.2 kg	Diameter, d_t	10 mm
Incl. mounted equipment mass	3.2 kg	Drag coefficient, $C_{\mathrm{D,t}}$	1.1
Frontal area KCU, A_kcu	0.25 m^2		
Drag coefficient KCU, $C_{\mathrm{D,kcu}}$	1.0		
KCU mass, m_kcu	25 kg		

Table 5.1: Properties of the airborne components used for the test flight on 8 October 2019.

aerodynamic coefficients, the calculation is reversed to identify the coefficients based on measured information about the state. Next, the statistics of the related measurements are evaluated to explain variations in the resulting coefficients. Ultimately, the identified coefficients are used to characterise the aerodynamics of the wing and identify an aerodynamic model.

5.5.1. Identifying the aerodynamic coefficients

The identification of the aerodynamic coefficients of the wing starts with approximating the aerodynamic force using the force equilibrium. After combining Equations 5.10 and 5.5 and rearranging, the aerodynamic wing force $\mathbf{F}_{\mathrm{a},\frown}$ is obtained:

$$\mathbf{F}_{\mathrm{a},\frown} = \mathbf{L}_{\frown} + \mathbf{D}_{\frown} + \mathbf{F}_{\mathrm{s},\frown} = -\left(\mathbf{W}_{\frown} + \mathbf{D}_\mathrm{kcu} + \mathbf{W}_\mathrm{kcu} + \mathbf{D}_{\mathrm{t}*} + \mathbf{W}_{\mathrm{t}*} + \mathbf{F}_{\mathrm{t}*}\right). \tag{5.14}$$

The right-hand side can be solved using the expressions presented in the previous section and inserting the measured kite position, tether force, and apparent wind velocity.

The aerodynamic coefficients are obtained by solving the lower system of equations:

$$\mathbf{F}^\mathrm{e}_{\mathrm{a},\frown} = \mathbb{T}_\mathrm{eb}\,\mathbb{T}_\mathrm{bf}\begin{bmatrix} D_{\frown} \\ F_\mathrm{s} \\ L \end{bmatrix} = \frac{1}{2}\rho\, v_\mathrm{a}^2\, S\,\mathbb{T}_\mathrm{eb}\begin{bmatrix}\mathbf{e}^\mathrm{b}_\parallel & -\mathbf{1}_y & \mathbf{e}^\mathrm{b}_\perp\end{bmatrix}\begin{bmatrix} C_D \\ C_S \\ C_L \end{bmatrix}, \tag{5.15}$$

in which the properties expressed in the earth and bridle reference frame are denoted by superscript e and b, respectively, and $\mathbf{e}^\mathrm{b}_\perp$ and $\mathbf{e}^\mathrm{b}_\parallel$ are the unit vectors perpendicular and parallel to the apparent wind velocity lying in the kite symmetry plane. Note that the matrix containing the unit vectors equals the rotation matrix $\mathbb{T}_\mathrm{bf}$ for the transformation from the flow to the bridle reference frame for the special case where the side slip $\beta_{\frown} = 0$. $\mathbb{T}_\mathrm{eb}$ is the rotation matrix for the transformation from the bridle to the earth reference frame.

Figure 5.2d shows that the lift and drag coefficients follow a cyclic pattern synchronous to the figure-of-eight manoeuvres during the reel-out phase of the reference cycle. Their mean values during reel-out are 0.70 and 0.19, respectively. A substantial drop in the lift coefficient and an increase in the drag coefficient is observed during the right and left turns, respectively. During the reel-in phase, the lift coefficient is gradually decreasing from approximately 0.45 to 0.33. The mean values of the lift and drag coefficients during reel-in are 0.36 and 0.11, respectively.

The angle of attack experienced by the wing $\alpha_{\frown}$ can be obtained by correcting the measured inflow angle for the pitch angle of the wing with respect to the front bridle:

$$\alpha_{\frown} = \alpha_{\mathrm{b}} - \alpha_{\mathrm{d}} \quad . \tag{5.16}$$

The depower angles α_{d} during the reel-out and reel-in phases are 3.2° and 9.8°, respectively. These are obtained using a geometric model of the bridle presented in Appendix A. Figure 5.2c shows both the measured flow angle and the inferred angle of attack of the wing. Due to the conservative operational approach, the difference in power setting between the phases is small.

The lift coefficient and angle of attack plots in Figure 5.2 do not show a clear relationship between the quantities, e.g., the valleys in the angle of attack during reel-out do not always coincide with valleys in the lift coefficient.

5.5.2. Distinctive modes in the measurements

The variations in measured tether force and air speed explain the variations in the aerodynamic coefficients to a large extent. As such, patterns in the measurements are likely to explain the characteristics of the aerodynamic coefficients. First, the statistics of the measurements of the full flight data are evaluated to gain valuable information for the aerodynamic characterisation.

The normalised tether force combines the measured tether force and airspeed in a simple expression:

$$C_{\mathrm{t}} = \frac{2\,\hat{F}_{\mathrm{tg}}}{\rho\, v_{\mathrm{a}}^2\, A}. \tag{5.17}$$

The air density and the projected surface area of the kite are considered constants. Consequently, the tether force coefficient only varies with the measured tether force and the airspeed. Figure 5.2d shows that the tether force coefficient roughly lies on top of the lift coefficient during the reel-out. This suggests that the variation in the lift coefficient is explained by the variations in measured tether force and airspeed.

Four subsets for different flight sections of the reel-out phase are extracted from the full dataset to study the variations in the tether force coefficient and the measurements. The four flight sections govern turns and the straight paths between them, see Figure 2.11b. Figure 5.8 shows the corresponding probability density distributions of the tether force coefficient and related measurements. The different shapes of the distributions indicate that different modes are active between the flight sections. Differences between turns and straight paths are expected. On top of that, also substantial differences between the left and right turns are observed. The distributions for the two types of straight paths show smaller differences.

Figure 5.8a shows that the tether force coefficients found for the right turns are substantially lower than those for the other reel-out flight sections. Despite that the same control approach is employed for the left and right turns, their difference in performance is larger than anticipated. The coefficients for the flight sections other than the right turns are normally distributed and peak approximately at the same (mean) value.

The difference between the tether force coefficient distributions of the two turns can be mostly explained by the air speed differences. The air speeds recorded during

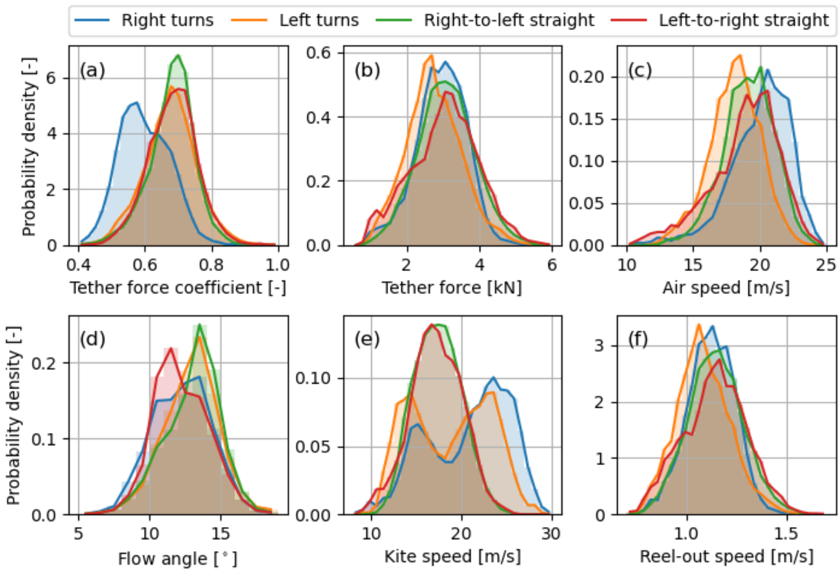

Figure 5.8: Marginal probability density distributions of the tether force coefficient and related measurements for four different flight sections of the reel-out phase.

the right turns are higher (Figure 5.8c), whereas the tether forces are more similar (Figure 5.8b). Surprisingly, the mean air speed of the left turns is the lowest of all flight sections despite relatively high kite speeds (Figure 5.8e). The shapes of the air speed distributions differ substantially between the two turns. The distribution shape for the right turn is right-skewed, while that for the left turn is near-symmetric. Note that the tether force coefficient is calculated with the airspeed squared, therefore, an airspeed difference leads to a relatively large tether force coefficient difference.

Figure 5.8e shows that the kite speed distributions for the turns exhibit two peaks. During the figure of eight, the speed of the kite is lowest before it starts turning. Subsequently, the speed peaks when the kite is pointing down. At the end of the turn, the kite is slightly decelerated but still flies at a higher speed than the speed at the valley of the distribution. In conclusion, the left and right peaks come from the start and end of the turn, respectively. The choice of the limits used for differentiating between straight paths and turns will greatly affect the shape of these distributions.

It is likely that the mismatch in kite speed between the left and right turns stems from a misalignment of the figure of eight with the wind direction discussed in Section 5.3. The kite can generate the most speed close to the direct downwind position. Both the wind misalignment and the kite speed mismatch cause a discrepancy in the air flow experienced by the kite. If the figure of eight were aligned with the wind, the same kite speed in both turns is expected given that they are controlled identically. A misalignment of the figure of eight can accumulate to large air speed differences.

Figure 5.9 depicts the joint distributions of some paired measurements to evaluate their relationships. For readability, only the distributions of the right and left turns are plotted, not those of the straight sections. Figure 5.9a shows the air speed-tether force distributions together with the expected quadratic trends. Variations in the vertical direction around this trend indicate variations in the tether force coefficient. The distribution for the right turn (partly) agrees with the quadratic trend, which suggests that the tether force coefficient is independent of the airspeed. In contrast, the tether force for the left turn scales with the airspeed to the power of 2.5 and thus suggests that the tether force coefficient increases with the airspeed.

Similar to the lift coefficient, the tether force coefficient is expected to vary with the flow angle. Indeed, Figure 5.9b shows that the flow angle and the tether force coefficient are weakly correlated for the left turns. Moreover, the flow angle tends to increase with the airspeed for the left turns as shown in Figure 5.9c. Combined, these observations could explain the high exponent with which the tether force scales with the airspeed for the left turns. Nevertheless, the latter relationships are not strong enough to draw a rigid conclusion.

The observed relationships for the left turns do not hold for the right turns. The distribution of the flow angle and the tether force coefficient for the right turn exhibits two peaks (Figure 5.9b) as opposed to one. The lowest peak of the two is stretched horizontally, indicating that there is no relationship between the tether force coefficient and the flow angle. Also, the correlation between the flow angles and airspeed is even weaker. The corresponding circular distribution is similar to a random distribution around a single point.

The most pronounced relationship in the measurements is found between the tether force and reel-out speed, shown in Figure 5.9e. The right and left turn distributions closely follow the same linear trend. This relationship is imposed by the winch controller as it tracks the ratio between the reel-out speed and the tether force during the reel-out phase.

5.5.3. Characterising the kite aerodynamics

The aerodynamic model is an important aspect of a flight operational model as it greatly affects the solution of the motion of the kite. These models commonly build upon the relationship between the angle of attack and the aerodynamic coefficient. To this end, a general relationship between the aerodynamic coefficients and the angle of attack is sought in the flight data. Ideally, a relationship can be identified that holds for both the powered and depowered kite.

Figure 5.10 depicts the relationship between the aerodynamic coefficients identified with the flight data and the angle of attack using joint distributions. Separate distributions are provided for the reel-out and reel-in phases plotted with different colours (Figure 5.10a and f). Additionally, distributions for each subset of the reel-out phase are provided to identify different flight behaviour between the different sections of the figure of eight. These are plotted on top of the reel-out and reel-in distributions to highlight possible differences.

For the reel-in phase, the distributions are more concentrated than those for the reel-out phase. The lift coefficient distribution during reel-out exhibits the least affinity with a linear relationship. The lift coefficient distribution for the right turns is the most apparent outlier and shows multiple peaks and relatively low values with respect to the other

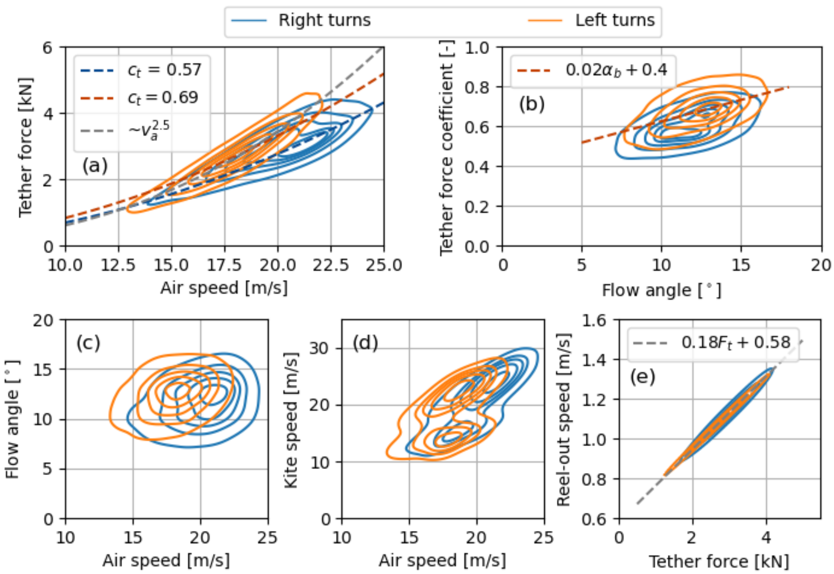

Figure 5.9: Joint probability density distributions showing a selection of pairwise relationships in the flight data for the turns during the reel-out phase. The plotted levels correspond to the 10%, 30%, 50%, 70%, and 90% iso-proportions of the density. The dashed lines show the fits to the data of each turn.

distributions. Also, the distribution for the left-to-right straight section differs from that of the reel-out phase as it peaks at lower angles of attack. Although the lift coefficient distribution for the left turns looks average, the drag coefficient distribution is an outlier as it exhibits a tail that extends to high values at high angles of attack. Again the peak of the distribution for the left-to-right straight section is shifted. This time not only to lower angles of attack but also to lower drag coefficients.

The attitude of the wing and flow measurements affect the identified lift coefficients but not the tether force coefficient. Consequently, the lift coefficient distribution for right turns in Figure 5.10b is differently shaped than the tether force coefficient distribution in Figure 5.9b. The additional mode in the lift coefficient distribution for right turns is most likely caused by the measured tether force and the airspeed input and not the identification routine since also the tether force coefficient distribution for the right turn exhibits two modes.

The distributions for the right-to-left straight section (Figure 5.10d and i) are more concentrated than those for the other reel-out flight sections. As such, the linear fits to these distributions are considered the most representative of the whole reel-out phase. The linear fits for the reel-in phase are directly obtained from the reel-in distributions. The predictive strengths of these relationships are quantified with the R^2 metric listed in Table 5.2.

Both reel-in distributions show a relatively strong dependence on the angle of attack: 38% and 63% of the variation can be explained by the angle of attack for the lift and drag coefficients, respectively. The same holds for the drag coefficient during reel-out: 38%

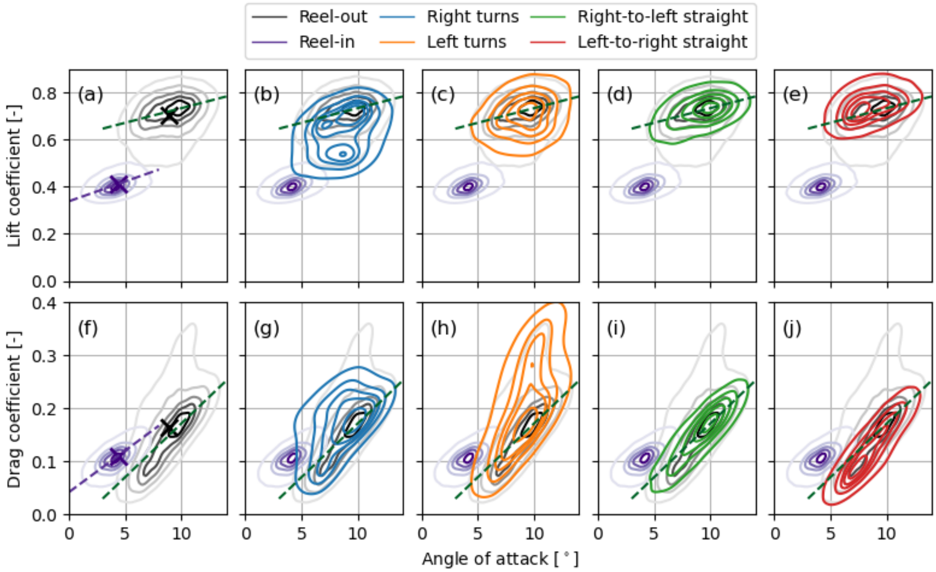

Figure 5.10: Joint probability density distributions showing the relationships between the identified aerodynamic coefficients using the steady single point mass model and the angle of attack for different flight sections. The plotted levels correspond to the 10%, 30%, 50%, 70%, and 90% iso-proportions of the density.

of the variation is explained by the angle of attack. In contrast, the relationship between the lift coefficient and angle of attack is weak: only 5% of the variation can be attributed to the angle of attack. Since the relationships for the reel-out phase are derived from the distributions of the right-to-left straight flight section, their R^2 values tend to be relatively high compared to the other reel-out flight sections. The negative R^2 values indicate that the variance around the mean coefficient is smaller than around the fits.

The linear relationships found for the reel-out and reel-in phases do not nicely align. This observation suggests that the coefficients of the powered and depowered wing cannot be captured with a single linear relationship. This is not unexpected, as even the lift curve for conventional airfoils is only linear at small angles of attack. The fluid-structure interaction of the membrane wing is likely to cause some non-linear effects during the reel-out phase. Therefore, making a linear fit to the reel-out data may not be the best choice. Moreover, the linear relationship found for the lift coefficient during reel-in does not align with the reel-out distribution. In other words, these observations suggest that the lift curve is not linear at small angles of attack.

The small R^2 values found suggest that more predictors than just the angle of attack are needed to predict the aerodynamic coefficients. Assessing the relationship with the side slip angle would be a logical follow-up. However, the side slip is not measured in the analysed test flight and would require a new test flight. Alternatively, the relationship with the steering input is evaluated.

Figure 5.11 depicts the relationship between the aerodynamic coefficients identified with the flight data and the steering input for the four flight sections of the reel-

Phase	Flight section	R^2 values	
		C_L	C_D
Reel-out		0.05	0.38
	Right turn	-0.76	-0.03
	Left turns	-0.02	-0.04
	Right-to-left straight	0.18	0.74
	Left-to-right straight	-0.08	0.61
Reel-in		0.38	0.63

Table 5.2: R^2 values of the identified coefficients with respect to the fits illustrated with the dashed lines in Figure 5.10.

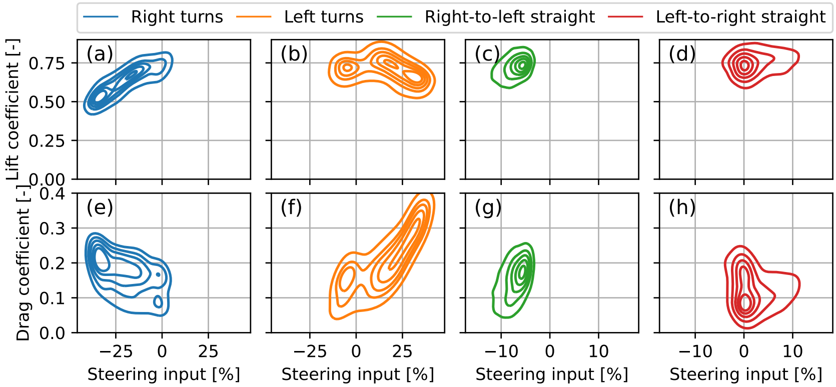

Figure 5.11: Joint probability density distributions showing the relationships between the identified aerodynamic coefficients using the steady single point mass model and the steering input for different flight sections in the reel-out phase. The plotted levels correspond to the 10%, 30%, 50%, 70%, and 90% iso-proportions of the density.

out phase. As expected, during the straight sections not much steering is applied and, consequently, the distributions do not indicate relationships between the aerodynamic coefficients and steering input. In contrast, the distributions for the turns do indicate relationships between the aerodynamic coefficients and steering input.

Steering tends to decrease and increase the lift and drag coefficient, respectively, as observed in the distributions for the turns. The distributions for the left turn (Figure 5.11b and f) are clearly bi-modal with a circular part below zero steering input and an elongated part for positive steering input. The circular part is similar to the random distributions for the right-to-left straight sections, whereas the elongated part suggests a near-linear positive relationship. The distributions for the right turn (Figure 5.11a and e) exhibit similar features with an overlap with the left-to-right straight sections, albeit less pronounced. For the right turns, the lift coefficient is lowered more due to steering than for the left turns. This explains the lower peak in Figure 5.10b. For the left turns, the two steering modes have similar lift coefficients and do not produce two modes in Figure 5.10c.

To conclude, the misaligned aerodynamic coefficient trends for the powered and depowered kite imply that the lift and drag coefficients should not be considered as functions of only the angle of attack in the aerodynamic model. Moreover, the angle of attack range covered in the test flight data is insufficient to establish separate aerodynamic models for the powered and depowered kite separately. To better reflect reality, the depower signal should be considered explicitly. Therefore, an aerodynamic model based on a linear relationship between the aerodynamic coefficients and the depower signal is proposed as a workable alternative and is evaluated in the following sensitivity study. The relationship with the steering input is not considered in the model. The mean values of the aerodynamic coefficients marked in Figure 5.10a and f are used as values for the powered and depowered kite in between which is interpolated.

5.6. Sensitivity of the steady state power estimation

The effect of the aerodynamic coefficients on the performance estimate is investigated with the steady flight state approximation. Determining the power output requires the reel-out speed as input. According to the straight and rigid tether assumption of the model, the reel-out speed equals the radial component of the kite velocity. The reel-out speed is thus a component of the kite velocity. Multiplying the reel-out speed by the tether force yields the power output. The power estimation is thus explicitly dependent on the kite velocity, whereas the force equilibrium in the previous aerodynamic identification only implicitly depends on the wind velocity through the apparent wind velocity.

By prescribing the tether force, the model developed by Schmehl et al. [102] finds the magnitude of the apparent wind velocity yielding a force equilibrium given the set value of the tether force at the ground station, the kite position, and the aerodynamic coefficients. Note that the previous identification works the other way around; the apparent wind velocity was fixed to find the coefficients. The model does not differentiate between the lift and side forces, therefore, the lift coefficient used in the model is denoted as $C_{L'} = \sqrt{C_L^2 + C_S^2}$. The direction of the apparent wind velocity and the kite velocity is determined by the wind velocity and the course angle input. With the other input fixed, multiple combinations of the wind velocity and course angle input exist that output the same apparent wind velocity.

Only if the radial component of the measured kite velocity equals the measured reel-out speed and the modelled and recorded kite and apparent wind velocity are the same, the calculated power will match the the measured power. This is unlikely to be the case as the tether is elastic and it is tilted upwind at the kite. The modelled kite and apparent wind velocity will match the measured velocities when using the previously reconstructed instantaneous wind velocity and the identified aerodynamic coefficients. Although using the reconstructed wind velocity ensures consistency with the flight data, the reconstructed wind velocity is not very accurate and, therefore, not considered a good option for assessing the power estimate. Although inferring the instantaneous wind at the test site from ERA5 wind data is also not very accurate, ERA5 is considered a better option for providing the wind input.

The sensitivity of the steady state power estimation is demonstrated by determining the evolution of the power output along the reference cycle for three cases with different aerodynamic coefficients and course angle input, see Table 5.3. First, the previously identified instantaneous aerodynamic coefficients are used (case 1). Next, the aerodynamic coefficients are modified to find a better match with the measured power (case 2). Finally, the aerodynamic coefficients are obtained from a simple aerodynamic model (case 3).

To verify the identification in Section 5.5.1, case 1 employs the steady flight state approximation to reproduce the earlier results. As depicted in Figure 5.12, a least squares optimiser is wrapped around the steady flight state approximation to find the aerodynamic coefficients and the course angle at each data record in the reference cycle. The objective of the optimisation is to match the modelled and measured wind velocities, see Equation 5.18 in Table 5.3.

Figure 5.13 shows the aerodynamic coefficients and course angles identified with the steady flight state approximation and related parameters. The aerodynamic coefficients exactly match the earlier identified coefficients. The difference between the computed and measured course angle fluctuates between approximately +/- 10 degrees. The corresponding reel-out speed and power mostly lie below the recorded values. The trapezoidal rule is used to approximate the energy output of the reference cycle based on the estimated power for each data record. The calculated mechanical energy is approximately only a quarter of the recorded energy in the flight data.

To evaluate where the large difference in calculated and recorded energy output comes from, case 2 employs a different optimisation objective in the identification of the aerodynamic coefficients and course angles. This time, only the magnitude of the apparent wind velocity is matched with the measurements and not the direction. Moreover, the reel-out speed and kite speed are matched, see Equation 5.19 (Table 5.3).

The resulting lift coefficients are similar to those found for case 1, whereas the drag coefficients slightly differ, see Figure 5.13. Substantial differences are found for the course angle, which may exceed 30 degrees. The reel-out speeds closely match those recorded in the flight data as imposed by the objective. Consequently, also the power and energy output match. Note that differences in the direction of the velocity are not penalised in the optimisation.

The large course angle differences required to match the calculated power with the measured power suggest that the steady flight state approximation is very sensitive to the course angle. Using the course angles recorded in the flight data would lead to underestimating the power output substantially. The observed pattern in the course angle mismatch suggests that the model deficiency cannot be attributed to the wind velocity input alone. It is expected that the discrepancy can be primarily attributed to the simplicity of the model.

For the third case, the power estimations are repeated using aerodynamic coefficients obtained from a simple aerodynamic model to quantify how such a model may impact a flight operational model. The coefficients are approximated as a linear function of the depower signal. The identified mean values of the coefficients for the reel-out and reel-in phase listed in Table 5.3 are used for the powered and depowered kite, respectively. Figure 5.13b shows that the corresponding coefficients approximately follow the same trend as the earlier identified coefficients. The recorded course angles from the flight data are used as input for the course angle parameter.

Figure 5.13d shows that the reel-out speed peaks in the middle of the turns for case 3.

Case	Input
1	$\min\limits_{C_{L'}, C_D, \chi} \; (\mathbf{v}_\mathrm{a} - \hat{\mathbf{v}}_\mathrm{a})^T \; (\mathbf{v}_\mathrm{a} - \hat{\mathbf{v}}_\mathrm{a})$ (5.18)
2	$\min\limits_{C_{L'}, C_D, \chi} \; (v_\mathrm{a} - \hat{v}_\mathrm{a})^2 + \left(v_\mathrm{k,r} - \hat{\dot{l}}_\mathrm{t}\right)^2 + (v_\mathrm{k} - \hat{v}_\mathrm{k})^2$ (5.19)
3	$\chi = \hat{\chi};\, C_{L'}, C_D = \begin{cases} 0.70, 0.16; & \text{if powered} \\ 0.41, 0.11; & \text{if depowered} \end{cases}$

Table 5.3: Input cases for evaluating the sensitivity of the power output of the steady flight state approximation.

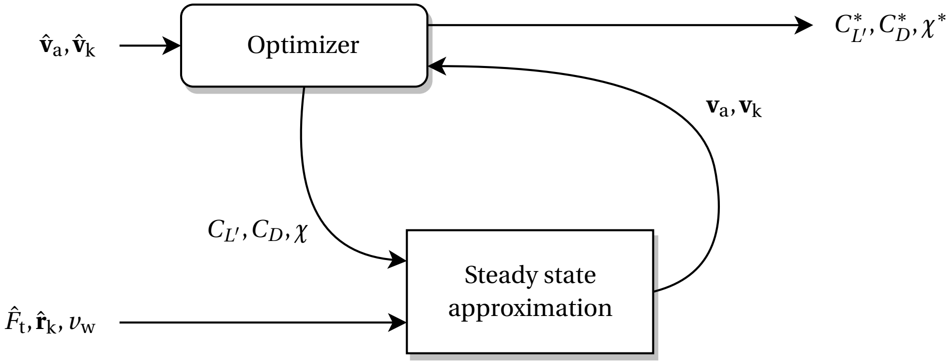

Figure 5.12: Optimiser wrapped around the steady flight state approximation for identifying the aerodynamic coefficients and course angle. The identification is conducted with multiple objectives which are functions of the apparent wind and kite velocities. The optimal parameters are denoted with asterisks.

This is explained by either a higher lift coefficient or a lower drag coefficient compared to the other cases. Also, the power output exhibits the same peaks while the power during the straight paths is only gradually changing. The power peaks during the turns result in a higher energy output than for case 1. Nevertheless, the energy output is still substantially lower than the measured energy output.

5.7. Conclusion

The wind velocity at the kite is reconstructed using the airflow and kite velocity measurements in an attempt to identify the vertical wind profile. The measurement error of the kite velocity and the zero side slip approximation introduce substantial errors during the turns. Consequently, the reconstruction during the straight paths and the reel-in are more reliable. The reconstruction suggests a lot more wind shear than is deemed realistic based on the profiles from ERA5. The reconstruction also suggests that the wind is slightly backing, i.e., turning counter-clockwise with height. This may result in asymmetric flight behaviour, as the flown figures of eight are aligned with the wind direction measured at the ground station. The reconstruction is believed to be too uncertain to identify the wind profile. Including side slip measurements could improve the reconstruction substantially.

The steady flight state approximation is used to infer the aerodynamic force gener-

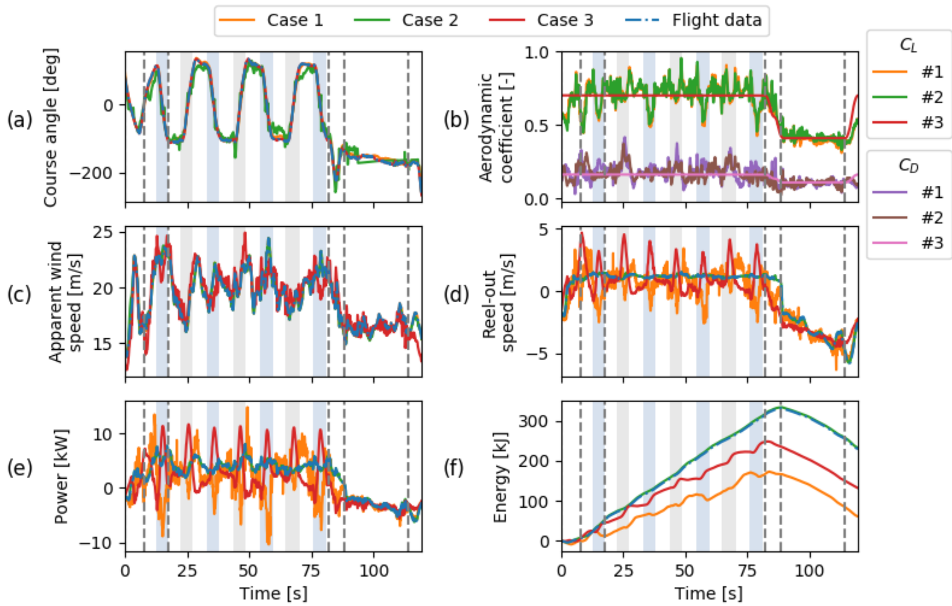

Figure 5.13: The results of the three cases used for studying the sensitivity of the steady state power estimation compared with flight data. (**a**–**b**) The course angle and aerodynamic coefficients input. (**c**–**d**) Apparent wind and reel-out speeds included in the identification objectives. (**e**–**f**) Corresponding power estimate and energy which follows from integrating the power.

ated by the kite. The approximation assumes a force equilibrium on the point mass at the position of the kite, which requires the forces acting on the tether to be lumped to the kite. Since the weight of the kite and tether is relatively small, the resulting aerodynamic force mostly counterbalances the tether traction force.

The attitude measurements of the kite are used to express the resultant aerodynamic force in the flow reference frame. Decomposing the force along the axes of the flow reference frame yields the lift, drag, and side force components. These forces are normalised to obtain the aerodynamic coefficients.

The comparison of the lift coefficient and normalised tether force suggests that most of the variation of the lift coefficient can be explained by the tether force and the apparent wind speed measurements. The statistical analysis of the measured quantities during the reel-out phase shows that the lower lift coefficients during the right turns can mostly be explained by the higher airspeeds as a result of higher kite speeds. One explanation for the higher kite speeds during right turns is that the figure of eight is not aligned with the wind. The kite is expected to fly closer to the direct downwind position during the right turns, where it can generate the most speed.

The joint distributions of the angle of attack and the lift and drag coefficients for the reel-out and reel-in phases exhibit misaligned trends. The linear relationship found for the lift coefficient during reel-in does not nicely align with the reel-out distribution. This suggests that the lift curve at small angles of attack is not strictly linear. Moreover, the drag coefficient distributions for the reel-out and reel-in phases cannot be captured with a parabolic drag curve. Examining the relationship between the steering input and the lift coefficient explains why the joint distribution of the angle of attack and the lift coefficient exhibits a second peak with low values. This peak can be attributed to a large steering input.

The implication to the aerodynamic model is that including the lift and drag coefficients as functions of only the angle of attack is inaccurate. Crucial parameters to incorporate explicitly in an aerodynamic model are the depower signal and steering input. As a workable alternative, an aerodynamic model based on a linear relationship between the aerodynamic coefficients and the depower signal is proposed.

The power output estimated with the steady flight state approximation directly depends on the wind velocity, whereas the force equilibrium only implicitly depends on the wind velocity through the apparent wind velocity. The ERA5 wind profile is used to obtain the wind speed input for evaluating the sensitivity of the power estimation. The sensitivity study shows that the calculated power is very sensitive to the course angle input and, to a lesser extent, to the aerodynamic coefficient input. Calculating the energy yield of a pumping cycle with slightly different course angle input may change the calculated yield with a factor of four.

6

Identifying the turning mechanism of a kite with suspended control unit

The steady flight state presented in the previous chapter does not account for the attitude of the kite. However, the turning of the kite and tether sag affects the attitude and, thereby, the system performance. The aim of this chapter is to investigate how the rotational motion of the kite and its turning mechanism can be efficiently modelled. Although the derived models are not directly applied in subsequent chapters, this chapter provides them with valuable insights for configuring and validating the flight operation model.

6.1. Introduction

The simpler models represent the kite as a single point mass or rigid body, assuming a straight tether with its mass and drag lumped to the kite. More refined models also resolve tether sag induced by lateral forces on the tether, such as gravity, centrifugal force, and aerodynamic drag. This can be conveniently done by discretising the tether. Typically, the tether is represented with lumped masses connected with rigid links or spring-damper elements [49, 53, 104–106]. Alternatively, Sánchez-Arriaga et al. [107] apply a multi-body approach using rigid rods. Fechner et al. [49] expand the discretisation approach to the kite. The kite is represented with five point masses; four point masses represent the wing, and one additional point mass represents the suspended KCU. A lumped-mass model with spring-damper elements yields a stiff system of differential equations. Solving this system requires a small time step and is prone to numerical instabilities. These models are considered too computationally costly for performance calculation but are used for control system design.

Parts of this chapter have been published in Wind Energy Science Discussions (2023)

To efficiently account for tether elasticity, Williams [48] solves the 'quasi-static' tether deformation as a subroutine to solving the motion of the kite. Consequently, the loaded tether shape due to gravity, centrifugal force and aerodynamic drag is considered, while the transient motion and longitudinal vibration are neglected. The discretised tether model assumes that the entire airborne system, including tether and kite, jointly rotates around the tether attachment point at the ground. This assumed kinematic relationship works well for near-straight flights but is not representative during turning manoeuvres.

The choice of the kite model determines the level of abstraction required to introduce steering forces as demonstrated in the work of Fechner et al. [49]. The work presents both a single-point and a five-point kite model of a LEI kite. By resolving the roll of the kite with respect to the upper tether segment, the five-point kite model allows for a realistic incorporation of the centripetal force acting on the relatively heavy control unit. Together with the kite, the lift force of the top wing surface may tilt into turns and pull the control unit along the same turn. Additionally, the lift forces of the wing tips contribute to the centripetal force. The lift coefficient of each wing surface depends on the local angle of attack and has a maximum of 1.1. The single-point model requires intricate centripetal force modelling because it lacks information about the attitude of the kite. To enable steering, it employs an artificial lateral force proportional to the steering input. The lift force of the top wing surface is assumed unaffected by this steering input and is approximately aligned with the upper tether segment. Due to this alignment, the lift force cannot exert a centripetal force on the control unit. This may explain why the lateral force coefficient of 2.59 is significantly higher than the lift force coefficient of the wing tip in the five-point kite model.

In reality, the deformation of the wing due to steering input is not as simple as suggested by the latter aerodynamic models. The LEI kite of Kitepower B.V. is steered by pulling the rear bridle lines attached to one side of the wing while loosening the lines on the other side. This asymmetric actuation of the bridle line system makes the wing deform and initiate a turn. Video footage of experiments shed some light on the aero-structural deformation due to steering [108]. Previous research on the topic has focused mainly on the interaction between the flow and the deforming bridled membrane wing [109–117]. The experimental data also indicates a pronounced dynamic interaction between the wing, the suspended KCU, and the tether during the turning manoeuvre. How much the swinging motion of the KCU relative to the wing affects the turning behaviour and the power generation of the kite has only recently been studied by Roullier [118]. An improved understanding of this effect would allow for enhancing performance models of flexible membrane kites, designing more precise control algorithms, and ultimately improving the system performance.

The goal of this paper is twofold: to study the dynamics that induce the observed characteristic pitch and roll swinging motion of the kite during sharp turning manoeuvres and discuss the implications to performance modelling. Pertaining to the first goal, this paper introduces a two-point kite model that is used together with a straight and discretised tether. Firstly, the motion is approximated as a transition through steady-rotation states with both tether representations. Subsequently, the motion is resolved dynamically with the discretised tether to study the impact of transient effects. Instead of resolving the translational motion of the wing, we prescribe a cross-wind flight path from the flight data of Kitepower B.V. This removes the dependency of the model on the aerodynamics of the kite and, thereby, reduces uncertainties. Pertaining to the second goal, this paper provides a breakdown of the mechanisms that initiate and drive a turn of a flexible kite system with a suspended control unit.

This chapter is organised as follows. Section 6.2 describes the cross-wind manoeuvre that is examined. In Section 6.3, the computational models are outlined. The results are presented in Section 6.4 and discussed in Section 6.5. Conclusions are presented in Section 6.6.

6.2. Cross-wind manoeuvre

The current analysis is illustrated using a figure-of-eight cross-wind manoeuvre of the wing shown in Fig. 6.1. This specific manoeuvre is part of the 65th pumping cycle of the test flight described in Chapter 2.5. Because of the high repeatability of the automatic flight manoeuvres, the other figures of eight of the dataset give similar results. Characteristic reference positions along this manoeuvre are designated to highlight the analysis, listed in Table 6.1. The kite flies along the trajectory in the direction of increasing reference numbers, i.e., flying upwards on the straight path segments and downwards during the turns. The tether is reeled out while the kite is flying cross-wind manoeuvres, increasing the radial position of the kite from 276 to 302 m at a height of 150–185 m. The asymmetry of the trajectory can be attributed to various factors, including misalignment with the wind velocity due to wind veer and imperfections within the system.

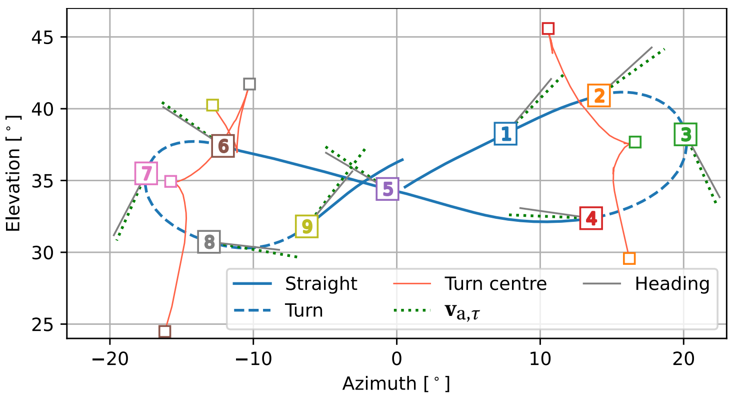

Figure 6.1: The studied figure-of-eight cross-wind manoeuvre of the wing depicted with respect to the wind reference frame, shown in Fig. C.1. The flight path comprises straight (solid blue) and turn (dashed blue) line segments. Reference positions 1 to 9 are designated along the path in flight direction. For the two turns, the changing position of the turn centre is tracked with the red lines. The turn-centre markers pair with the numbered path markers of the same colour. The dotted lines depict the modelled tangential apparent wind velocity. Alongside the apparent wind velocity lines, the solid lines depict the heading inferred from the attitude measurements of sensor 1.

Figure 6.2a–d depict the conditioned position data of the figure of eight. The position data is based on measurements of sensor 0, which have been processed using the default Kalman filter implementation of the Pixhawk®. The velocity measurements used in the

Table 6.1: Timestamps of the reference positions along the figure-of-eight path shown in Fig. 6.1, starting at 29.9 s and ending at 51.2 s in the 65th pumping cycle.

Instance label	1	2	3	4	5	6	7	8	9
Time [s]	31.9	33.9	35.6	37.5	41.0	44.5	46.2	47.6	49.1

present analysis come from the same sensor. The tangential and radial components of these measurements are depicted together with those measured by sensor 1 in Fig 6.2e–f (decomposition shown in Fig. 6.3). For an unknown reason, sensor 0 did not measure acceleration. Therefore, the acceleration measured with sensor 1 is used in the analysis and is depicted in Fig. 6.2g–i.

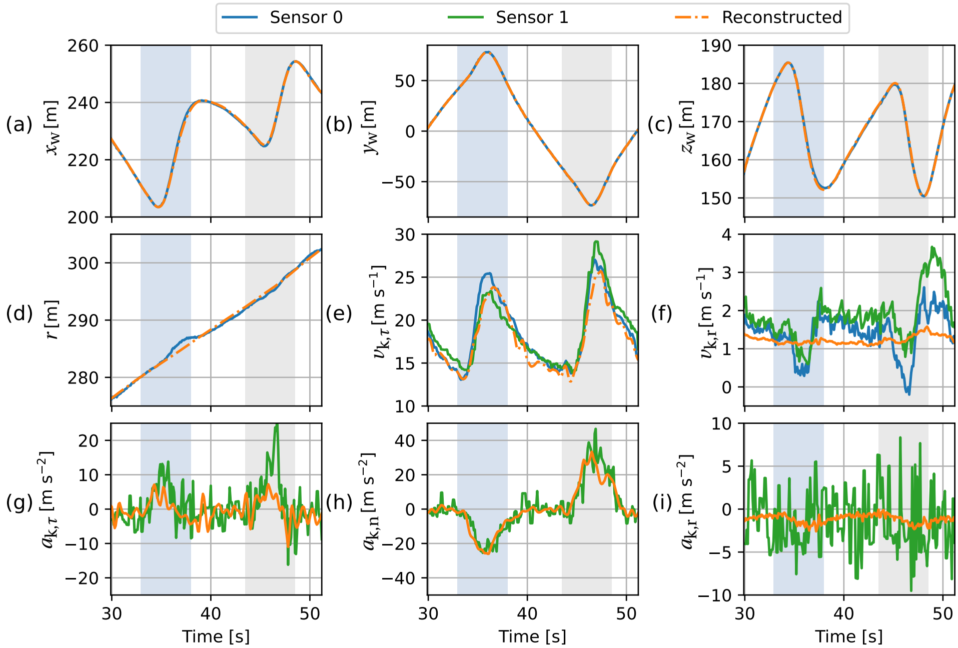

Figure 6.2: Kinematics of the studied figure-of-eight manoeuvre measured with the two Pixhawk® sensor units and the kinematics obtained with the flight trajectory reconstruction described in Appendix B. The intervals shaded blue and grey indicate right and left turns, respectively, from a downwind perspective. (**a**–**c**) Kite position coordinates of the wind reference frame (sensor data is Kalman filtered). (**d**) Radial position coordinate of the kite. (**e**–**f**) Tangential and radial kite velocity. (**g**–**i**) Tangential, normal, and radial kite acceleration.

Comparing the tether reel-out speed to the position of the wing indicates anomalies in the recorded wing position that manifest as unrealistically large jumps in radial position predominately occurring during right turns, as can be observed in Fig. 6.2d. These anomalies are removed using a discrete-time optimisation problem that minimises the error between the modelled radial wing speed and recorded tether reel-out speed while limiting the bias between the modelled and recorded wing position. The flight trajectory reconstruction might not be strictly valid. Nevertheless, it serves the higher aim of this study by providing a consistent kinematic input for the dynamic simulation. The identification of these anomalies and the optimisation details are described in Appendix B.

For simplicity, the present study assumes that the wind velocity is uniform and constant. The average wind speed measured at the ground for the reference pumping cycle is approximately 7 m s^{-1}. Based on the estimated wind shear, the wind speed at the kite is assumed to be 10 m s^{-1}. The grey lines in Fig. 6.1 show the heading of the kite at the reference positions inferred from sensor 1. The dotted green lines show the projection of the apparent wind velocity approximated with

$$\mathbf{v}_a = \mathbf{v}_\mathrm{w} - \mathbf{v}_\mathrm{k}\,, \tag{6.1}$$

in which $\mathbf{v}_\mathrm{w} = [10\,0\,0]^\top$ is the wind velocity and $\mathbf{v}_\mathrm{k}$ is the measured kite velocity. The side slip angle is the angle between the heading of the kite and the apparent wind velocity. The approximation of the apparent wind velocity lacks the necessary precision to assess the side slip. Moreover, the side slip angle was not measured during the flight test, and assessing the side slip is out of scope.

6.3. Modelling the motion of the tether and kite

The flight behaviour along the figure of eight described in the previous section is analysed with two different methods for solving the motion of the two-point kite model with a discretised tether model. First, this section discusses the tether-kite model configuration. Next, the two methods for solving the motion are discussed. The first approximates the tether-kite motion as a transition through steady-rotation states. The second solves the motion directly with dynamic equations of motion.

6.3.1. Tether-kite model

The two-point kite model accounts for the two distinct mass concentrations of the wing and the KCU. During cross-wind flight, the bridle line system is tensioned by the aerodynamic force acting on the wing. Accordingly, the two point masses stay at a constant distance, considering that the effect of wing actuation, including deformation, is negligible. From a modelling perspective, the two point masses at a constant distance are similar to a rigid body model, with rotational inertia in pitch and roll but not in yaw. The yaw motion is irrelevant to the present analysis due to the exclusion of the wing aerodynamics. This would not be the case when solving the full, unconstrained kite motion.

The two-point kite model developed for the present analysis can be added in a straightforward way to a discretised tether model as an additional final segment. An example with five tether segments of equal length l_j and a kite segment of length l_b is shown in Fig. 6.3.

To account for a varying length l_t and mass m_t of the deployed tether, the segment lengths and point masses are updated every instance according to

$$l_\mathrm{j} = \frac{l_\mathrm{t}}{N}, \tag{6.2}$$

$$m_\mathrm{j} = \frac{m_\mathrm{t}}{N}, \tag{6.3}$$

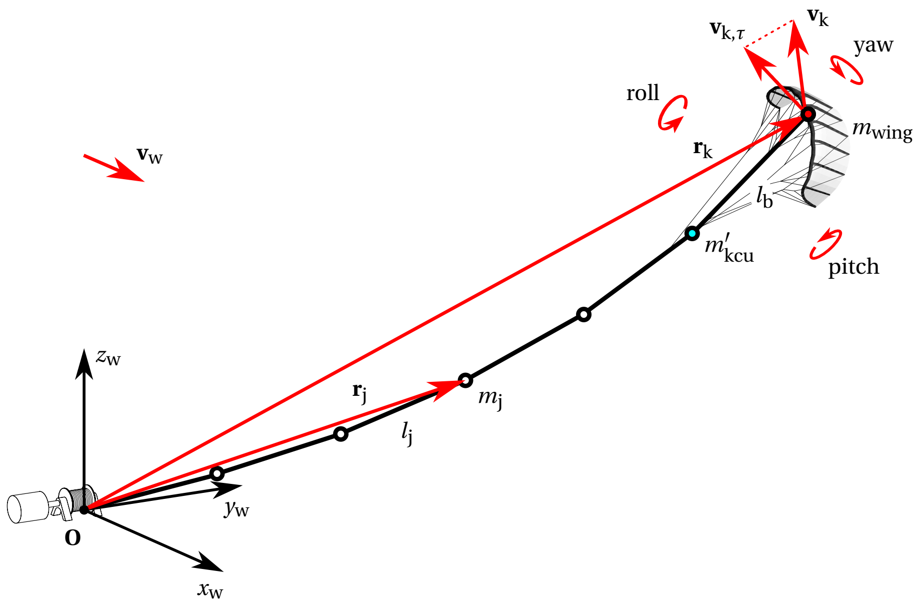

Figure 6.3: Two-point model of the kite added to a tether discretised by $N = 5$ tether elements. The position $\mathbf{r}_k$ and flight velocity $\mathbf{v}_k$ of the kite are defined as the position and velocity of the point **S** where the sensor units are attached to the wing, see Figure 4. Also shown is the tangential kite velocity component $\mathbf{v}_{k,\tau}$ (perpendicular to $\mathbf{r}_k$) and the wind reference frame x_w, y_w, z_w with origin at the tether attachment point **O** on the ground and x_w-axis aligned with the wind velocity vector.

where N is the constant number of tether elements. The point mass representing the KCU is determined as

$$m'_{kcu} = m_{kcu} + \frac{m_j}{2}. \tag{6.4}$$

The tether and bridle segments are assumed to be rigid. Moreover, variations in the lengths of these segments due to elasticity are neglected. The effect of tether elasticity on the swinging motion of the kite is expected to be negligible as long as the modelled tether length agrees with the effective real-world tether length.

Aerodynamic drag is one of the forces considered to act on the point masses representing the tether. The drag is calculated as

$$\mathbf{D}_{t,j} = \frac{1}{2}\rho \left\|\mathbf{v}_{a\perp,j}\right\| \mathbf{v}_{a\perp,j}\, C_{D,t}\, l_j\, d_t, \qquad \text{with} \quad j = 1, \ldots, N, \tag{6.5}$$

where ρ is the air density, $\mathbf{v}_{a\perp,j}$ is the local apparent wind velocity perpendicular the tether segment below the j^{th} point mass, $C_{D,t}$ is the tether drag coefficient, and d_t is the tether diameter.

Two aerodynamic forces are acting on the KCU point mass below the wing: the drag of the KCU itself $\mathbf{D}_{kcu}$ and half the drag of the upper tether element. Consequently, the total drag acting on the KCU point mass is

$$\mathbf{D}'_{kcu} = \mathbf{D}_{kcu} + \mathbf{D}_{t,kcu} = \frac{1}{2}\rho \left\|\mathbf{v}_{a\perp,kcu}\right\| \mathbf{v}_{a\perp,kcu}\, C_{D,kcu}\, A_{kcu} + \frac{\mathbf{D}_{t,N}}{2}, \tag{6.6}$$

in which $\mathbf{v}_{a\perp,kcu}$ is the perpendicular component of the apparent wind velocity at the KCU. The frontal area of the KCU is denoted as A_{kcu} and the drag coefficient as $C_{D,kcu}$. The chosen value of 1.0 for the drag coefficient is within the common range for a blunt body. The bridle and ram-air turbine drag are not included as separate terms but are considered accounted for by the KCU and wing drag. The values of physical parameters are listed in Table 6.2.

Table 6.2: Physical parameters of the airborne system model.

v_w	m_{kcu}	m_{wing}	l_b	ρ	d_t	$C_{D,t}$	A_{kcu}	$C_{D,kcu}$
$10\,\mathrm{m\,s^{-1}}$	25 kg	14.2 kg	11.5 m	$1.225\,\mathrm{kg\,m^{-3}}$	10 mm	1.1	$0.25\,\mathrm{m^2}$	1.0

Equation 6.5 does not account for any variation of the apparent wind velocity along the tether element and is only a reasonable approximation when using many tether elements. For single-element use, the alternative expression for the tether drag contribution (last term in Eq. 6.6) better preserves the moment of the tether drag around the ground station

$$\mathbf{D}_{t,kcu} = \frac{1}{8}\rho \left\|\mathbf{v}_{a\perp,kcu}\right\| \mathbf{v}_{a\perp,kcu} C_{D,t}\, l_t\, d_t\,. \tag{6.7}$$

6.3.2. Steady-rotation state

The subroutine for solving the ‘quasi-static’ tether shape proposed by Williams [48] is adopted in the present analysis to assess the swinging motion of the kite. With an initial guess of the tether length and orientation of the lower segment, the corresponding tether shape is determined using a shooting method. The positions of the point masses are determined one by one, starting with the lowest point mass and moving up towards the last point mass located at the tether end. From the pseudo force balance on a particular point mass (at the intersection of two tether elements), the position of the next point mass is inferred. This balance considers the tensile forces, drag, weight, and centrifugal force. Given the tensile force acting on the tether element below the point mass, only the tensile force acting on the tether element above remains unknown and is solved. The direction of this force dictates the axial direction of the corresponding tether element. Together with the length of a tether element, the axial direction yields the position of the next point mass. By repeating this calculation for each point mass, the position of the kite is obtained given the measured tether force at the ground. A least squares optimisation is employed to find the tether length and shape for which the upper tether end coincides with the position of the wing. Consult Williams [48] for more details.

To facilitate the calculation of loads, the velocities and accelerations of the point masses are approximated by assuming that they collectively rotate around the tether attachment point at the ground with a constant angular velocity $\boldsymbol{\omega}$, treating the point masses as particles lying on a rigid body. According to this kinematic assumption, the velocity and acceleration of each point mass depend solely on the angular velocity and its respective position. The velocity $\mathbf{v}_j$ and acceleration $\mathbf{a}_j$ for the j^{th} point mass are

$$\mathbf{v}_j = \boldsymbol{\omega} \times \mathbf{r}_j\,, \qquad \text{with} \quad j = 1,\ldots,N \tag{6.8}$$

$$\mathbf{a}_j = \boldsymbol{\omega} \times \mathbf{v}_j\,, \qquad \text{with} \quad j = 1,\ldots,N \tag{6.9}$$

where $\mathbf{r}_\mathrm{j}$ is the position of the point mass. This kinematic assumption is referred to as the steady-rotation assumption throughout this paper.

Prior to calculating the kinematics of the point masses, the angular velocity needs to be determined. Williams approximates the rotational velocity with

$$\boldsymbol{\omega}_\mathrm{straight} = \frac{\mathbf{r}_\mathrm{k} \times \mathbf{v}_\mathrm{k}}{\|\mathbf{r}_\mathrm{k}\|^2} = \frac{\mathbf{r}_\mathrm{k} \times \mathbf{v}_{\mathrm{k},\tau}}{\|\mathbf{r}_\mathrm{k}\|^2}, \tag{6.10}$$

in which $\mathbf{r}_\mathrm{k}$ and $\mathbf{v}_\mathrm{k}$ are the position and velocity of the kite, respectively, and $\mathbf{v}_{\mathrm{k},\tau}$ is the tangential component of the kite velocity, shown in Fig. 6.3. The resulting rotational velocity yields a rotation along a great circle on the surface of a sphere, as shown in Fig. 6.4. This rotational velocity is labelled as 'straight' because the great-circle rotation produces the straight path segments of a figure-of-eight manoeuvre. Note that this rotational velocity is perpendicular to the position and the (tangential) velocity of the kite, i.e., it points in the normal direction.

A shortcoming of this great-circle angular velocity approximation is that it does not yield an acceleration representative of a turning kite. Calculating the corresponding acceleration according to the steady-rotation assumption (Eqs. 6.8 and 6.9) will yield an acceleration that is aligned with the position vector and, thus, no lateral acceleration. The lateral acceleration, however, is important to consider as it is the dominant component during turns, as can be observed in Fig. 6.2h. The kinematic assumption does allow a lateral acceleration; however, this requires that the angular velocity has a radial component. Note that the steady-rotation assumption cannot produce a tangential acceleration.

The addition of a radial component to the great-circle angular velocity approximation enables producing a rotation along a small circle on the surface of a sphere coinciding with the turn of the figure-of-eight manoeuvre as shown in Fig. 6.4. Similar to the derivation of the normal angular velocity from Eq. 6.8, the radial angular velocity is derived from Eq. 6.9 and can be calculated with the normal component of the acceleration $\mathbf{a}_{\mathrm{k},n}$:

$$\boldsymbol{\omega}_\mathrm{r} = \frac{\mathbf{v}_{\mathrm{k},\tau} \times \mathbf{a}_{\mathrm{k},n}}{\|\mathbf{v}_{\mathrm{k},\tau}\|^2}. \tag{6.11}$$

The newly proposed rotational velocity approximation for turns reads as:

$$\boldsymbol{\omega}_\mathrm{turn} = \boldsymbol{\omega}_\mathrm{straight} + \boldsymbol{\omega}_\mathrm{r}. \tag{6.12}$$

The wing kinematics resulting from the flight path reconstruction are used to calculate the rotational velocity for turns. Figure 6.5a shows that the normal component of the turn rotational velocity is much smaller than the radial component. Figure 6.5b-c show the kinematics back-calculated with the steady-rotation assumption. The back-calculated wing velocity is solely produced by the normal component of the turn rotational velocity and only has a tangential component. Although the original wing velocity does have a radial component (smaller than 1.6 $\mathrm{m\,s^{-1}}$) and the back-calculated speed does not, their magnitudes are virtually the same. In contrast, the back-calculated wing acceleration is solely produced by the large radial component of the turn rotational velocity. The back-calculated wing acceleration also shows a very good match with the original wing acceleration despite the fact that it does not have a tangential component.

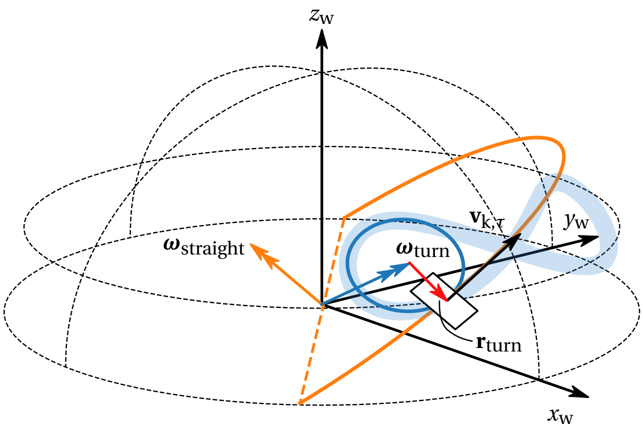

Figure 6.4: Two possible angular velocities, $\boldsymbol{\omega}_{\text{straight}}$ and $\boldsymbol{\omega}_{\text{turn}}$, that can be deduced from the tangential kite velocity $\mathbf{v}_{k,\tau}$. Their respective steady-rotation flight paths comprise a great circle (orange) and an instantaneous turn circle (blue) that approximately coincides with the turn of the figure-of-eight manoeuvre. The yawed tangential plane perpendicular to the position vector of the kite is depicted as a rectangle and represents the kite.

In conclusion, these results show that the steady-rotation assumption yields a very good approximation of the kite kinematics.

To conclude, we incorporate the following model modifications with respect to the model of Williams [48]:

- The elasticity of the tether elements is not considered as ;
- We add a radial component to the great-circle angular velocity;
- A different lumping approach is used for the uppermost tether point mass than for the other tether point masses, i.e., the mass and drag of half a tether element are allocated to the former instead of the mass and drag of a full element;
- We add an extra element (rigid link) to represent the kite as described in Section 6.3.1.

6.3.3. Dynamic equations of motion

The proposed dynamic model is a derivative of the generic model for multiple kite system architectures with fixed tether lengths introduced by Zanon et al. [53]. This model uses Cartesian coordinates to reduce the non-linearity of the model formulation. Although the model allows for complex systems, we only consider a simple single-tether, single-kite configuration. To limit the dimensions of the presented system of equations, we introduce a two-point kite model formulation with only two tether elements, in contrast to the 30 tether elements used for generating results. The first tether element connects the ground station to the only designated tether point mass m_1, and the second tether element connects m_1 to the point mass of the control unit m'_{kcu}, in a similar arrangement as the configuration depicted in Fig. 6.3.

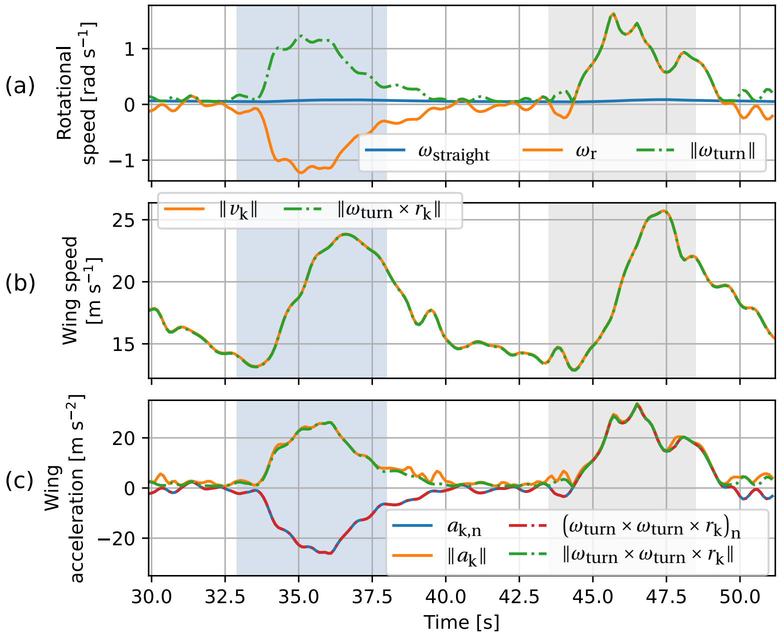

Figure 6.5: Assessing the steady-rotation assumption with the rotational velocity for turns. (**a**) The normal (straight-path) and radial rotational speed inferred from the reconstructed wing kinematics. (**b**, **c**) The wing speed and acceleration back-calculated with the turn rotational velocity (using Eqs. 6.8 and 6.9) compared to the wing speed and acceleration from the flight trajectory reconstruction used to calculate the rotational velocity. The shaded intervals indicate the turns.

The model is described by a differential-algebraic system of equations (DAE), with constraints originating from the use of non-minimal coordinates. The differential states $\mathbf{x}$, algebraic states $\mathbf{z}$, and control inputs $\mathbf{u}$ of the two-point model are

$$\mathbf{x} = [\mathbf{r}_1, \mathbf{r}_{\text{kcu}}, \mathbf{r}_{\text{k}}, \mathbf{v}_1, \mathbf{v}_{\text{kcu}}, \mathbf{v}_{\text{k}}, l_{\text{t}}, \dot{l}_{\text{t}}], \quad \mathbf{z} = [\mathbf{a}_1, \mathbf{a}_{\text{kcu}}, \lambda_1, \lambda_2, \lambda_{\text{b}}], \quad \text{and} \quad \mathbf{u} = [\mathbf{a}_{\text{k}}, \ddot{l}_{\text{t}}]; \quad (6.13)$$

in which subscript kcu refers to the kite control unit, k refers to the top wing surface of the kite, t denotes tether, b denotes bridle, and the numbers refer to the tether point masses and elements. The state variables are the positions and velocities of the point masses and the tether length and reel-out speed. The algebraic variables include the acceleration of the control unit point mass and Lagrange multipliers λ. The Lagrange multipliers enforce the constraints and have a close relationship with the forces acting in the tether and bridle elements. The control variables are the wing acceleration and the reel-out acceleration of the tether.

Without imposing the translational motion of the wing, the dynamics of the two-point kite model with two tether elements read as:

$$\begin{bmatrix} \begin{bmatrix} m_1\,\mathbf{I}_3 & \mathbf{0}_3 & \mathbf{0}_3 \\ \mathbf{0}_3 & m'_{\mathrm{kcu}}\,\mathbf{I}_3 & \mathbf{0}_3 \\ \mathbf{0}_3 & \mathbf{0}_3 & m_{\mathrm{k}}\,\mathbf{I}_3 \end{bmatrix} & \mathbf{G}_X^\top \\ \mathbf{G}_X & \mathbf{0}_3 \end{bmatrix} \begin{bmatrix} \mathbf{a}_1 \\ \mathbf{a}_{\mathrm{kcu}} \\ \mathbf{a}_{\mathrm{k}} \\ \lambda_1 \\ \lambda_2 \\ \lambda_{\mathrm{b}} \end{bmatrix} = \begin{bmatrix} \mathbf{D}_{\mathrm{t},1} - m_1\,g\,\mathbf{1}_z \\ \mathbf{D}'_{\mathrm{kcu}} - m'_{\mathrm{kcu}}\,g\,\mathbf{1}_z \\ \mathbf{F}_{\mathrm{a}} - m_{\mathrm{k}}\,g\,\mathbf{1}_z \\ -\mathbf{v}_1^\top\mathbf{v}_1 + \frac{1}{N^2}\left(\dot{l}_{\mathrm{t}}^2 + l_{\mathrm{t}}\,\ddot{l}_{\mathrm{t}}\right) \\ -\left(\mathbf{v}_{\mathrm{kcu}} - \mathbf{v}_1\right)^\top\left(\mathbf{v}_{\mathrm{kcu}} - \mathbf{v}_1\right) + \frac{1}{N^2}\left(\dot{l}_{\mathrm{t}}^2 + l_{\mathrm{t}}\,\ddot{l}_{\mathrm{t}}\right) \\ -\left(\mathbf{v}_{\mathrm{k}} - \mathbf{v}_{\mathrm{kcu}}\right)^\top\left(\mathbf{v}_{\mathrm{k}} - \mathbf{v}_{\mathrm{kcu}}\right) \end{bmatrix}, \tag{6.14}$$

in which

$$\mathbf{G}_X = \begin{bmatrix} \mathbf{r}_1 & \mathbf{0}_{1\times3} & \mathbf{0}_{1\times3} \\ \left(\mathbf{r}_1 - \mathbf{r}_{\mathrm{kcu}}\right)^\top & \left(\mathbf{r}_{\mathrm{kcu}} - \mathbf{r}_1\right)^\top & \mathbf{0}_{1\times3} \\ \mathbf{0}_{1\times3} & \left(\mathbf{r}_{\mathrm{kcu}} - \mathbf{r}_{\mathrm{k}}\right)^\top & \left(\mathbf{r}_{\mathrm{k}} - \mathbf{r}_{\mathrm{kcu}}\right)^\top \end{bmatrix}, \tag{6.15}$$

$\mathbf{I}_3$ and $\mathbf{0}_3$ are the 3×3 identity and zero matrix, respectively, $\mathbf{F}_{\mathrm{a}}$ is the aerodynamic force acting on the wing, g is the gravitational constant, and $\mathbf{1}_z = [0\ 0\ 1]^\top$. The equations of motion for the point masses are described in the upper three rows. The constraint equations described in the lower three rows represent the links between the point masses.

The constraint equations in the lower three rows of Eq. 6.14 are inferred from the constraints on the distances between linked point masses. The distance between the control unit and the top wing surface point masses is constrained by the constant bridle length l_{b}:

$$c_{\mathrm{b}} = \frac{1}{2}\left(\left(\mathbf{r}_{\mathrm{k}} - \mathbf{r}_{\mathrm{kcu}}\right)^\top\left(\mathbf{r}_{\mathrm{k}} - \mathbf{r}_{\mathrm{kcu}}\right) - l_{\mathrm{b}}^2\right) = 0\,. \tag{6.16}$$

The relative distances between the remaining linked point masses are constraint by the instantaneous tether length l_{t}:

$$c_1 = \frac{1}{2}\left(\mathbf{r}_1^\top\mathbf{r}_1 - \left(\frac{l_{\mathrm{t}}}{N}\right)^2\right) = 0 \tag{6.17}$$

and

$$c_2 = \frac{1}{2}\left(\left(\mathbf{r}_{\mathrm{kcu}} - \mathbf{r}_1\right)^\top\left(\mathbf{r}_{\mathrm{kcu}} - \mathbf{r}_1\right) - \left(\frac{l_{\mathrm{t}}}{N}\right)^2\right) = 0\,. \tag{6.18}$$

These constraints are differentiated twice to yield an index-1 DAE, enabling more efficient integration. As a consequence of the index reduction, the tether length acceleration and the accelerations of the point masses appear in the constraint equations. The initial states must satisfy two consistency conditions per constraint to ensure consistent kinematics of the tether and point masses in the simulation. The original expressions for the constraints serve as consistency conditions. Moreover, the time derivatives of these expressions are required as consistency conditions:

$$\dot{c}_{\mathrm{b}} = \left(\mathbf{r}_{\mathrm{k}} - \mathbf{r}_{\mathrm{kcu}}\right)^\top\left(\mathbf{v}_{\mathrm{k}} - \mathbf{v}_{\mathrm{kcu}}\right) = 0\,, \tag{6.19}$$

$$\dot{c}_1 = \mathbf{r}_1^\top\mathbf{v}_1 - \frac{l_{\mathrm{t}}\,\dot{l}_{\mathrm{t}}}{N^2} = 0 \tag{6.20}$$

and

$$\dot{c}_2 = \left(\mathbf{r}_{\mathrm{kcu}} - \mathbf{r}_1\right)^\top\left(\mathbf{v}_{\mathrm{kcu}} - \mathbf{v}_1\right) - \frac{l_{\mathrm{t}}\,\dot{l}_{\mathrm{t}}}{N^2} = 0\,. \tag{6.21}$$

To prevent inaccuracies of an aerodynamic model of the wing from interfering with the simulation, we do not resolve the dynamics of the point mass of the wing. Instead, the acceleration of the wing is prescribed and used as input. The wing acceleration is inferred from a cross-wind flight path from the flight data of Kitepower B.V., as described in Appendix B. Consequently, the equation of motion of the wing (third row in Eq. 6.14) becomes redundant and is dropped for this analysis:

$$\begin{bmatrix} \begin{bmatrix} m_1\,\mathbf{I}_3 & \mathbf{0}_3 \\ \mathbf{0}_3 & m'_{\mathrm{kcu}}\,\mathbf{I}_3 \end{bmatrix} & \mathbf{G}'^{\top}_X \\ \mathbf{G}'_X & \mathbf{0}_3 \end{bmatrix} \begin{bmatrix} \mathbf{a}_1 \\ \mathbf{a}_{\mathrm{kcu}} \\ \lambda_1 \\ \lambda_2 \\ \lambda_{\mathrm{b}} \end{bmatrix} = \begin{bmatrix} \mathbf{D}_{\mathrm{t},1} - m_1\, g\, \mathbf{1}_z \\ \mathbf{D}'_{\mathrm{kcu}} - m'_{\mathrm{kcu}}\, g\, \mathbf{1}_z \\ -\mathbf{v}_1^\top \mathbf{v}_1 + \frac{1}{N^2}\left(\dot{l}_\mathrm{t}^2 + l_\mathrm{t}\,\ddot{l}_\mathrm{t}\right) \\ -\left(\mathbf{v}_{\mathrm{kcu}} - \mathbf{v}_1\right)^\top \left(\mathbf{v}_{\mathrm{kcu}} - \mathbf{v}_1\right) + \frac{1}{N^2}\left(\dot{l}_\mathrm{t}^2 + l_\mathrm{t}\,\ddot{l}_\mathrm{t}\right) \\ -\left(\mathbf{v}_\mathrm{k} - \mathbf{v}_{\mathrm{kcu}}\right)^\top \left(\mathbf{v}_\mathrm{k} - \mathbf{v}_{\mathrm{kcu}}\right) - \left(\mathbf{r}_\mathrm{k} - \mathbf{r}_{\mathrm{kcu}}\right)^\top \mathbf{a}_\mathrm{k} \end{bmatrix} \tag{6.22}$$

in which $\mathbf{G}'_X$ is Eq. 6.15 with the third column removed. Moreover, the term with the wing acceleration in the algebraic equation of the bridle element is moved to the right-hand side.

Incorporating the acceleration of the control unit point mass as an algebraic state allows the DAE of the full model to be expressed in a semi-explicit form. The time derivatives of the differential states are

$$\dot{\mathbf{x}} = [\mathbf{v}_1,\, \mathbf{v}_{\mathrm{kcu}},\, \mathbf{v}_\mathrm{k},\, \mathbf{a}_1,\, \mathbf{a}_{\mathrm{kcu}},\, \mathbf{a}_\mathrm{k},\, \dot{l}_\mathrm{t},\, \ddot{l}_\mathrm{t}] \tag{6.23}$$

and the algebraic equations are obtained by rearranging Eq. 6.22. The DAE is solved with the IDAS integrator in CasADi [119]. IDAS employs the backward differentiation formula (variable-order, variable-coefficient) for implicit integration to solve the system. The motion is resolved at a fixed time step of 0.1 s. The solver produces a consistent simulation with insignificant drift in the consistency conditions, i.e., the distance between the wing and the KCU drifts with 0.0001 m in 24.2 s.

In contrast to the steady-rotation state calculation in Sec. 6.3.2, drag is calculated directly with the local apparent wind velocity $\mathbf{v}_{\mathrm{a,j}}$ instead of its normal component $\mathbf{v}_{\mathrm{a}\perp,\mathrm{j}}$ (Eqs. 6.5, 6.6, and 6.7) to limit the non-linearity of the model. To sum up, we incorporate the following model modifications with respect to the work of Zanon et al. [53]:

- The tether length time derivatives are added to the dynamic equations to enable modelling pumping AWE systems;
- Drag is computed directly at the point masses instead of being computed at the centres of the tether elements and then lumped to the adjacent point masses;
- The acceleration of the top of the kite (wing point mass) is not solved for. Instead, the wing acceleration inferred from measurements is directly imposed;
- Also here, we add an extra element (rigid link) to represent the kite as described in Section 6.3.1.

6.4. Induced swinging motion of the kite

Firstly, the steady-rotation-state approximation is used to study the motion of the tether and kite along the figure-of-eight manoeuvre. A discretisation by 30 tether segments is compared with a minimal discretisation using only a single tether segment. Secondly, the motion is simulated with the dynamic model using 30 tether segments. Subsequently, the resulting roll and pitch along the figure of eight from the different models are compared with measurements. Finally, the motion of the tether and kite along a full pumping cycle is studied.

6.4.1. Tether-kite lines computed with steady-rotation states

The steady-rotation-state approximation uses the measured tether force, wing position, and optimised angular velocity to determine the instantaneous positions of the point masses. The line formed by the segments between these point masses is referred to as the tether-kite line. Figure 6.6 shows the resulting tether-kite lines with 30 tether elements at the reference instances.

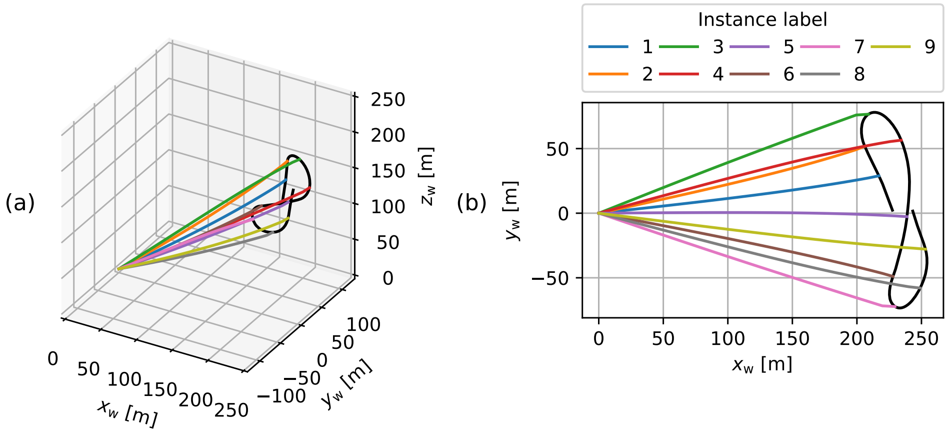

Figure 6.6: Tether-kite lines for the nine reference instances resulting from the steady-rotation-state approximation with the tether discretised by 30 segments in 3D (**a**) and top-view (**b**).

Variations in the deformation of the tether-kite line are hard to identify with the naked eye in the previous plots. Therefore, the cross-axial displacement is plotted against the radial position for the first five reference instances with the solid lines in Fig. 6.7. The displacement is expressed with respect to the tangential apparent wind velocity of the kite. The largest displacements are found in the down-apparent-wind direction, which can be attributed to the tether drag. The direction in which gravity contributes to the displacement varies depending on the position along the figure of eight. Table 6.3 specifies in which direction gravity acts for the first five reference instances. For all instances except for the third, gravity acts in the down-apparent-wind direction. The cross-apparent-wind displacement contribution of gravity changes sign after the third instance. Finally, the resistance to turn, or the inertia, mostly contributes to the displacement in the positive cross-apparent-wind direction, as can be inferred from the high positive values in the last column of Table 6.3.

The discontinuities in the tether-kite lines at the KCU indicate that it has a substan-

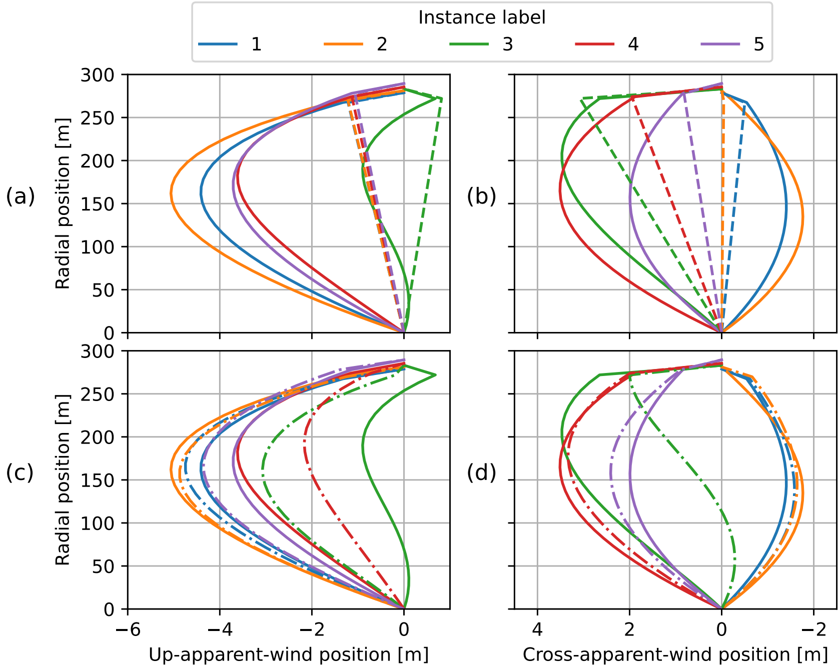

Figure 6.7: Tether-kite lines with cross-axial displacement decomposed with respect to the tangential apparent wind velocity of the wing (see Fig. 6.1). Steady-rotation states with 30 tether elements (solid lines in **a**, **b**, **c**, and **d**), with a single tether element (dashed lines in **a** and **b**), and the dynamic solution with 30 tether elements (dash-dotted lines in **c** and **d**) for the first five reference instances. Note that the x- and y-axes have different scales and that the x-axes are flipped in the second column.

tial effect on the attitude of the kite element. The high mass and drag lumped to the KCU point relative to the mass and drag lumped to the tether points cause these discontinuities.

To illustrate the imposed kite attitude more clearly, it is quantified using the pitch and roll of the kite element with respect to the tangential plane (perpendicular to the position vector of the kite). The exact definitions are given in Appendix C. Figure 6.8a shows that the pitch is roughly constant during the straight flight path sections and drops below zero during the turns (blue line). The negative pitch is confirmed by the tether-kite line plot of the 3$^{\text{rd}}$ instance in Fig. 6.7a, where the upper kite element is tilted backwards. Note that this depiction changes when plotting the tether-kite line with respect to, e.g., the vertical instead of the apparent wind velocity. The KCU is actually positioned higher than the wing and can be considered to be pulled along by the wing.

Figure 6.8b shows a distinct pattern for the roll of the kite along the figure of eight (blue line). The roll is slightly negative, roughly constant at the first straight section flying to the right, whereas it is slightly positive at the subsequent straight section flying to the left. In between, during the right turn, the roll peaks in the middle of the turn at 36.2 s. The left turn shows an opposite pattern. Note that the model does not account

Table 6.3: The negated vertical unit vector $-\mathbf{1}_z$ and the negated centripetal unit vector $-\mathbf{e}_{\text{centripetal}}$ decomposed in the up-apparent-wind and cross-apparent-wind direction experienced by the wing. The centripetal unit vector is determined by the approximated centripetal acceleration at the kite $\mathbf{e}_{\text{centripetal}} = \frac{\boldsymbol{\omega}_{\text{turn}} \times (\boldsymbol{\omega}_{\text{turn}} \times \boldsymbol{r}_{\text{k}})}{\|\boldsymbol{\omega}_{\text{turn}} \times (\boldsymbol{\omega}_{\text{turn}} \times \boldsymbol{r}_{\text{k}})\|}$. The listed fractions help to explain the contributions of gravity and turn inertia to the cross-axial displacement of the tether-kite lines in Fig. 6.7.

	$-\mathbf{1}_z$		$-\mathbf{e}_{\text{centripetal}}$	
Instance label	Up	Cross	Up	Cross
1	-0.56	-0.55	0.09	0.41
2	-0.44	-0.62	0.28	0.95
3	0.72	-0.32	0.25	0.97
4	-0.03	0.84	-0.23	0.96
5	-0.46	0.68	-0.27	0.84

for transient effects, which are expected to be substantial during the turns.

The rolling motion of the kite during the turns can be predominantly attributed to the resistance to turn, or inertia, of the KCU. The inertia of the tether has a much smaller effect on the roll. This stresses the need for including a separate point mass for the KCU when assessing the kite attitude.

The analysis is repeated using a single tether element. Figure 6.7a and b show the resulting tether-kite lines with the dashed lines. As expected, this minimal model is not able to give a good estimation of the maximum displacements. Nevertheless, the resulting kite elements align well with the results of the model with 30 tether elements. Figure 6.8 confirms this alignment as both the pitch and roll are similar for the two discretisations.

6

6.4.2. Cross-check with dynamic results

The dynamic simulation requires the wing acceleration, imposing the flight path, and the tether reel-out acceleration as input. The flight trajectory is reconstructed as described in Appendix B to ensure a running simulation and ensure that the inputs are consistent with the studied figure-of-eight manoeuvre. The intensive reconstruction yields a slightly adapted tether reel-out speed with respect to the measured speed and imposes a nearly constant difference between the tether length and radial kite position in the simulation. In this paper, we refer to this difference as tether slack. The initial tether length of the simulation is chosen such that the tether slack is 0.28 m, which is the mean value observed in the steady-rotation-state results.

Figure 6.9 shows the tether force evolution that results from the dynamic simulation. Since the force is sensitive to the choice of the tether slack, the agreement with measurements during the straight sections confirms that the choice for the constant tether slack is reasonable. During the turns, the calculated tether force does not agree well with the measurements. The simulated force shows distinct peaks, whereas the measured force shows a more gradual increase. These differences, however, are not specific to the dynamic model but are expected to be artefacts of the wing and tether acceleration control input.

The resulting tether-kite lines are plotted in Fig. 6.7c and d. Most shapes of the refer-

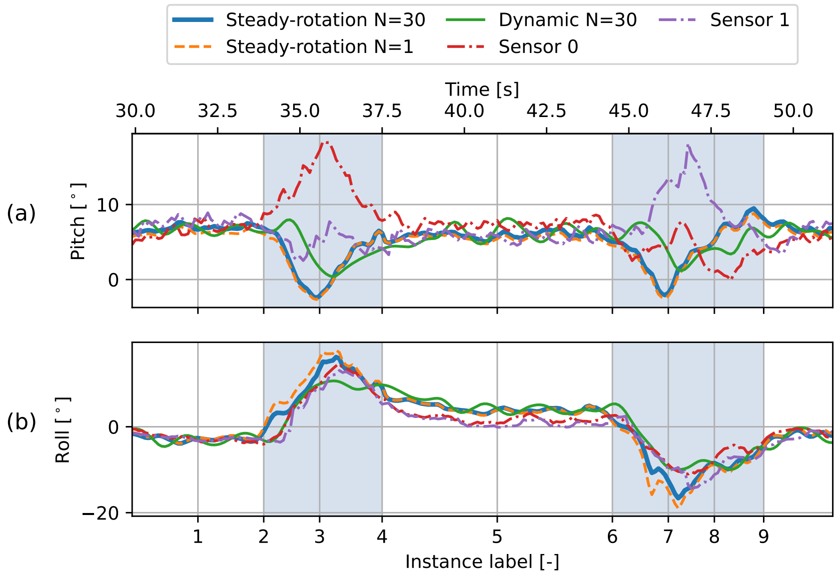

Figure 6.8: The pitch and roll of the kite derived from the attitude of the bridle element (with respect to the tangential plane) along the figure of eight. The results of the steady-rotation-state and dynamic analyses are depicted alongside the pitch and roll inferred from attitude measurements of the two sensors mounted to the wing, which include local effects of wing deformation. The shaded intervals indicate the turns.

ence instances show a reasonable agreement with the steady-rotation-state results. An apparent outlier is the 3$^{\text{rd}}$ reference instance, which occurs at the outside of the turn. This discrepancy can also be observed in Fig. 6.8a, in which the pitch resulting from the dynamic simulation closely follows the steady-rotation-state model results, except for the middle of the turn.

Dynamic models are necessary to account for the transient effects on the tether-kite line, which arise from the highly dynamic flight behaviour during turns. These transient effects are likely to explain why in Fig. 6.7d the lower end of the tether of the 3$^{\text{rd}}$ reference instance still has a negative cross-apparent-wind displacement like its predecessor, while the corresponding steady-rotation-state result is positive over the full length. Note that the current dynamic model does not necessarily enhance accuracy by considering transient effects as it requires different assumptions, e.g., on the tether reel-out acceleration, with associated uncertainties.

6.4.3. Kite attitude validation

The available measurements useful for validating the motion of the tether-kite line are the wing attitude measurements. These allow for estimating the actual pitching and rolling motion of the kite and, thereby, can help with validating the models. Validating the rotational motion of the kite is particularly important for performance model development, as accurate descriptions of this motion are essential for incorporating the aerodynamics and the turning mechanism. The tether motion cannot be validated as no

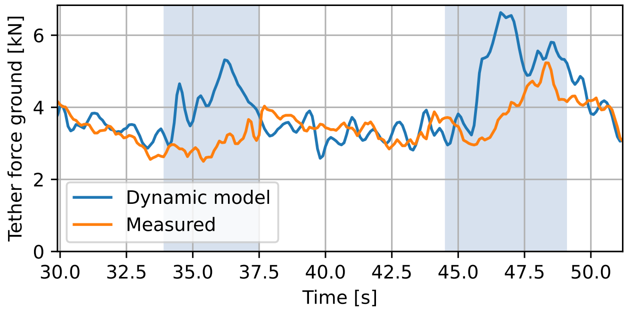

Figure 6.9: Tether force evolution along the figure of eight resulting from the dynamic simulation and from the flight data. The shaded intervals indicate the turns.

measurements are taken directly from the tether. Validating the tether motion is considered less important for performance model development.

Figure 6.8 compares the modelled pitch and roll angles of the kite element with measurements from two different sensors mounted to the wing. The same pitch and roll definitions are used to express the wing attitude measurements, provided in Appendix C. The kite attitude is inferred from these measurements by assuming that the kite is fully rigid and that the orientation of the wing relative to the bridle is defined by the depower angle α_d shown in Fig. 2.13. Moreover, the measurements are corrected for misalignments with the wing reference frame. 7° is added to the measured pitch of both sensors to correct for the sensor misalignment. Similarly, 8.5° is subtracted from the roll of both sensors to correct for sensor misalignment.

6

Both sensors measure a similar roll along the whole figure of eight, as shown in Fig. 6.8. However, the pitch measured with the two sensors differs substantially during the turns. Investigating the root cause revealed a strong relationship between the difference in pitch and the steering input. Fig. 6.10c illustrates their relation within the 65$^{\text{th}}$ pumping cycle with a Pearson correlation coefficient of -0.96. A steering input causes the steering tape to pull in on one side and give slack on the other. As a result, the wing twists around the leading edge with a zero twist at the centre. The high correlation found suggests that the twist between the struts on which the sensors are mounted is measured with high precision. The pitch at the centre of the wing is assumed to be the average of the two measurements.

Figure 6.8 shows that the differences in pitch and roll resulting from the models and the measurements are small during the straight sections. The computed pitch and roll angles match the measurements within three degrees. Contrastingly, the two models exhibit systematic differences during the turns. In particular, the pitch exhibits larger differences during the turns. Although the dynamic result lies closer to the average measured pitch during the turns, it does not exhibit a similar peak. This discrepancy could be attributed to the high uncertainty of the position measurement during the turns, resulting in large modifications to the flight trajectory by the reconstruction. Thereby, the actual wing motion that is causing the peak in pitch might have gotten lost in the reconstruction or was not properly measured. Moreover, the steady-rotation states might not

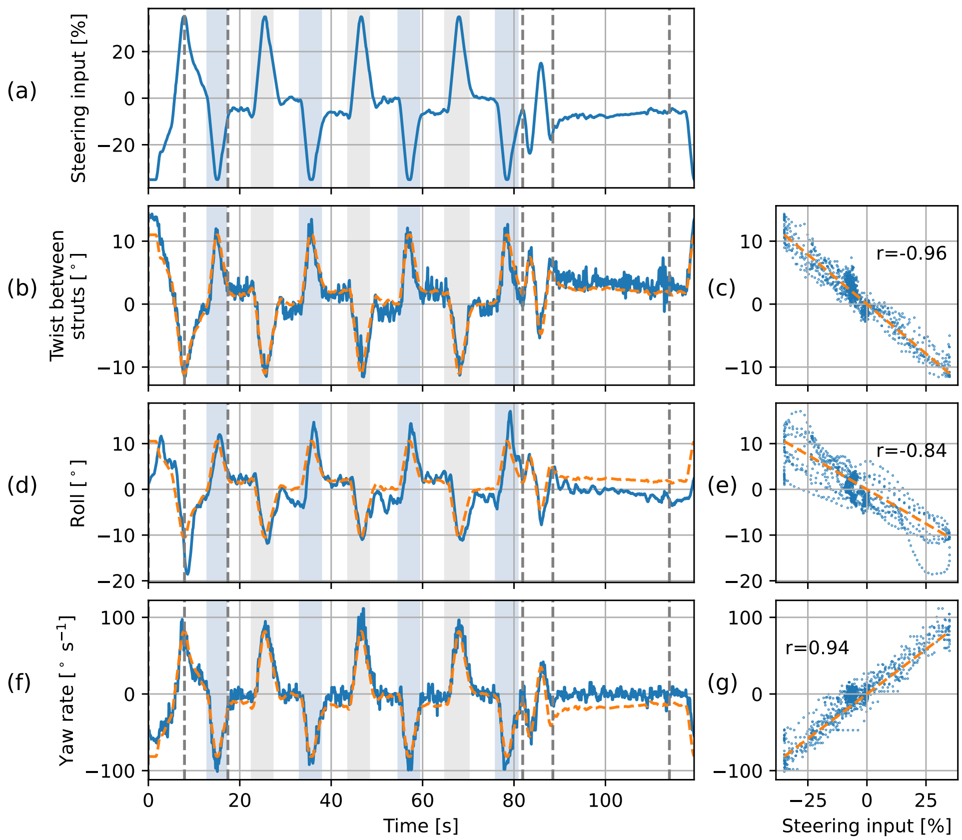

Figure 6.10: Relations between (**a**) the steering input and (**b, c**) the difference in pitch of the two sensors, (**d, e**) roll of the kite, and (**f, g**) yaw rate of the kite in the 65th pumping cycle. The orange dashed lines in the left column depict the steering input scaled with the slope found in the linear fit shown with the orange dashed lines in the right column.

accurately capture the kite attitude during turns because they do not consider transient effects.

In general, the steady-rotation states perform reasonably well in estimating the kite attitude, both with a single tether element and 30 tether elements. This suggests that both discretisations can capture the inertial effect of the KCU during turns. The dynamic model does not necessarily produce more accurate results than the steady-rotation-state model. This can be explained by inaccuracies in the input causing errors, e.g., due to imperfections of the flight trajectory reconstruction.

6.4.4. Pitching motion along a full pumping cycle

To study the pitching motion of the kite outside the reel-out phase, we zoom out and evaluate multiple pumping cycles, including the 65th cycle, which contains the previously investigated figure-of-eight manoeuvre. During the reel-in phase, the kite turns less, and the associated rolling motion is small. In contrast, the pitching of the kite induced by the tether sag is more pronounced as the tether tension reduces and the weight and drag of the tether are relatively large. Both the weight and drag of the tether result in the tether sagging downwards.

Figure 6.11 shows the kite pitch inferred from the wing measurements and the kite pitch resulting from the steady-rotation-state analysis with 30 tether elements. The results of ten consecutive pumping cycles are depicted, starting with the 65$^{\text{th}}$ pumping cycle. Each cycle starts with the transition into the reel-out phase, followed by approximately four figures of eight. Subsequently, the kite is pointed towards the zenith, depowered, and reeled back in (after the last shaded interval). The cycle ends after powering up again in preparation for a new cycle.

Each cycle shows an increase in pitch after the last turn in the reel-out phase as the kite transitions into the reel-in phase. The model overestimates the pitch at the start of the reel-in and underestimates it towards the end but gives a good overall agreement. There are many factors that may cause this discrepancy. One plausible explanation is that the reduced load during the reel-in phase leads to the deformation of the kite struts on which the sensors are mounted. The deformation is measured but not accounted for in the model and, thus, not incorporated in the computed results. Note that during the reel-in, the steering input is non-zero, as shown in Fig. 6.10a. This causes a pitch offset between the two sensors.

6.5. Turning mechanism of the kite

In this section, we discuss the implications of the observed swinging motion for the performance modelling of a kite system. Different mechanisms initiate and drive a turn of a flexible kite system with a suspended control unit.

The initiation mechanism for turning flexible kites with a suspended control unit relies on twisting the wing tips. A steering input causes the wing to twist, which increases the angle of attack at the wing tip at the inside of the turn and decreases it at the outside wing tip. This creates an aerodynamic side force component perpendicular to the kite symmetry plane and pointing towards the turn centre. The introduction of a side component effectively rolls the resultant aerodynamic force acting on the whole kite without rolling the kite itself. In contrast to flexible kites with a suspended control unit, multiline flexible kites that are actuated from the ground employ this mechanism to drive the whole turn; the side force is dominant in providing the centripetal force.

The driving mechanism for turning flexible kites with a suspended control unit is the rolling of the kite. As soon as the turn is initiated, the kite will roll into the turn to exert a centripetal force on the relatively heavy KCU, pulling it along. Together with the kite, the lift force generated by the top wing surface rolls into the turn and contributes to the centripetal force. The higher the mass of the KCU, the more roll is required to execute the same turn. Consequently, a smaller fraction of the lift is available to carry the weight of the airborne components and pull the tether. While the aerodynamic side force is still necessary to maintain turning, it is the roll of the kite that accommodates the largest contribution to the centripetal force and is thus considered to drive the turn.

To incorporate this turning mechanism, a single-point kite model would need the

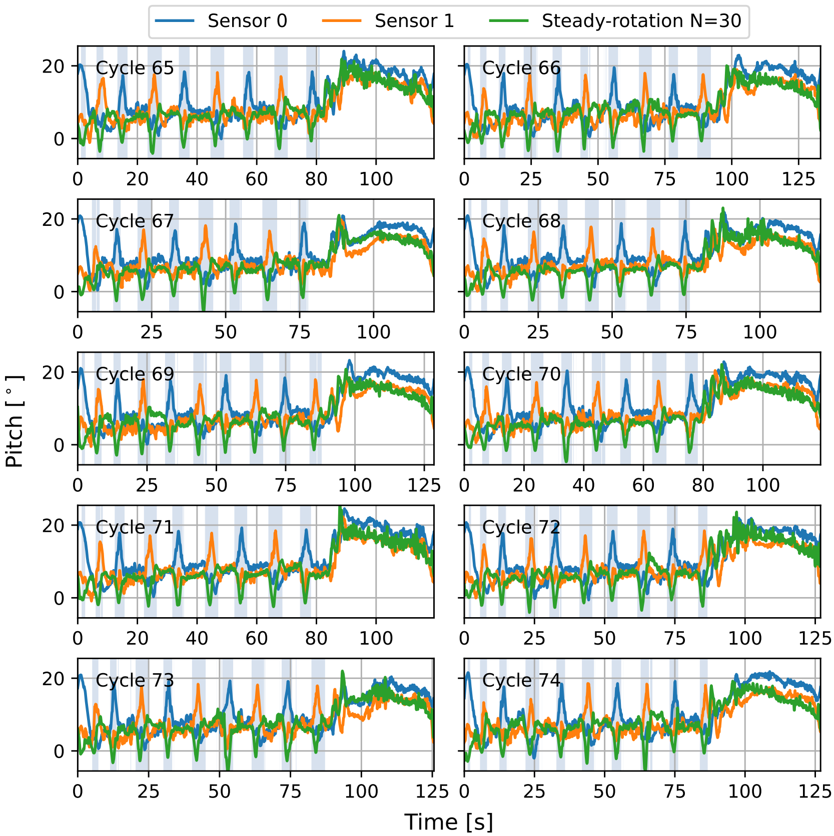

Figure 6.11: The pitch of the kite element with respect to the tangential plane along ten pumping cycles resulting from the steady-rotation-state analysis using 30 tether elements (T-I N=30), together with the kite pitch inferred from the wing attitude measured with two sensors. The shaded intervals indicate the turns during the reel-out phase. After the turns, the system transitions into the reel-in phase.

roll of the kite as an input, relying on the user to provide realistic roll angles. Another option is modelling the roll, e.g., using an empirical relationship between the roll and the steering input, as shown in Fig. 6.10e. However, with little extra computational cost, the roll can be resolved by modelling the kite with at least two point masses: one for the wing and one for the KCU. Thereby, it no longer needs to rely on system-specific empirical relationships to include the steering mechanism. Instead, the aerodynamic side force needed to initiate and maintain the turn can be calculated based on the deformation of the kite tips and associated aerodynamics.

Although the kite pitch does not change substantially during the reel-out phase, the tether-kite motion causes it to change substantially outside this phase. The sag-induced pitch concerns performance modelling as it affects the angle of attack experienced by the wing, which in turn affects the generated aerodynamic forces. Resolving the pitch also requires modelling the kite with at least two point masses and enables incorporating an aerodynamic model for the wing with a dependency on the angle of attack.

Complemented with an aerodynamic model of the kite, the dynamic model no longer relies on prescribing the wing acceleration. In contrast to the simulations conducted in the current analysis, simulations that solve the wing motion will be very sensitive to the wind input, which poses a large challenge to the validation of the model.

6.6. Conclusions

The inertia of the suspended control unit has a large effect on the roll of a flexible kite during turns in the reel-out phase. During the reel-in phase, the pitch of the kite changes due to the weight and drag of the control unit and increased tether sag. These effects are not resolved when the kite is modelled with a single point mass. With two point masses, one at the wing and one at the control unit, the steady-rotation-state model performs reasonably well in capturing the pitch and roll with little extra computational effort. A two-point model of the kite can thus be a powerful tool for the performance modelling of flexible kite systems.

The swinging motion of a kite with a suspended control unit is assessed with two approaches: approximated as a transition through steady-rotation states and solved dynamically. In contrast to the dynamic model, the steady-rotation-state model neglects transient effects. Both approaches employ a two-point kite model extending a discretised tether model using an additional rigid element for the kite. By prescribing the cross-wind flight path of the wing, no aerodynamic model of the kite is required.

An alternative expression for the angular velocity underlying the steady-rotation assumption is derived that accounts for the turning of the kite. This angular velocity expression accommodates lateral accelerations on the point masses and, thereby, allows studying the lateral swinging motion of the kite. The angular velocity for turns is approximated with flight data and shows good agreement with the kite kinematics. Unlike the original angular velocity expression, the proposed expression yields a good approximation of not only the wing velocity but also of the wing acceleration.

The tether-kite lines resulting from the steady-rotation states show discontinuities at the junction between the tether and the kite. These indicate that the control unit has a substantial effect on the attitude of the kite and stress the need for including a separate point mass for the control unit in performance models for flexible kite systems. The steady-rotation states perform reasonably well in estimating the roll of the kite, both with a single and 30 tether elements. The computed pitch and roll angles match the measured angles within three degrees during the straight sections of the figure-of-eight manoeuvre. During the turns, the peaks in the roll are overestimated, and the instantaneous differences in roll may exceed five degrees, whereas the pitch exhibits more systematic differences. These systematic differences could partially be explained by the fact that the model did not account for transient effects. However, drawing a definite conclusion is challenging, as the measurements include steering-induced pitch, making the wing measurements a poor reference.

Although the dynamic model considers transient effects, it is not more accurate in capturing the roll and pitch behaviour during turns than the steady-rotation states. This is expected to be primarily caused by inaccuracies in the wing acceleration and tether reel-out acceleration inputs. Due to anomalies in the flight trajectory measurements, a reconstruction was necessary to generate consistent inputs, enabling a running simulation. The reconstruction assumes that the tether slack length, defined as the difference between the tether length and radial position of the kite, remains constant. The intensity of the reconstruction adds further uncertainty to the results. Moreover, since the employed model is designed for computational efficiency, it does not capture non-trivial aspects such as tether elasticity.

Two separate mechanisms have been identified that initiate and drive a turn of a flexible kite system with a suspended control unit. A steering input causes an aerodynamic side force that initiates the turn. As soon as the turn is initiated, the kite starts to roll as it needs to pull the relatively heavy control unit into the turn. The rolled lift force provided by the top wing surface of the kite provides the largest contribution to the centripetal force and is said to drive the turn. Since a two-point kite model resolves the roll, the lift force may tilt along with the kite to drive turns. Hence, it avoids intricate centripetal force modelling, as seen in a single-point kite model. Furthermore, by resolving the pitch, it allows computing the angle of attack of the wing. The angle of attack is an important input to the aerodynamic model required when solving the wing motion instead of prescribing a flight path, as done in the current study.

The results of this study could be significantly improved with better quality flight data, more raw data, and information about how measurements are conditioned and calibrated. Currently, the sensor units are mounted to the flexible wing. As a result, wing deformation and actuation of the depower angle of the wing are also measured. This could be prevented by mounting the sensor units to the kite control unit. To find a better match between the measured and simulated tether forces, it would be interesting to incorporate variable tether slack and account for stretching in the dynamic simulation. A stepping stone could be to wrap the simulation in an optimisation problem to find the tether acceleration input that produces the measured tether force and cross-check the results with the tether lengths resulting from the steady-rotation states. More accurate tether length information in the experimental data would greatly help such analysis. Moreover, the flight trajectory reconstruction could be enhanced with this information, as well as with more advanced state estimation techniques. Finally, both the steady rotation state and dynamic model could still benefit from refining the wind modelling and fine-tuning the model parameters.The flight behaviour along the figure of eight described in the previous section is analysed with two different methods for solving the motion of the two-point kite model with a discretised tether model. First, this section discusses the tether-kite model configuration. Next, the two methods for solving the motion are discussed. The first approximates the tether-kite motion as a transition through steady-rotation states. The second solves the motion directly with dynamic equations of motion.

7

Modelling the flight operation of the pumping cycle

In the last two chapters, the analyses were limited to actual flight paths. This chapter abstracts the flight path as part of a basic flight operation model to compute the mean power of the pumping cycles flown in the test flight. The aim of this chapter is to validate the predicted performance using the instantaneous wind profiles estimated in Chapter 5.

7.1. Introduction

A wide variety of flight mechanical models for AWE systems exist in the literature. Most models have been developed with a specific system architecture and application in mind. In long-term performance assessments, these models are typically applied in optimisations to evaluate the full potential of a system. A classification of flight mechanical models and their application is depicted in Figure 2.8. The performance models reviewed hereafter are low-fidelity models for pumping systems.

The quasi-steady models (QSM) extend steady flight frameworks to encompass the full flight path. In contrast to dynamic models, they do not resolve the flight path with great detail and they rely to a large extent on a prescribed flight path. A large variety of flight path representations are used between different QSMs; ranging from abstract representations to finely discretised flight paths. The level of discretisation of the traction and retraction phases of four existing QSMs are plotted against each other in Figure 7.1. All four models assume a straight rigid tether.

Luchsinger [120] represents the pumping cycle with just two points: one for the traction and one for the retraction phase. The idealised flight states with the kite flying direct downwind yield simple expressions for the mean cycle power. Fechner and Schmehl [62] and Van der Vlugt et al. [46] account for elevated flight and use simple integration schemes to solve an idealised trajectory of the pumping cycle without solving the crosswind motion explicitly. Van der Vlugt et al. [46] account for the crosswind

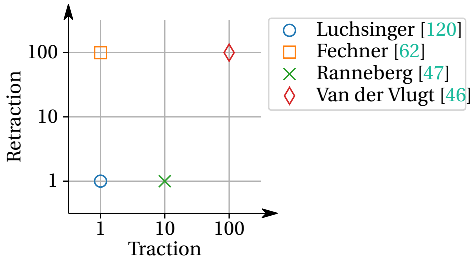

Figure 7.1: QSMs mapped according to the order of magnitude of the number of points used to represent the traction and retraction phases of the pumping cycle.

motion with a surrogate flight state at representative elevation and azimuth angles that coincide with the middle of a figure-of-eight pattern.

The model proposed by Ranneberg et al. [47] is developed for flight operation optimisations. The flight trajectory is to a large extent prescribed and thus not solved. For the traction phase, flight states are evaluated at collocated points lying on the surface described by a prescribed figure of eight swept along an elevated axis. The retraction phase is represented with a single flight state yielding a fixed retraction slope typical to rigid kite systems. Optimisation is used to find the tether force, reeling speed, and roll angle that yield a force equilibrium at each point and at the same time maximise the mean cycle power.

In contrast to QSMs, dynamic models resolve the flight path. Optimal control can be used offline in flight path optimisation. For this application, the dynamic model does not necessarily produce realistic control actions such as by a flight control computer. Alternatively, the controlled variables may consist of high-level kinematic quantities such as the roll of a kite.

Gros and Diehl [51] present a dynamic point mass model and rigid body model in Cartesian coordinates of a single kite system that can efficiently be applied in optimal control problems. Malz et al. [24] conducted a validation study that uses the rigid body model with empirically obtained stability and control derivatives of the aerodynamic loads of the Ampyx Power AP2 prototype in an optimal control problem that fits the simulation to flight data. Zanon et al. [53] formulated a general description for deriving models of multi-kite systems. The procedure also allows discretising the tether as lumped masses connected with rigid links.

Validation of the mentioned high-level flight operational models is challenging since the simplifications needed to reduce computational cost increases the gap between the actual measured quantities in a flight test and the model results. Moreover, dealing with the uncertainty about the wind conditions during the flight test poses a large challenge. Malz et al. [24] solve this challenge by allowing changes to the wind velocity with respect to the measurements.

Flexible kites are less well captured with a rigid body model than rigid kites due to the more complex fluid-structure interaction. Furthermore, flexible kites are less well characterised due to the complexity of carrying out system identification experiments with flexible kites. Consequently, stability and control derivatives are readily available

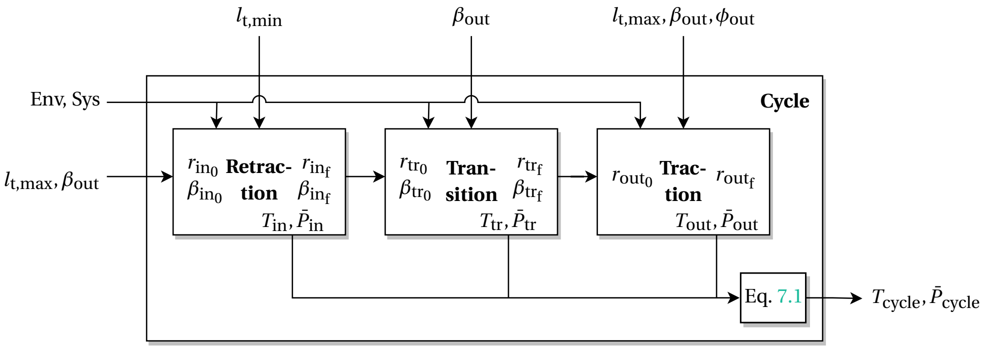

Figure 7.2: Flow diagram in the QSM of Van der Vlugt et al. [46] showing the interdependency of the phase modules. The quantities specifying the environment and system are shared between the phases.

for rigid kites, but not for flexible kites. This makes it even more challenging to acquire faithful dynamic models for flexible kites and validate them.

QSMs are more suitable for flexible-kite systems than for rigid-kite systems, since their high surface-to-mass ratio yields faster dynamics [46]. Thereby, the state of the kite at a given instance does not depend much on its preceding trajectory. Since the model of van der Vlugt et al. include the most detailed flight path representations relative to the other QSMs, it can be more directly compared with the high-resolution flight data. For this reason, the model of van der Vlugt et al. is used as starting point to compare modelling results against flight data.

7.2. Quasi-steady model framework

The aim of the QSM developed by van der Vlugt et al. is to assess the performance of the pumping flight operation of a flexible kite system. The QSM divides the pumping cycle into three phases: the traction, retraction, and transition phases. Each phase has a separate module in the model, as depicted in Figure 7.2. Within each phase, the trajectory of the kite is approximated as a transition through steady flight states, which are calculated using the steady flight state approximation evaluated in Section 5.6. The mean cycle power is calculated using the average power and duration output of the phase modules:

$$\bar{P}_{cycle} = \frac{\bar{P}_{in}\, T_{in} + \bar{P}_{tr}\, T_{tr} + \bar{P}_{out}\, T_{out}}{T_{in} + T_{tr} + T_{out}} \quad . \tag{7.1}$$

in which T is a duration and $\bar{P}$ an average power. The subscripts in, tr, and out refer to the retraction, transition, and traction phases, respectively.

The Euler method is used to solve the displacement along one or two spherical coordinates depending on the phase, see the displacement equations in Table 7.1. The velocity is not solved through integration but is calculated with the steady-state approximation. Note that the azimuthal displacement is not solved in any of the phases. This is a reasonable approximation when the cross-wind figure-of-eight manoeuvre of the kite is not very wide such as for the retraction trajectory flown by Kitepower in their flight

Phase	Retraction	Transition	Traction
Fixed states	$\phi = 0$	$\phi = 0$	$\beta_{\mathrm{out}}, \phi_{\mathrm{out}}$
Free states	r, β	r, β	r
Initial values	$l_{\mathrm{t,max}}, \beta_{\mathrm{out}}$	$l_{\mathrm{t,min}}, \beta_{\mathrm{in_f}}$	$r_{\mathrm{tr_f}}$
Displacement equations	$\Delta r = v_{\mathrm{t}} \Delta t$ $\Delta \beta = -\frac{v_\tau \cos\chi}{r} \Delta t$	$\Delta r = v_{\mathrm{t}} \Delta t$ $\Delta \beta = -\frac{v_\tau \cos\chi}{r} \Delta t$	$\Delta r = v_{\mathrm{t}} \Delta t$
End criterion	$r = l_{\mathrm{t,min}}$	$\beta = \beta_{\mathrm{out}}$	$r = l_{\mathrm{t,max}}$
Controlled parameter	$\begin{cases} F_{\mathrm{t,in}} \\ -v_{\mathrm{t,min}}, \text{ if } v_{\mathrm{t}} > -v_{\mathrm{t,min}} \\ -v_{\mathrm{t,max}}, \text{ if } v_{\mathrm{t}} < -v_{\mathrm{t,max}} \end{cases}$	$\begin{cases} v_{\mathrm{t}} = 0 \\ F_{\mathrm{t,min}}, \text{ if } F_{\mathrm{t}} < F_{\mathrm{t,min}} \\ F_{\mathrm{t,max}}, \text{ if } F_{\mathrm{t}} > F_{\mathrm{t,max}} \end{cases}$	$\begin{cases} F_{\mathrm{t,out}} \\ v_{\mathrm{t,min}}, \text{ if } v_{\mathrm{t}} < v_{\mathrm{t,min}} \\ v_{\mathrm{t,max}}, \text{ if } v_{\mathrm{t}} > v_{\mathrm{t,max}} \end{cases}$

Table 7.1: The differences in how the displacement of the kite is solved within each of the phase modules of the QSM of Van der Vlugt et al. [46].

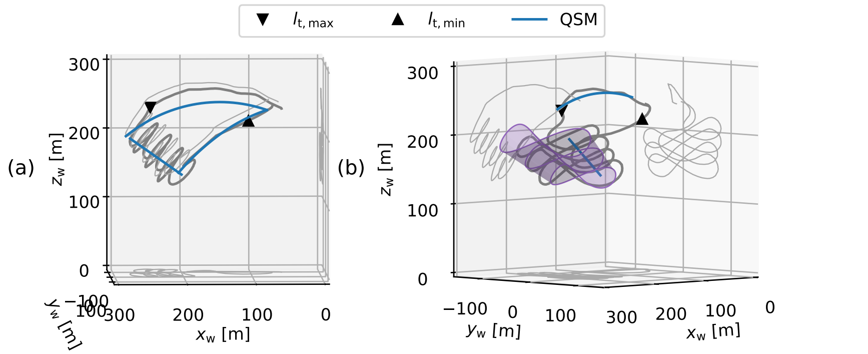

Figure 7.3: Trajectory representations in the QSM framework. (**a**) Closed cycle trajectory of the model of Van der Vlugt et al. [46] . (**b**) Separate traction and retraction phase trajectories used for cross-checking with flight data. The purple surface illustrates the space on which the kite trajectory lies in the surface representation of the traction phase.

test, see Figure 2.11. However, the approximation may be less suitable when alternative retraction approaches are used, e.g., flying the kite to the side of the wind window as done by SkySails [10].

Figure 7.3a shows a trajectory resulting from the QSM. The trajectories of the retraction and transition phases are represented with curved paths, whereas the traction phase trajectory is represented with a straight path. The tether length at the end of a certain phase is the starting tether length of the succeeding phase. During the transition phase, the tether force limits are commonly exceeded. Consequently, the kite may be reeled out and will start the retraction phase with a larger tether length than the transition phase.

The model does not resolve the cross-wind flight; not even for the traction phase to which cross-wind flight is inherent. Instead, the cross-wind flight is approximated with steady flight states at constant elevation β_{out}, azimuth ϕ_{out}, and course angle χ_{out} representative for the figure-of-eight manoeuvre. Only the radial displacement of the kite is solved yielding the reel-out path represented as a straight line. The line representation of the traction phase is only applicable when the controller is designed to fly a constant tangential pattern, i.e., the cross-wind flight pattern projected onto the azimuth-elevation plane stays more or less the same during the traction phase.

Alternatively, the cross-wind flight can be accounted for by prescribing a figure-of-eight pattern as done in the surface representation of the traction phase. The surface describes the space on which the kite trajectory is considered to lie, depicted in Figure 7.3b. By allocating points on this surface in a smart way, the steady states at these positions together give a good representation of the traction phase and thus of the mean traction power. Note that this spatial discretisation does not necessarily require solving a realistic kite trajectory conform equations of motions. The steady-state results need to be weighted proportional to the kite speed to determine the phase averages. This approach has previously been proposed by Ranneberg et al. [47].

7.3. Figure-of-eight power estimation

The kite is flying figures of eight for more than half of the time in a pumping cycle. Therefore, it is crucial to the mean cycle power estimation that the power output of a figure of eight can be predicted accurately. Assessing the power estimation for figure-of-eight manoeuvres is the first step towards assessing the mean cycle power estimation.

The aim of this section is to evaluate if the variation in power output of the figures of eight recorded in the flight data can be explained using simplified traction phase modelling approaches. The radial position of the kite is fixed in the calculations, which is analogous to only considering a cross-section of the line- and surface representation. First, the common trends and differences between the figures of eight in the flight data are identified. These are then used to infer the required input for both modelling approaches. Finally, the tether force and generated power calculated with the two modelling approaches are cross-checked and compared with flight data.

7.3.1. Trends in the flight data

Trends of the cross-wind flight are identified from the flight data to make informed decisions on what simulation approaches to take such that the trends are adequately captured. To do so, the individual figures of eight are extracted from the test flight data introduced in Section 5.2. These comprehensible subsets enable identifying the common trends and differences between the figures of eight.

In Figure 7.4, the evolutions of important quantities of all figures of eight are plotted next to each other. The good alignment of the azimuth angle at the start in Figure 7.4a stems from the figure-of-eight extraction routine. This routine imposes that the figure of eight starts at zero azimuth and ends when the kite passes the zero azimuth plane for the second time. Consequently, all the figures of eight start at the same point, after which the spread of the lines increases with time due to slight differences in the flight trajectory. The mean duration of a figure of eight is roughly 21 s.

Although for some quantities, larger differences between individual figures of eight are observed, all quantities exhibit a distinct pattern. To identify the common trend, each evolution is expressed with respect to time normalised by their respective duration. Next, all evolutions are resampled with the same resolution to calculate the median evolution. Subsequently, the median evolution is de-normalised with the average figure-

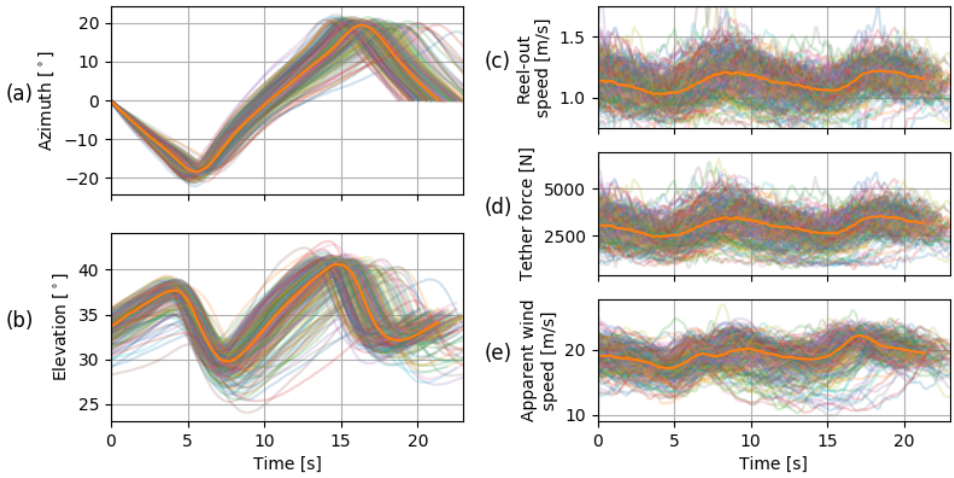

Figure 7.4: The evolutions of quantities measured during each figure of eight flown in the flight test together with the inferred common trends.

of-eight duration. The resulting de-normalised median evolutions are plotted with the orange lines in Figure 7.4.

Plotting the elevation trend against the azimuth trend gives the average figure of eight shown in Figure 7.5. The figures of eight are not symmetric and the right lobe of the figure of eight typically reaches a higher elevation angle.

The linear tether speed-force relationship imposed by the controller results in the medians of both quantities exhibiting the same trends, as shown in Figure 7.4c and d. The medians of the tether speed and force exhibit two cycles within one figure of eight with a 7% lower depth of the first valley. In Figure 7.6a, the colour scale illustrates how the tether force trend evolves along the figure of eight. It clearly shows that the valleys coincide with the start of the turns (top of the figure of eight) and the peaks coincide with the start of the straight sections (bottom of the figure of eight).

The apparent wind speed evolutions shown in Figure 7.4e are fluctuating more heavily than the other quantities. The median evolution is nevertheless still able to adequately capture a common trend. Figure 7.6b shows that the measured apparent wind speed is highest during the right turns. In the lower right corner the apparent wind speed peaks at roughly 22 m/s while the apparent wind speed is only roughly 19 m/s in the lower left corner.

Unexpectedly, Figure 7.6 shows a misalignment between the apparent wind speed and tether force. As the force is measured at the ground station, it is not directly affected by asymmetries in the kite. On the contrary, the airspeed is measured off-centre by the flow measurement setup mounted to the left of the bridle. Thereby, the Pitot tube travels a shorter distance during the left turns than during the right turns. The observed difference in airspeed between the turns approximately agrees with the turning-imposed speed difference of the Pitot tube. Therefore, the true apparent wind speed at the centre of the kite is expected to evolve symmetrically.

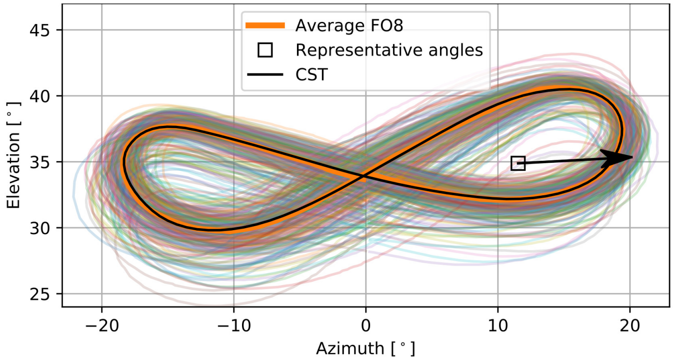

Figure 7.5: The cross-wind flight trajectories of each figure of eight flown in the flight test with on top the fitted curve found with the CST method. Also depicted are the representative figure-of-eight angles of the line representation of the traction phase inferred from the flight trajectories. The direction of the arrow illustrates the representative coarse angle.

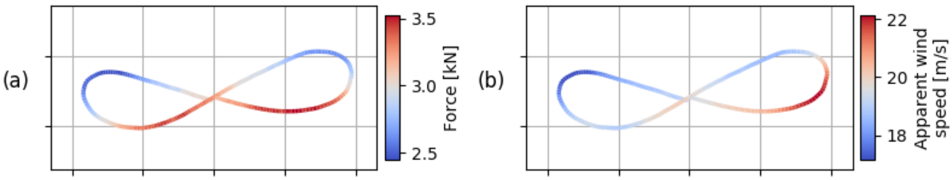

Figure 7.6: The evolutions of the common trends found for the measured tether force and the apparent wind speed along the figure-of-eight pattern.

The figures of eight flown in the flight test with a fixed control approach exhibit substantial variations in performance. These differences can mostly be explained by the wind variability. Figure 7.7a and b show the relationships between the average tether force and the apparent and reconstructed wind speeds. The reconstructed wind speed follows from the reconstruction in Section 5.3. The apparent wind speed has a strong relationship with the tether force. The relationship between the tether force and reconstructed wind speeds is less direct, which manifests by a relatively large vertical spread of the data points around the trend line. Moreover, the imprecision of the wind reconstruction will contribute to this spread and obscure the strength of the relationship.

Exponential curves $F_\mathrm{t} = c\,v^p$ are fitted to the data in Figure 7.7. In this expression, c is a constant, p is the effective exponent, and, depending on the panel, v is the apparent or the reconstructed wind speed. The resulting curves are plotted with the dashed black line.

The fitted exponential curve of Figure 7.7a has an effective exponent $p = 2.73$. The value of the effective exponent is high compared to the quadratic relationship between the tether force and apparent wind speed that follows from the massless kite theory. This theory assumes that the tether force counteracts the aerodynamic force, imposing the

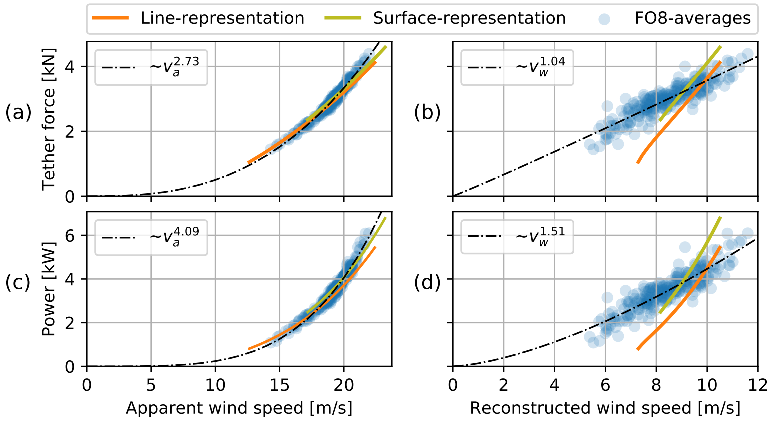

Figure 7.7: The variation of the averaged tether force and power with the averaged apparent and reconstructed wind speed between the figures of eight flown in the flight test.

quadratic relationship. In reality, not all the generated aerodynamic force can be utilised to generate power, leading to an increase in the effective exponent. In particular, at low airspeeds, a relatively large portion of the aerodynamic force is needed to counterbalance the weight of the kite. Note that this curve is not expected to perfectly capture the relationship but acts merely as a reference.

Figure 7.7c and d show that the relationship of the power with the apparent wind speeds has a higher effective exponent than the relationship with the wind speed, which is almost linear. This does not align with theoretical studies such as that of Luchsinger [120] which find a cubic relationship with wind speed. These studies typically assume a constant reeling factor while in the flight test the tether speed-force ratio is kept constant. The resulting variation of the reeling factor will affect the effective exponent and may partly explain the difference in the effective exponent.

7

7.3.2. Line representation

Utilising the line representation to abstract crosswind flight could provide a useful tool for performance analysis depending on the magnitude of the associated error. By cross-checking the modelling approach with other approaches, the error associated with disregarding the specifics of the figure-of-eight manoeuvre is quantified.

Priorly, the line representation is compared with the figure-of-eight performance recorded in flight data. In the calculation, the tether length variation within a figure of eight is neglected by estimating the power output of a figure of eight with the steady state calculation at a single point lying on the line. This calculation is repeated for the range of wind speeds encountered during the reel-out phases in the flight test. The tether length is fixed to 285 m in the calculation. This tether length is representative of the mean tether length of the second figure of eight of a pumping cycle. The flight data also includes preceding and succeeding figures of eight and, therefore, some variation in power output recorded in the flight data can be attributed to the variation in tether length. This contribution is however not considered because it is small with respect to the power output variation due to the wind speed.

Van der Vlugt et al. [46] hypothesise that the line representation gives a good approximation of the mean traction power when the representative angles are inferred from the time averages of the cosines of the angles recorded in flight data or from a more faithful model of the cross-wind flight:

$$\cos\phi_{\text{out}} = \frac{\int_0^{T_{\text{out}}} \cos\phi \, dt}{T_{\text{out}}} \quad , \tag{7.2}$$

$$\cos\beta_{\text{out}} = \frac{\int_0^{T_{\text{out}}} \cos\beta \, dt}{T_{\text{out}}} \quad . \tag{7.3}$$

Note that the time-averaged angles depend on the tangential speed of the kite, which is varying along the figure of eight. An exact match in power requires that the power is a linear function of the cosines of these angles. However, this is not even the case in the massless kite theory (see Equation 5.2), let alone for a more realistic power calculation. The same method can be used to determine the representative coarse angle. The square markers in Figure 7.5 show the resulting representative figure-of-eight angles of the line representation. The corresponding point in the tangential plane lies within the figure-of-eight lobe. Although, for the more asymmetric figures of eight the point may also lie outside the lobe.

The lift and drag coefficients of the wing are considered constant and the mean values that were found in the identification of Section 5.5 are used: $C_L = 0.7$ and $C_D = 0.16$. Moreover, the tether speed-force ratio control is imposed by wrapping the steady state calculation in an optimiser, which searches for a solution that lies on the linear curve found in Section 5.5.2:

$$v_{\text{t}} = 0.18\text{e}{-3}\, F_{\text{t}} + 0.58 \quad , \tag{7.4}$$

which yields the tether reel-out speed v_{t} in m/s given the tether force F_{t} in N.

The orange lines in Figure 7.7 show the line-representation results. Below 7.5 m/s wind speed, no steady state solution is found with the representative figure-of-eight angle, while the recorded reconstructed wind speed goes down to almost 5 m/s. At 7.5 m/s wind speed, the steady state solution yields an apparent wind speed of roughly 13 m/s, which is lower than the lowest recorded figure-of-eight-average apparent wind speed. At 10.5 m/s wind speed, the apparent wind speed of the steady state solution exceeds 22 m/s, which is roughly the highest recorded figure-of-eight-average apparent wind speed.

The resulting tether force-apparent wind speed curve shown in Figure 7.7a is nearly quadratic and only slightly increases the effective exponent with respect to the massless kite theory by accounting for weight. The effective exponent found with the line representation is substantially lower than the effective exponent $p = 2.73$ found in the flight data, suggesting that not all factors increasing the effective exponent are considered. Chapter 6 showed that a part of the generated aerodynamic force is needed to turn the kite. Thereby, turning will also result in an increase in the effective exponent. However, the effect of turning cannot be studied with the original steady-state calculation as it does not include turning dynamics.

The discrepancy between the line representation and flight data appears a lot larger based on the tether force-wind speed relationships in Figure 7.7b. Although both show near-linear trends, their slopes are not similar. The line representation appears to underestimate the tether force below 10 m/s wind speed.

Assuming that the wind speed reconstruction is correct, the power production efficiency decreases with increasing wind speed. One possible explanation is that the kite aerodynamics are more efficient at low speeds, while in the calculation the aerodynamic coefficients are considered constant. An alternative explanation is that a bias in the reconstructed wind speeds leads to a wider range of wind speeds than the actual range. Thereby, the slope of the relationship between force and wind speed appears to be smaller than the actual slope. The variations in tether length and representative figure-of-eight angles are only expected to cause small changes to the power output recorded in the flight data.

7.3.3. Surface representation

In contrast to the line representation, the surface representation does include specifics of the figure-of-eight flight path and, thereby, enables a more direct comparison with flight data. Note that it does not resolve the cross-wind motion but requires prescribing the figure-of-eight pattern.

To estimate the performance of a figure-of-eight manoeuvre, the radial position is virtually fixed while solving the cross-wind displacement. The figure-of-eight pattern is evenly discredited, and the steady states at the corresponding positions are solved. The resulting evolution of the tangential kite speed is accounted for when calculating the quantities averaged over a figure of eight.

First, a new parameterisation method is presented that can produce realistic figures of eight. The description representative of the figures of eight in the flight data is then used to evaluate how well the variation in power output is captured in flight data by the surface representation.

figure-of-eight parameterisation

When assuming tangential pattern control, the pattern can conveniently be described in the azimuth-elevation plane. A commonly used parameterisation is a special case of a Lissajous curve:

$$\phi_{\text{Lissa}}(s) = \frac{w}{2}\sin(2\pi s) + \phi_c \quad , \tag{7.5}$$

$$\beta_{\text{Lissa}}(s) = \frac{h}{2}\sin(4\pi s) + \beta_c \quad , \tag{7.6}$$

in which s is the independent variable connecting the two equations (no to be confused with the distance), ϕ_c and β_O set the center, w is the width, and h the height of the figure of eight. The curve is described in the tangential plane. One figure of eight is completed in $s \in [0, 1]$. The Lissajous parameterisation is too limited to describe a realistic cross-wind pattern of a kite. The Lissajous figure of eight is both horizontally and vertically symmetric, whereas a real-world figure of eight is typically asymmetric. Therefore, a more flexible parameterisation is needed to facilitate using more realistic figure-of-eight patterns in the QSM.

The class/shape function transformation (CST) method is adapted to manipulate the Lissajous curve and thereby describe realistic figures of eight. The CST method is originally developed for efficiently describing airfoil geometries and that of other aircraft components [121]. The class function describes the general desired features of the geometry. The shape function is a smooth function close to one and facilitates manipulating the geometry locally. The product of the two functions yields the final geometry. The method only requires a limited number of parameters to generate a wide variety of figures of eight.

A Bernstein polynomial, which is a linear combination of Bernstein basis polynomials, is originally used as the shape function. An important feature of this polynomial is that is smooth. However, it does not preserve periodicity, which is a requirement for the figure-of-eight pattern. The 2$^{\text{nd}}$basis polynomial of the 4$^{\text{th}}$order Bernstein polynomial:

$$p(x) = 6x^2(1-x)^2 \quad , \tag{7.7}$$

has a value and slope of zero at $x = 0$ and $x = 1$. These properties are utilised to mimic a periodic smooth function by pairing two of these basis polynomials:

$$x_{\text{left}} = \frac{1}{w_p}(s - s_0 + 1) \quad , \tag{7.8}$$

$$x_{\text{right}} = \frac{1}{w_p}(s - s_0) \quad , \tag{7.9}$$

$$p_{\text{left}}(s) = \begin{cases} p(x_{\text{left}}), & \text{if } 0 \leq x_{\text{left}} \leq 1 \\ 0, & \text{otherwise} \end{cases} \quad , \tag{7.10}$$

$$p_{\text{right}}(s) = \begin{cases} p(x_{\text{right}}), & \text{if } 0 \leq x_{\text{right}} \leq 1 \\ 0, & \text{otherwise} \end{cases} \quad , \tag{7.11}$$

$$p_{\text{pair}}(s) = p_{\text{left}}(s) + p_{\text{right}}(s) \quad , \tag{7.12}$$

in which p_{left} is shifted by 1 to the left with respect to p_{right}, $w_p = 0.4$ sets the width of the these polynomials, and p_{right} starts being non-zero from s_0.

The sum of $\mathrm{N_p} = 10$ evenly shifted and weighted polynomials pairs added up by 1 is used as the shape function:

$$S(s, \mathbf{b}) = \sum_{i=1}^{\mathrm{N_p}} \left(b_i\, p_{\text{pair},i}(s)\right) + 1 \quad , \tag{7.13}$$

in which b_i is the weighing factor of the i$^{\text{th}}$ polynomial pair. The set of basis polynomials pairs is depicted in Figure 7.8a. Only the first three polynomial pairs are non-zero at both sides of the interval. For the others, one of the pairs falls outside the relevant interval. By adjusting the weighting factors assigned to each polynomial pair, a large variety of smooth, periodic functions can be generated.

The Lissajous curve is a logical candidate as the class function when using the CST method to generate realistic figures of eight. However, the shape function is not effective if the class function is close to zero. For the special case where $\phi_c = \beta_c = 0$ this results in

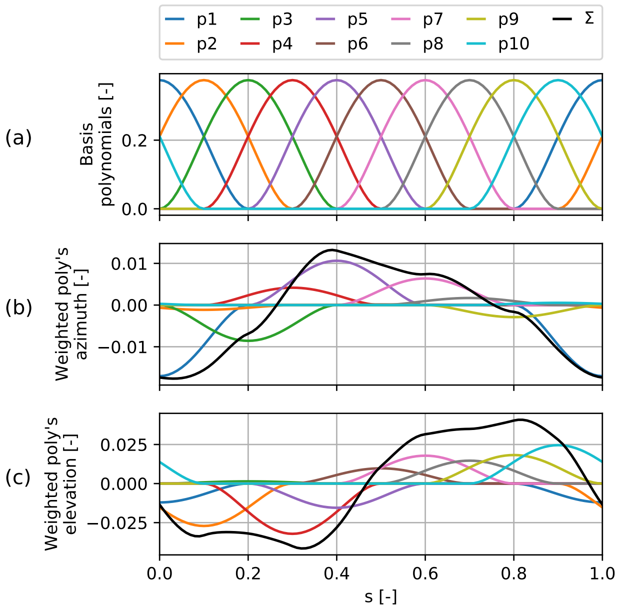

Figure 7.8: (**a**) Set of polynomial pairs that enable smooth and local adjustments along the curve obtained with the CST method. (**b**, **c**) The basis polynomial pairs weighted with the values found for the fit shown in Figure 7.5. Adding 1 to the sum of the weighted polynomial pairs yields the shape functions for the azimuth and elevation position.

the final geometry intersecting the origin twice as dictated by the class function. To assure the effectiveness of the shape function for $s \in [0,1]$, the Lissajous curve is displaced first before multiplied with the shape function and displaced back afterwards, i.e., 10 is added up to both Lissajous curve expressions and subtracted again after being multiplied with the shape function:

$$\phi_{\mathrm{CST}}(s,\mathbf{b}_\phi) = \left(\phi_{\mathrm{Lissa}}(s)+10\right) S_\phi(s,\mathbf{b}_\phi) - 10 \quad , \tag{7.14}$$

$$\beta_{\mathrm{CST}}(s,\mathbf{b}_\beta) = \left(\beta_{\mathrm{Lissa}}(s)+10\right) S_\beta(s,\mathbf{b}_\beta) - 10 \quad . \tag{7.15}$$

The CST method is used to produce a figure of eight representative for the whole test flight. The weighing factors $\mathbf{b}_\phi$ and $\mathbf{b}_\beta$ that yield a good fit of the adapted Lissajous curve to the representative figure of eight are obtained using optimisation:

$$\min_{\mathbf{b}_\phi,\mathbf{b}_\beta,\mathbf{s}} \quad \sum_{i=1}^{\mathrm{N_s}} \left(\phi_{\mathrm{CST}}(s_i,\mathbf{b}_\phi) - \hat{\phi}_i\right)^2 + \left(\beta_{\mathrm{CST}}(s_i,\mathbf{b}_\beta) - \hat{\beta}_i\right)^2 \quad , \tag{7.16}$$

in which $\mathrm{N_s}$ is the number of points used for discretising the representative figure of eight and $\hat{\phi}_i$ and $\hat{\beta}_i$ are its coordinates at the $\mathrm{i^{th}}$ point. Note the optimisation problem positions the control points of Equations 7.14 and 7.15 at the same value of the independent variable s. The optimal shape function converges to approximately the same solution when further increasing $\mathrm{N_s}$ after a sufficiently fine discretisation is obtained. The width,

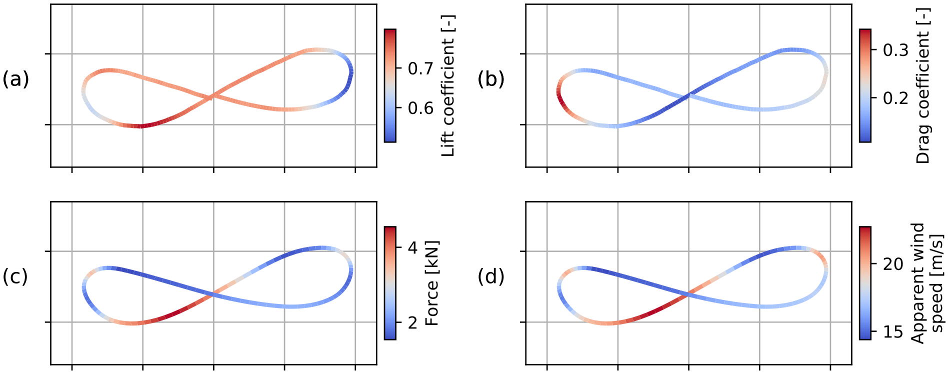

Figure 7.9: (**a**, **b**) The evolutions of the common trends found for the identified lift and drag coefficients along the figure-of-eight pattern. (**c**, **d**) The evolutions of the tether force and apparent wind speeds along the figure-of-eight pattern computed with the QSM for a wind speed of 8 m/s.

height, and centre coordinates of the Lissajous curve act as global decision variables, whereas the weighing factors only affect the shape locally.

The fitted adapted Lissajous curve is shown in Figure 7.5. It lies directly on top of the figure of eight representative for the full test flight. For the current fitting problem, the Lissajous curve parameters are fixed to $\phi_c = 0°$, $\beta_c = 35°$, $w = 37°$, and $h = 8°$. Given the quality of the fit, it suffices to fix these parameters.

Simulation results

Similar as for the line-representation assessment, the simulations are conducted for the range of wind speeds found in the flight data. Moreover, the same tether speed-force ratio is imposed and the same representative tether length is used.

Instead of employing constant lift and drag coefficients in the simulation, the lift and drag coefficient are adjusted along the figure of eight to enable approximating the measured tether force in Figure 7.6a more closely. The measured tether force shows that the highest tether force is only attained after steering is released, whereas, with constant aerodynamic coefficients, the highest tether force is expected when the kite is accelerated by the weight of the kite when flying down. This suggests that the drag of the kite is increased in the turns due to the steering-induced deformation of the kite.

For simplicity, the lift and drag coefficients are implemented as function of the position of the kite along the figure of eight, instead of directly as function of the angle of attack and steering input. Similar to the earlier identification of trends along the figure of eight, common trends can be identified for the evolutions of the lift and drag coefficients using the coefficients identified in Section 5.5. The resulting evolutions are shown in Figure 7.9a and b. Clear drops in the lift coefficient and peaks in the drag coefficient can be observed during the turns. The identification of the aerodynamic coefficients is based on the biased measured apparent wind velocity. Consequently, the evolutions of the identified aerodynamic coefficients contain errors due to this turning bias which in turn introduce errors in the simulation results.

Complementary to the original QSM, the current modelling approach solves the crosswind motion with the displacement equation: $\Delta\tau = \frac{v_\tau}{r}\Delta t$ in which τ is the distance in the azimuth-elevation plane. In the current analysis, the radial displacement is not solved. The kite is moved along the prescribed figure of eight at a fixed radial position with equal step lengths. Fixing the radial position enables a fair comparison with the line representation.

Figure 7.9c and d show the results of the simulation for 8 m/s wind speed. The simulated tether force shows a mismatch with the trend identified from the measurements, see Figure 7.6a. The simulated tether force shows a single, relatively large peak when exiting the left turn, as opposed to peaks at the exits of both turns found in the measurements. Moreover, the simulated tether force evolves asymmetrically along the figure of eight, whereas the evolution of the measured tether force is more symmetric. As expected, the simulated tether force and apparent wind speed evolutions along the figure of eight nicely align, in contrast to the trends of the respective measurements.

Figure 7.7d shows that the surface representation finds a higher power production efficiency than the line representation. The power curve found with the surface representation intersects the flight data trend at 8.8 m/s wind speed, which is roughly the average reconstructed wind speed. This suggests that the surface representation yields a better power estimate.

With the current approach, the surface representation cannot explain the variation in power output substantially better than the line representation. The surface representation considers the effect of turning on the aerodynamic coefficients. Still, it does not consider the turning dynamics that are expected to further increase the effective exponent of the tether force-apparent wind speed relationship in Figure 7.7a.

Similar trends of the tether force-wind speed relationship are obtained with both the line- and surface representations as shown in Figure 7.7b. While the computed trends as a function of the apparent wind speed exhibit reasonable agreement with flight data, the agreement significantly diminishes when the trends and flight data are expressed against the reconstructed wind speed. This suggests that there may be inaccuracies either in the conversion of the modelled apparent wind speed to the free wind speed, in the conversion of the measured apparent wind speed to the reconstructed wind speed, or in both processes. It is most likely that a deficiency in the wind speed reconstruction explains the discrepancy in agreements. The turning bias in the apparent wind speed measurement due to off-centre flow measurements could contribute to the faulty reconstructed wind speeds. However, this effect is expected to be small as the turns are disregarded in determining the figure-of-eight-average reconstructed wind speeds.

7.4. Mean cycle power estimation

The previous section assessed how well the power production of a figure-of-eight manoeuvre is estimated using a quasi-steady approach. This section zooms out and assesses how this translates to the traction phase and ultimately to the full pumping cycle. The aim of this section is to identify the modelling approach that yields the best estimate of the mean power of the pumping cycles recorded in the flight data.

7.4.1. Traction phase

The traction phase in real flight is not as clearly delimited as in the QSM. At the start of the pumping cycle, the kite is pointed away from the zenith and employs two downloops to set up the traction phase. At this stage, the kite is already reeled out and producing power. However, only the part of the pumping cycle where the kite is flying figures of eight around the reel-out path with constant elevation angle (pink trajectory in Figure 2.11) is considered in the QSM.

The figure of eight is the primary element of the traction phase, but power is also generated before and after the figure-of-eight flight which consists of three to four figures. Also, these transitions should be considered in the final mean cycle power calculation. For now, the initial downloops and final uploops are disregarded and only the flight along the elevated reel-out path is analysed. The averaging period is increased from roughly 21 s for figures of eight to more than a minute for the traction phase.

Traction phase simulations are carried out with both the line- and surface representations and are repeated for the range of wind speeds encountered during the traction phase in the flight test. For simplicity, a uniform wind profile is assumed in the simulations. Each simulation covers a range of tether lengths, whereas previously only a single tether length was used. As the variation caused by the reeling length differences are minor, the same starting and ending tether lengths are used for every simulation, irrespective of the wind speed. The performance is evaluated on a linear grid at 10 radial positions between 252 m and 338 m.

The simulations with the surface representation solve the cross-wind and reeling motion separately. The cross-wind motion is solved first and yields average performance quantities over a figure of eight at a virtually fixed radial position. Next, the reel-out speed evolution along the elevated path is calculated by repeating the figure-of-eight evaluation for a number of tether lengths. The reel-out speed evolution is used to average the performance quantities over the full traction phase. A lower tether force limit is introduced to the simulations to yield better convergence at large tether lengths for which the drag is relatively large.

Figure 7.10 illustrates how the simulation compares with flight data for the reference pumping cycle. For this comparison, the tether length range is directly inferred from the flight data of the cycle. The line representation shows a lower tether force and reel-out speed compared to the surface representation and thus a longer duration with the actual duration lying in between. The respective average powers are 3.5 kW and 4.6 kW, while the average power recorded in the flight data is 4.1 kW.

The same possible explanations of the disagreement between simulated and observed figure-of-eight power apply to the full traction phase. Consequently, the simulated power-wind speed relationship exhibits a larger slope than the trend observed in the flight data, see Figure 7.11a. A linear curve is used to depict the trend of the power output against the average reconstructed wind speed during the traction phases in the flight data.

Besides the power, the duration of the traction phase is needed to calculate the mean cycle power calculation. Figure 7.11b shows that the simulation predicts the duration of the traction phases poorly for low wind speeds. However, a reasonable match with the trend in the flight data is observed for higher wind speeds.

7

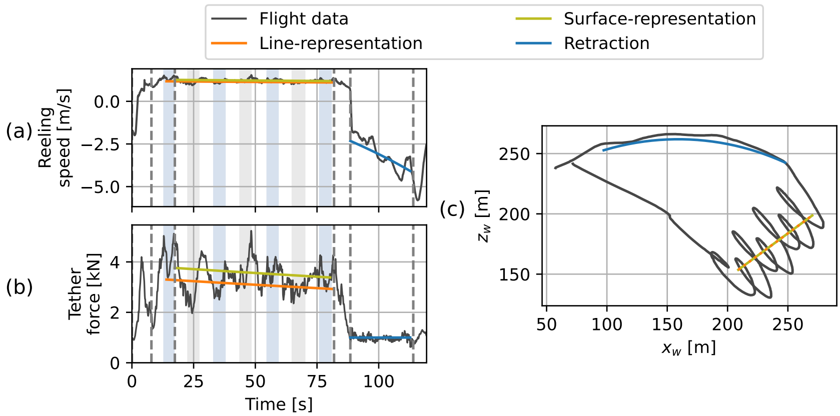

Figure 7.10: (**a**, **b**) Time series of the controlled properties. (**c**) Side-view of the trajectory for the reference pumping cycle. The equivalent QSM simulation results are plotted alongside the flight data.

7.4.2. Retraction phase

The original QSM starts the retraction phase straight after finishing the traction phase as shown in Figure 7.3a. In reality, the kite first needs to decelerate from the fast cross-wind flight before it can be efficiently reeled in while flying towards the zenith. Only after reaching the position depicted with the triangular marker, the system stops producing power and the kite is reeled in. The kite is powered up again after having approached the ground station up to roughly 100 m downwind. For now, only the flight between the latter two points is analysed (brown trajectory in Figure 2.11).

To evaluate how well the recorded variation in power consumption is captured by the simulation, simulations are conducted for the range of wind speeds encountered during the retraction phases in the flight test. The simulations assume a uniform wind profile. The transition phase is disregarded and the retraction phase is simulated starting from an elevation angle of 45°. The lift and drag coefficients of the wing are considered constant and the mean values that were found in the identification of Section 5.5 are used: $C_L = 0.41$ and $C_D = 0.11$. Also, the starting and ending tether lengths are fixed to 341 m and 269 m, respectively, as they introduce only little variation.

The controller tries to reel in with a constant tether force of roughly 1 kN while securing a minimal reel-in speed of roughly 1.75 m/s. This approach results in an increasing reel-in speed because the kite depowers as it approaches zenith. For a few pumping cycles, the reel-in speed is kept constant at the lower limit at the start of the reel-in phase if the wind is too strong to lower the tether force to 1 kN, e.g., due to a gust. Most pumping cycles, such as the reference pumping cycle, are however able to immediately reach the target tether force.

Figure 7.10 illustrates how the simulation compares with flight data for the reference pumping cycle. For this comparison, the tether length range and starting elevation angle are directly inferred from the flight data of the cycle. The simulated reeling speed agrees well with the trend observed in the flight data. Consequently, the duration is only slightly underestimated. The calculated average power is -3.2 kW against the -2.9 kW recorded in

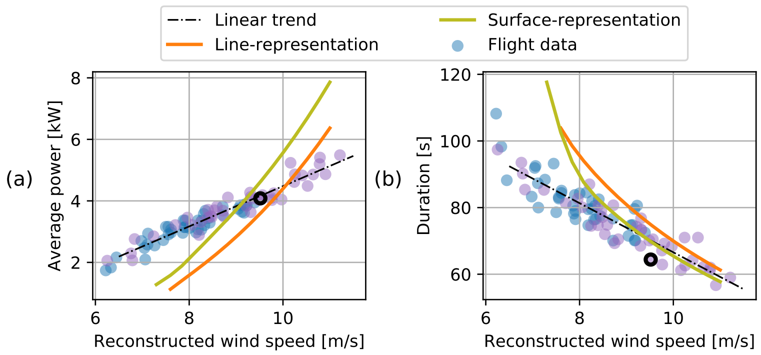

Figure 7.11: The inter-cycle variation of the average power and duration of the traction phase with the average reconstructed wind speed in the flight data. The blue markers belong to the first half of the flight test and the purple to the second. The reference pumping cycle is indicated with the black marker. The flight data is compared to the QSM results obtained with the two traction phase representations.

the flight data. The final position of the kite in the simulation and flight data also agree well, showing that the motion of the kite is reasonably solved.

Figure 7.12a shows that simulated power matches the trend in flight data at a reconstructed wind speed of roughly 10.2 m/s. Just as for the traction phase, the line that follows from the simulation has a larger slope than the trend in the flight data. In contrast, the duration of the retraction phase is approximated very well with the QSM, see Figure 7.12b.

7.4.3. Pumping cycle

The aim of the QSM is not to solve the motion realistically but to estimate performance indicators for longer periods such as a phase and ultimately the full pumping cycle. The original QSM includes a transition phase to close the flight trajectory after the retraction phase and before the traction phase. The modelling approach does however not capture the real transition flight well and thus is not a valuable addition to the mean cycle power estimation. Therefore, an alternative is sought to account for the transitions, including the transition after the traction phase.

Other models use a constant dead time where no power is produced or consumed to account for the transitions [47, 62]. As the total pumping cycle duration is decreasing with wind speed, the penalty on the mean cycle power due to a constant dead time increases with wind speed. Including a dead time effectively decreases the slope of the mean cycle power-wind speed relationship and leads to a mismatch with the slope of the mean cycle power trend observed in the flight data.

As an alternative to the dead time, empirically obtained corrections are introduced to better match the slope of the mean cycle power trend observed in the flight data. The pumping cycle is divided in two: the power production and consumption segments which are delimited by the triangular markers in Figure 7.3. The corrections to the power and duration of the traction phase are obtained with the help of the flight data of the production segment and those for the retraction phase using the flight data of the con-

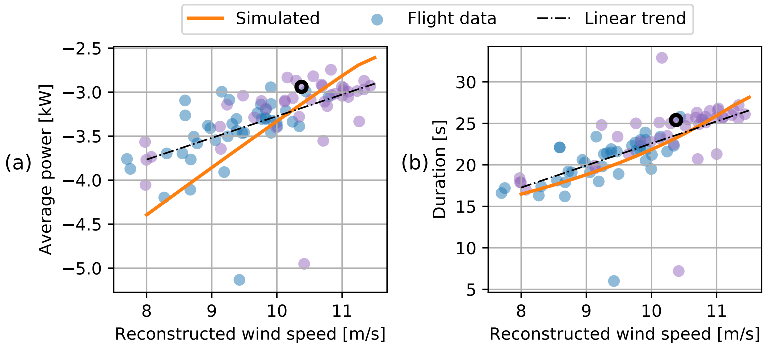

Figure 7.12: The inter-cycle variation of the average power and duration of the retraction phase with the average reconstructed wind speed in the flight data. The blue markers belong to the first half of the flight test and the purple to the second. The reference pumping cycle is indicated with the black marker. The flight data is compared to the QSM results of the retraction phase.

sumption segment.

Figure 7.13 shows how the previously obtained linear trends in the flight data of the traction and retraction phases capture the trend in the production and consumption segments well after being shifted vertically. Only changing the intercepts of the trends while keeping the slopes fixed leads to the best agreement with the flight data of the segments. The traction phase trends are shifted with -0.7 kW and 28.4 s and the retraction phase trends with -67 W and 7.6 s.

7

The summation of the segment duration trends yields the cycle duration trend shown in Figure 7.13f and the weighted average of the segment power trends yields the mean cycle power trend shown in Figure 7.13c. Although including a dead time provides a more general method, the proposed corrections provide a better agreement with the mean cycle power recorded in the flight data. The

The mean cycle power and cycle duration are obtained using the previous phase results after applying the corrections. The cycle curves are determined using both traction phase modelling approaches. The resulting curves are shown in Figure 7.13c and f. Both the mean cycle power curves obtained with simulations and from flight data are less steep than the traction phase curves. The traction phase has a dominant effect on the mean cycle power curves as it takes up most of the pumping cycle. Consequently, the mismatch between the slopes of the simulated and flight data curves persists. Below 8 m/s wind speed, the simulated curves drop below zero while the flight data curve is positive for the full range of evaluated wind speeds.

Based on the reconstructed wind speeds, the corrections seem valid for a substantial range of wind speeds. However, this may only appear to be the case due to the possible bias in the wind speed reconstruction. Data from more flight tests are needed to assess how valid the corrections are for other wind conditions. If the relative penalty would increase with wind speeds as suggested by the constant dead time approach, the calculated power in Figure 7.13c would be zero for smaller wind speeds.

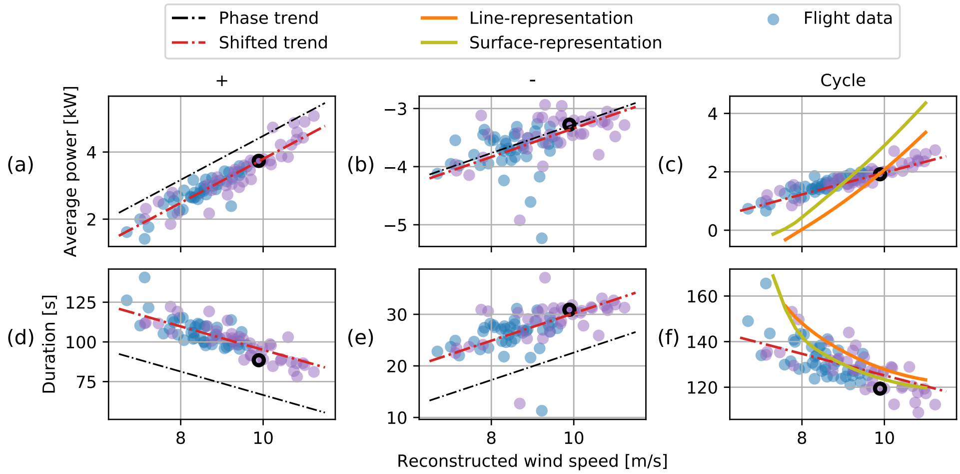

Figure 7.13: The inter-cycle variation of the average power and duration of the generation segment (**a**, **d**), consumption segment (**b**, **e**), and full cycle (**c**, **f**) with the average reconstructed wind speed in the flight data. The blue markers belong to the first half of the flight test and the purple to the second. The reference pumping cycle is indicated with the black marker. The average power and duration of the full cycle are compared to the QSM results.

7.5. Conclusion

The trend of the measured mean cycle power with wind speed is not well explained using the quasi-steady model (QSM). The dissimilarity of the trends of the modelled and flight data also emerges when assessing the performance estimation on the level of the figure-of-eight manoeuvres. Although the modelling of the figures of eight will cause some mismatch, it is expected that the observed discrepancy is predominantly explained by a deficiency in the wind speed reconstruction. The erroneous wind speed observations may yield a faulty trend between power and wind speed and, if overlooked, yield a distorted view of the model validity.

To independently validate power estimations for different sections of the pumping cycle, the cycle is dissected into smaller parts, with the smallest parts being the individual figures of eight. Since the kite is flying figures of eight for more than half of the time in a pumping cycle, the accuracy of the mean cycle power estimation relies heavily on the accuracy of the power estimation for figures of eight. The figures of eight flown in the flight test exhibit substantial variations in recorded power output, which is most likely explained by variations in the wind speed. Exponential functions are fitted to the figure-of-eight-average tether force observations. This exponent gives an indication of how much of the aerodynamic force effectively transfers to the tether. The trend of the observations with apparent wind speeds shows an increased effective exponent compared to the quadratic relationship prescribed by the massless kite theory.

The figure-of-eight power output as function of the wind speed is estimated using two representations of the traction phase in the QSM: the line- and surface representations. The line representation does not consider the cross-wind motion of the kite and, thereby, does not account for any variation within the figure of eight. The surface representation on the other hand moves the kite along a prescribed figure of eight which resembles those flown in the test flight. This cross-wind pattern is obtained with a newly developed realistic figure-of-eight parameterisation based on a class/shape function transformation. To include the effect of increased drag of the kite during the turns, the aerodynamic coefficients are incorporated as function of the position of the kite along the figure of eight. The power curve found with the surface representation intersects the flight data trend at the average reconstructed wind speed and 3.8 kW power. The line representation underestimates the power with 0.9 kW.

While the computed trends as a function of the apparent wind speed exhibit reasonable agreement with flight data, the agreement significantly diminishes when the trends and flight data are expressed against the reconstructed wind speed. This effect is explained by a faulty conversion of the measured apparent wind speed to the reconstructed wind speed, resulting from deficiencies in the wind speed reconstruction. The turning bias in the apparent wind speed measurement due to off-centre flow measurements is expected to be too small to independently induce the discrepancy.

The figure-of-eight power estimation is extended to cover the full traction phase and the resulting power curves are compared with flight data. Furthermore, simulations of the retraction phase are conducted to generate a power curve which is compared to flight data. The mean cycle power calculation accounts for the transitions with empirically obtained corrections instead of solving the relatively complex flight during the transitions. The mismatch of the calculated and observed power curve of the figures of eight also persists in the mean cycle power curve.

8

Optimising the flight operation for long-term performance assessment

The previous chapter assessed the accuracy of the flight operation model based on the limited range of wind conditions experienced during the 3-hour test flight. Despite a slight increase in wind speed, the operational settings were not altered once reliable flight behaviour was obtained. However, when operating year-round, the system encounters a large variety of wind profiles, requiring adjustments to the pumping flight operation to maximise energy yield. This chapter utilises the wind climate description from Chapter 4 to assess the significance of wind profile variability on the annual energy production.

8.1. Introduction

The economic viability of a wind energy converter highly depends on the energy yield, which in turn is highly sensitive to the site-specific wind climate. The cost is to a lesser extent dependent on the wind climate. To consider this in more detail, a specific wind energy converter is selected, and a siting study is carried out as part of the economic viability assessment.

In a typical preliminary siting study, the energy yield of a specific wind energy converter is estimated using a statistical summary of the wind climate at a specific location. This may be calculated by adapting regional wind data from a wind atlas or inferred from nearby long-term measurements to include the effect of the local terrain. Local corrections are made using short-term on-site measurements or microscale models like WAsP [36]. It is common practice for conventional wind turbines only to consider the wind climate at hub height. However, AWE systems have access to a larger height range beyond the reach of wind turbines. Therefore, the wind climate should be extended to account for the specifics of the variable height operation.

In addition to the wind climate, a characterisation of the system performance is needed to estimate the energy yield. For conventional wind turbines, the performance is commonly described with a curve of the power output as a function of the wind speed at hub height: the power curve. The power output of AWE systems relies more heavily on the wind profile to which the operation is tailored. Besides the magnitude of the wind profile, its shape can also vary substantially. Therefore, the power output is less well captured as a function of a single variable, i.e., the wind speed at a single height. Consequently, the spread of the power output around a power curve of an AWE system will be larger than for conventional wind turbines and may demand a more detailed energy yield calculation.

Luchsinger [120] applies the massless kite theory to derive a theoretical power curve of a pumping AWE system flying direct downwind. As for conventional wind turbines, the power curve can be divided into three regions: the region where power has a cubic dependency on wind speed, the region limited by the maximum allowable tether force, and the linearly decreasing region limited by the maximum power of the generator. The theory is expanded to account for elevated flight but does not cover the wind profile.

More detailed models of the pumping cycle flight [47, 62, 122] employ either the power law or the logarithmic wind profile to account for the variation of the wind with height. The resulting power curves are expressed as a function of the wind speed at a reference height or average operational height. Since the operational time of development systems and flight data is limited, the published power curves have only been compared with short-term measurements, which exhibit a substantial spread around the average power curves [47, 61].

Ranneberg et al. [47] study the effect of terrain roughness on the power curve by assessing the sensitivity of the power output to the roughness length parameter of the logarithmic wind profile relationship for a given wind speed at operating height. The study does not consider site-specific variations of the wind profile shape, e.g., due to atmospheric stability. The annual energy production (AEP) is estimated by combining hourly wind profile data with power curves for five operational heights instead of summarising the wind climate statistically. By assuming that the power output is independent of the exact wind profile shape, the height that yields maximum power output is determined for each point in time. This is a valid approach if the system sweeps a small height range. However, the exact wind profile shape becomes more relevant to consider when the system sweeps a larger height range, as is the case for flexible-kite AWE systems [46, 123].

Malz et al. [58] additionally refrain from using a power output characterisation when estimating AEP and optimise the performance of a fly-gen AWE system to produce maximum power output for each 3-hourly wind profile using 3 months of global reanalysis data. Arranging the optimal power outputs in descending order yields the power duration profiles. Additionally, a black box model was trained on the optimisation results from 15 locations to predict the power output with reasonable accuracy based on the wind profiles at another location. In follow-up work [59], the power duration profiles for the full year of 2016 are calculated for 20 locations. The four evaluated designs yield rather low capacity factors due to a low cut-in wind speed.

Instead of assuming an analytical wind profile, Sommerfeld et al. [57] infer power curves based on a small sample of one year of wind profiles simulated with the Weather

Research and Forecasting model with a 10-minute resolution. A targeted sampling approach is used based on the clustering of the simulated wind profiles. The operation of a pumping AWE system is optimised for each of the profiles in the sample. The optimal power outputs are mapped to the mean wind speed within the operating height range, and curve fitting is used to infer a power curve. Despite this fitting approach, the power curve shows significant fluctuation, suggesting an insufficient sample size or the requirement for a more detailed power output parameterisation.

This chapter hypothesises that characterising the wind climate and power output in terms of the wind profile shape produces an accurate AEP estimation for AWE systems. To test this hypothesis, the power output is characterised using multiple power curves for different wind profile shapes, and the representative wind profile shapes and magnitude distributions are identified from wind atlas data using the data-driven approach presented in Chapter 4. The power output of an AWE system is characterised by optimising the operation for the different wind profile shapes using the basic pumping flight operation model of Chapter 7. Combining the resulting power curves with the wind profile distributions enables the calculation of the AEP. To test the hypothesis, the AEP estimation is compared against a more computationally intensive brute-force approach which optimises the power output for each hourly wind profile in the dataset.

The following sections of this chapter outline the process of developing an efficient AEP estimation approach for an AWE system based on wind atlas data. Section 8.2 discusses how the basic pumping flight operation model should be employed in an optimisation approach to find what operating height range should be swept by the kite in order to maximise the power output of the AWE system. In Section 8.3, a series of optimisations are carried out for a logarithmic wind profile with increasing magnitudes to generate a baseline power curve. Section 8.5 quantifies the power output variation around the baseline power curve and illustrates the necessity of using multiple power curves. Subsequently, the prevailing onshore and offshore wind profile shapes are identified in Section 8.3. These are then used to generate the other power curves to finally calculate the AEP.

8.2. Pumping flight operation optimisation

To maximise energy production, the operation of a pumping AWE system should be tailored to the wind conditions. The system is controlled by steering the kite and controlling the torque acting on the winch at the ground station. The winch controller modulates the tether force, which affects the flight of the kite, including the reeling speed. To obtain a positive mean cycle power, the system operates at a high and low traction power during the traction and retraction phases, respectively. Increasing the percentage of time spent in the traction phase increases the mean cycle power. Additionally, the mean cycle power can be increased by planning the flight trajectory carefully such that the height range with the highest power potential is tapped by the kite.

Numerical offline optimisation can be used to determine the operational strategy that maximises the AEP of an AWE system. This requires planning the pumping cycle for many wind conditions, but most importantly for the prevalent wind conditions at a particular site. An operational strategy can be conceived by optimising the pumping cycle for each of these wind conditions. The computational cost of the AEP estimation can be reduced substantially by limiting the number of optimisations and reducing the computational cost of a single optimisation.

To find the optimal pumping flight operation, it suffices to employ a basic pumping flight operation model which is sensitive to the specifics of the input wind profile. Although a more detailed model may increase the accuracy and provide more details about the flight, a simpler model can still provide a good indication of how to plan the pumping cycle. The quasi-steady model (QSM) of Chapter 7 solves the motion of the kite with sufficient detail with respect to the wind profile and, thereby, is a good candidate for assessing the pumping flight operation in the AEP estimation.

The QSM is slightly modified to make it more suitable for optimisation. A fixed number of points is used to represent the flight trajectory. Five points are used to discretise the traction phase and ten for the retraction phase. The optimisation maximises the mean cycle power $\bar{P}_{\text{cycle}}$ using the following problem formulation:

$$\begin{aligned}
\max_{T_{\text{out}},\, T_{\text{in}},\, \beta_{\text{out}},\, l_{0,\text{out}},\, \mathbf{F}_{\text{t}},\, \kappa} \quad & \bar{P}_{\text{cycle}} \\
\text{s.t.} \quad & l_{0,\text{out}} = l_{f,\text{in}} \\
\text{For every point:} \quad & \sum F = 0 \\
& z_{\min} < z_{\text{k}} < z_{\max} \\
\text{During reel-out:} \quad & 0 < v_{\text{t}} < v_{\text{t,max}} \\
& v_{\tau} > v_{\tau,\min} \\
& P < P_{\max} \\
\text{During reel-in:} \quad & v_{\text{t}} > -v_{\text{t,max}} \\
& v_{\tau} > 0 \\
& P > -P_{\max}
\end{aligned} \quad . \tag{8.1}$$

Besides maximising power, the optimisation finds the steady state for every path point, i.e., $\sum F = 0$, by varying the kinematic ratios κ.

The duration of the traction and retraction phases, T_{out} and T_{in}, respectively, are used as optimisation variables, while the time step between the steady state points is kept uniform for each phase. A constraint is added to the optimisation to ensure tether length periodicity, i.e., the same starting length at the start of the traction phase $l_{0,\text{out}}$ and end of the retraction phase $l_{f,\text{in}}$.

The traction phase in the flight operation model is modelled with the line representation described in Section 7.3.2. The elevation angle of the traction phase β_{out} is included as an optimisation variable. Together with the tether length, the elevation angle of the reel-out path controls which height range the kite taps into. Thereby, it can be used to fly the kite in the height range with the highest power potential. Alternatively, it can also be used to depower the kite and alleviate the tether force. Controlling the elevation angle can thereby expand the wind speed range offering safe operation.

A lower limit $z_{\min}$=100 m is imposed to the kite height z_{k} to account for the safety margin that ensures sufficient clearance with respect to the ground. Moreover, a maximum height limit $z_{\max}$=500 m is imposed. The first commercial AWE systems envisage this ceiling to comply with airspace legislation.

The tangential speed v_τ is constrained to ensure that the kite flies cross-wind with sufficient speed along the full traction phase, i.e., completing at least 300 m of cross-wind flight which is representative for one cross-wind pattern. Therefore, the lower limit is expressed as a function of the duration of the traction phase: $v_{\tau,\text{min}} = \frac{300}{T_\text{out}}$ and is included in the form of an inequality constraint. Moreover, a positive tangential speed is imposed during the retraction phase to prevent the kite from flying backwards.

The aerodynamic coefficients of the wing are fixed to the values: $C_L = 0.7$ and $C_D = 0.16$ for powered flight and $C_L = 0.22$ and $C_D = 0.11$ for depowered flight. These values, except for the depowered lift coefficient, were previously found in Section 5.5. The depowered lift coefficient is lowered to a value considered more realistic for a less conservative operational approach than that of the test flight.

The use of a fixed number of trajectory discretisation points enables the tether forces $\mathbf{F}_\text{t}$ to be controlled for each point separately. This accommodates directly changing the tether force gradually along the traction and retraction phases. Note that this control approach is different from the approach in the test flight, where the tether speed-force ratio was kept constant. Imposing a lower bound of 750 N on the tether force ensures that the kite stays tensioned, as required for a flexible kite. The upper bound is imposed by the thickness and material strength of the tether and should guarantee a reasonable lifetime of the tether. A tether diameter of 6 mm is used for the baseline system configuration with a maximum allowable tether force of 8.2 kN.

The reeling speed v_t and power output P follow from each steady state calculation by controlling the tether forces. Inequality constraints are added to the optimisation to ensure that the resulting reeling speeds do not violate the maximum reeling speed limit $v_{\text{t,max}}$=10 m/s of the system. Moreover, a maximum power limit P_max=40 kW is used for the baseline system configuration. This limit is taken lower than the 100 kW rated power of the development platform of Kitepower to match the kite size and generator capacity better. The 100 kW limit will be addressed when analysing the sensitivity of the power curve to the generator's rated power. The transition phases are not solved explicitly in the optimisation model but are simply accounted for by a constant dead time where no power is produced or consumed. A larger dead time results in a lower mean cycle power. Based on the test flight data, a representative dead time of 17 s is employed. The addition of a dead time promotes relatively long traction phases that increase the duty cycle, i.e., the traction phase relative to the cycle duration.

The optimisation problem is solved using a gradient-based optimiser which in general, efficiently and robustly, finds a solution to the problem at hand. The Sequential Least SQuares Programming (SLSQP) algorithm of pyOpt [124] is used. SLSQP is a good general-purpose method for differentiable constrained non-linear problems.

8.3. Power curve for a logarithmic wind profile

Typically, the performance of AWE systems is characterised by assuming a neutral logarithmic wind profile (Equation 2.3) throughout the full height range of operation. Accordingly, the mean cycle power is expressed as a function of the wind profile magnitude with a single curve. To demonstrate the flight operation optimisation, the optimisation results are presented for a neutral logarithmic wind profile shown in Figure 8.1a. A roughness length z_0=0.1 m is used, which is a representative value for farmland with low crops. Flight operation optimisations are conducted for a series of neutral logarithmic wind profiles of varying magnitudes to generate the power curve.

Figure 8.2a shows the power curve resulting from the flight operation optimisations as a function of the wind speed at a 200 m reference height. The wind speed at any height and the wind profile shape together fully define the steady wind profile considered by the pumping flight operation model. In the case of no variation in wind speed profile shape, the choice for the reference height determines the scale on the x-axis but does not affect the shape of the power curve. The choice for the reference height becomes non-trivial when variations in the wind profile shape are considered, as is the case in the AEP estimation. The reference height choice is more extensively discussed later on in this chapter when discussing the AEP estimation.

The lowest and highest wind speeds for which the optimisation finds a feasible solution indicate the cut-in and cut-out wind speeds for the AWE system. For the investigated system, these are $v_{\mathrm{w,200m}}$ = 5.8 m/s and $v_{\mathrm{w,200m}}$ = 24.5 m/s, respectively. Outside this range, the system does not produce any power. The pumping AWE system only just generates a net positive power output at cut-in. The power curve of a pumping AWE system does not exhibit a plateau above the rated wind speed that is imposed by the generator capacity like the power curve of modern utility-scale tower-based wind turbines. For a pumping AWE system, the generator capacity only indirectly limits the mean cycle power through the traction power.

Figure 8.1b shows the optimal pumping cycle trajectories. The elevation angle of the traction phase (straight path) increases with wind speed. At low wind speeds, the reel-out elevation angle increases slowly, and, at the same time, the range of elevation angles swept during the retraction phase increases. Moreover, the pumping cycle grows taller until just before cut-out. The upper curved path at cut-in indicates that the kite is climbing in height throughout the full retraction phase, whereas the kite continuously descends at cut-out. At high wind speeds ($v_{\mathrm{w,200m}}$ >14.8 m/s), the kite reaches a fixed elevation retraction path towards the end of the phase.

8

Figure 8.2c and d show that traction phase duration decreases with wind speed and the retraction phase duration peaks around $v_{\mathrm{w,200m}}$=15 m/s after which it mostly decreases again with higher wind speeds. The peak of the retraction phase duration occurs at the lowest wind speed at which the kite reaches its maximum height, see Figure 8.3i. For higher wind speeds, the kite climbs faster at the start of the retraction phase, and the duration is lowered to prevent the kite from exceeding the maximum height limit. The lower limit imposed on the retraction phase duration prevents finding an unrealistically small pumping cycle at cut-in.

The duty cycle shown in Figure 8.2b follows from the traction and retraction phase durations combined with the dead time. At cut-in, the kite can be reeled in very quickly during the retraction phase while being reeled out slowly during the traction phase yielding a duty cycle of almost 90%. The duty cycle decreases with wind speed due to the increasing reel-out speed during the traction phase and the decreasing reel-in speed during the retraction phase, see Figure 8.3c and d.

Figure 8.2e and f show that employing a relatively large tether length with associated drag is more optimal at cut-in than a high elevation angle with associated wind misalignment. Above cut-in, this trade-off changes and results in a higher elevation angle as the tether drag penalty increases with the increase in tangential speed of the kite (Figure 8.3g). Below the wind speed at which the maximum power is reached, the trade-off aligns with the maximum traction power point depicted with the pink lines in Figure 8.2e and Figure 8.3i.

After the system runs into some of its operational limits, depowering mechanisms are employed to expand the wind speed range allowing safe operation. The following mechanisms may be used by the optimiser for depowering the kite during the traction phase:

1. Decreasing/increasing the tether force away from the tether force that maximises power output;
2. Misaligning the pulling force of the kite and wind velocity by increasing the elevation angle;
3. Avoiding the height range with the highest power potential;
4. Flying with a longer tether to increase tether drag.

In reality, there are more ways to depower the kite. A very effective depowering mechanism is to pitch the kite nose down. The model does not consider gradual depowering by pitching during the traction phase. It does differentiate between the pitch during the traction and retraction phases using two combinations of aerodynamic coefficients.

At $v_{\mathrm{w,200m}}$=13 m/s, the tether force at the start of the traction phase reaches its maximum limit, and the controller stops tracking the tether force that maximises the power output. This is equivalent to employing the first depowering mechanism because the tether force is kept below the optimal value. For higher wind speeds, the depowered fraction of the traction phase increases until the tether force along the full traction phase is at its maximum limit at $v_{\mathrm{w,200m}}$=16 m/s. The depowering results in a more rapid increase in reeling speed (see the blue line in Figure 8.3c for 13<$v_{\mathrm{w,200m}}$<16 m/s) and a decrease in tangential speed (Figure 8.3g).

Similar to the plateau of the power curve of a conventional wind turbine, the traction power curve in Figure 8.3e exhibits a plateau above $v_{\mathrm{w,200m}}$=16 m/s. The corresponding traction power limit is imposed by the rated power of the generator. Coincidentally, the wind speed at which the power limit is reached coincides with the wind speed at which the tether force at the end of the traction phase reaches its maximum limit. Once the maximum traction power limit is reached, the kite is depowered by increasing the elevation angle more rapidly (Figure 8.2e). This is equivalent to employing the second depowering mechanism to actively misalign the pulling force of the kite and wind velocity.

The last depowering mechanism employed just below the cut-out is adapting the trajectory to try to avoid flying through the height range with the highest power potential. This provides a more optimal way of depowering than reducing the tether force below its maximum limit. The height with the highest power potential is specific to the system and changes with the wind profile magnitude. For the optimal elevation angle, the maximum power point height is depicted with the brown line in Figure 8.3i. Above $v_{\mathrm{w,200m}}$=22 m/s, the pumping cycle is shrunk by reducing the traction phase duration in an attempt to fly below the maximum power point. At the same time, the tether force needs to be lowered below its maximum limit for additional depowering.

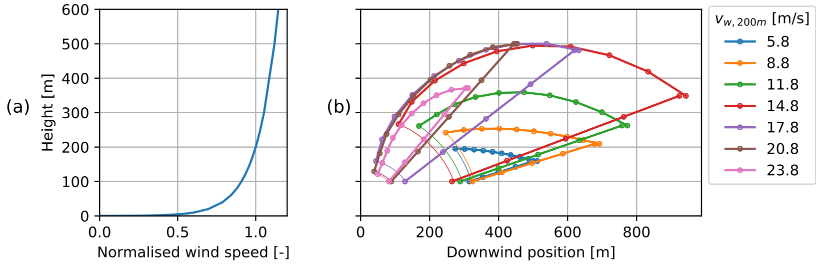

Figure 8.1: (**a**) Neutral logarithmic wind profile shape for an onshore location with roughness length z_0=0.1 m used to generate the baseline power curve. (**b**) Side-views of the optimal pumping cycle trajectories for a range of wind speeds. Only the outline of the trajectory is modelled in the QSM.

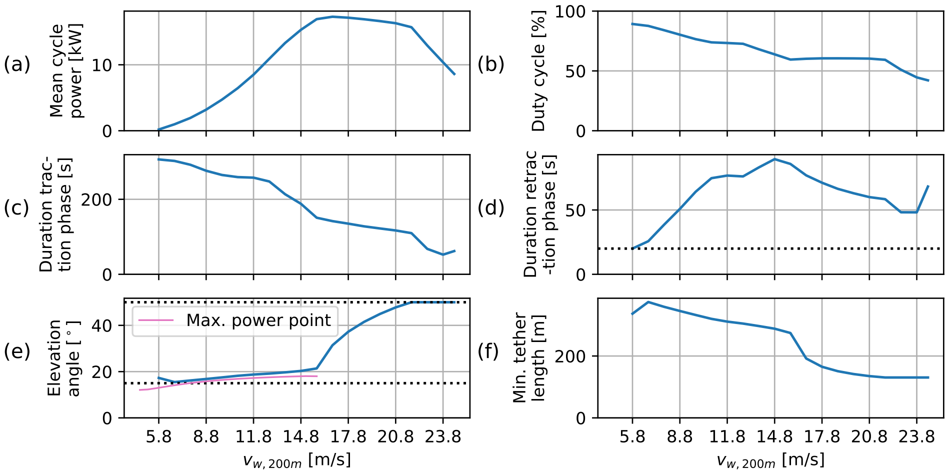

8

Figure 8.2: (**a**) Power curve and (**b**) optimal duty cycle as a function of the wind speed found with the baseline optimisations. (**c**–**f**) Optimal values found for the optimisation variables characterising the pumping flight trajectory. The dotted lines depict the variable bounds used in the optimisations.

The lower limit on the tangential speed of the kite increases towards cut-out due to the decreasing traction phase duration, as depicted in Figure 8.2g. At the same time, the depowering mechanisms result in a decreasing tangential speed. After the tangential speed of the kite at the end of the traction phase reaches the lower limit, feasible operation of the pumping cycle is no longer possible and, thereby, the cut-out is reached.

Note that increasing the tether force is not employed as depowering mechanism for the investigated system. Instead, the tether force is decreased for depowering. If the tether force limit would be reached after reaching the power limit (e.g. for a system with a thicker tether), depowering by increasing the tether force becomes an option and may yield a higher mean cycle power. However, this approach is not optimal in practice as the kite cannot easily be further depowered once the maximum tether force limit is reached and compromises operational safety.

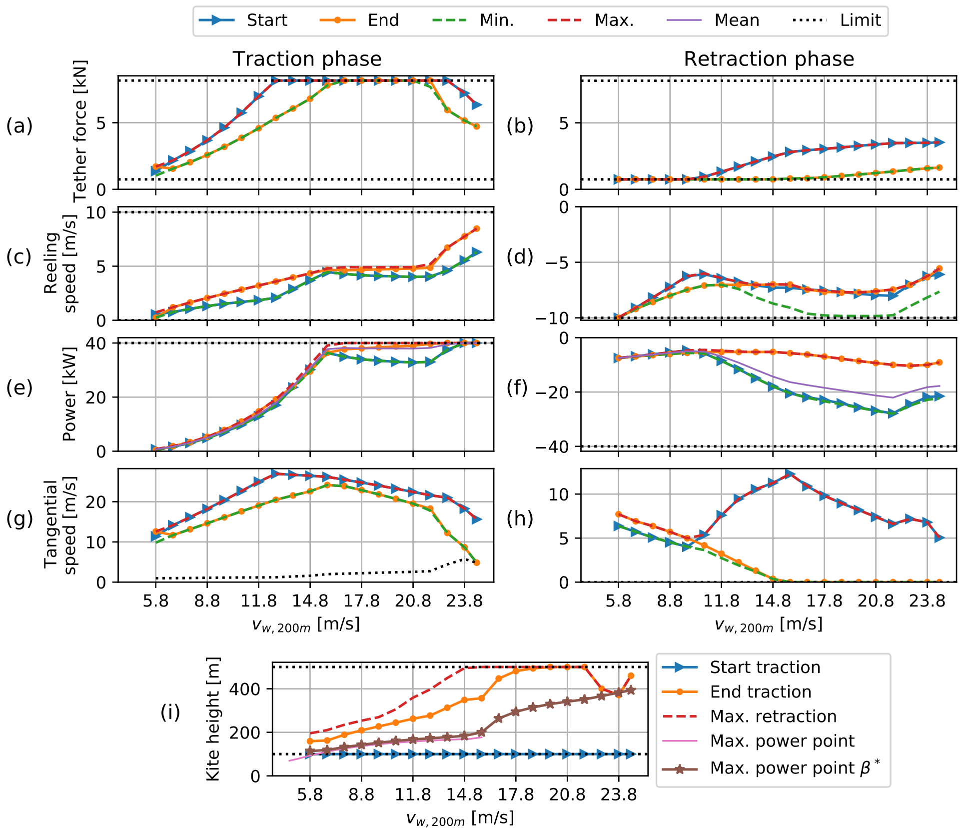

Figure 8.3: Flight properties during the traction and retraction phase for each optimal pumping cycle trajectory resulting from the baseline optimisations. (**a**–**h**) Optimal tether force, reeling speed, power, and tangential speed as a function of reference wind speed during the two phases. The tether forces are included as optimisation variables and thus controlled directly. (**i**) Optimal height range swept by the kite.

8

8.4. Variation around the power curve

The power output of a wind energy converter is typically presented as a function of a single variable. However, in reality, the power output depends on many variables. Therefore, the actual output will never strictly adhere to a fixed power curve. A substantial part of the power output variation around the power curve can be attributed to the variation of the wind conditions. AWE systems are expected to encounter larger variations in wind conditions than conventional wind turbines as the kites sweep a larger volume than wind turbine blades.

A preliminary power curve can be produced by making assumptions and simplifying the wind conditions. The power curves of AWE systems commonly assume a neutral logarithmic wind profile to describe the wind conditions up to heights exceeding the surface layer, even though the logarithmic wind profile is not strictly valid above the surface layer. Consequently, the spread around the power curve of an AWE system is expected to be substantial.

Alternatively, the power curve can be approximated without assuming a wind profile relationship with the use of historical wind data. Power output statistics are obtained by optimising the power output for every wind profile in the dataset. The mean power output as a function of the wind speed yields the power curve. This approach is used to generate a power curve based on the Dutch Offshore Wind Atlas dataset for the on- and offshore reference locations introduced in Section 4.2. Figure 8.4 shows the resulting statistics for 2008 as a function of the wind speed at 200 m and 300 m height.

Figure 8.4a compares the statistically obtained power curve against the power curve assuming a logarithmic wind profile with roughness length z_0=0.03 m for the onshore location. This value falls in the lower end of the measurement-inferred roughness lengths for the area surrounding the mast Cabauw [101]. Up to the nominal wind speed around $v_{\mathrm{w,200m}}$=16 m/s, the logarithmic power curve slightly overestimates the power output. Above the nominal wind speed, the logarithmic power curve shows a stronger decline in power, and the statistical power curve shows a higher power output above $v_{\mathrm{w,200m}}$=22 m/s.

The mismatch between the curves is largest directly above the nominal wind speed, where the kite sweeps a larger height range during the traction phase. Thereby, an erroneous assumption with respect to the wind profile shape increasingly contributes to larger wind speed modelling errors and consequently power output errors. In summary, a high power output error is expected when the kite sweeps a large height range, and the wind profile does not resemble a neutral logarithmic wind profile.

The power output statistics also allow a study of the variation in the power output. The variation around the mean curve is modest, indicating a high correlation with $v_{\mathrm{w,200m}}$. In particular, at low and high wind speeds the average operational height in the traction phase is around 200 m making the power output less sensitive to the wind profile shape. To conclude, $v_{\mathrm{w,200m}}$ is a good predictor of the power output given the way the system is operated. However, this may be different for different AWE systems.

Figure 8.4c shows the same statistics but plotted against the wind speed at 300 m height. Note that the logarithmic power curve is slightly moved to the right as the wind speed is generally higher at 300 m than at 200 m. Also, the shape of the mean curve has changed, leading to a higher mismatch with the logarithmic power curve below the nominal wind speed due to the different binning of the wind profiles (with respect to

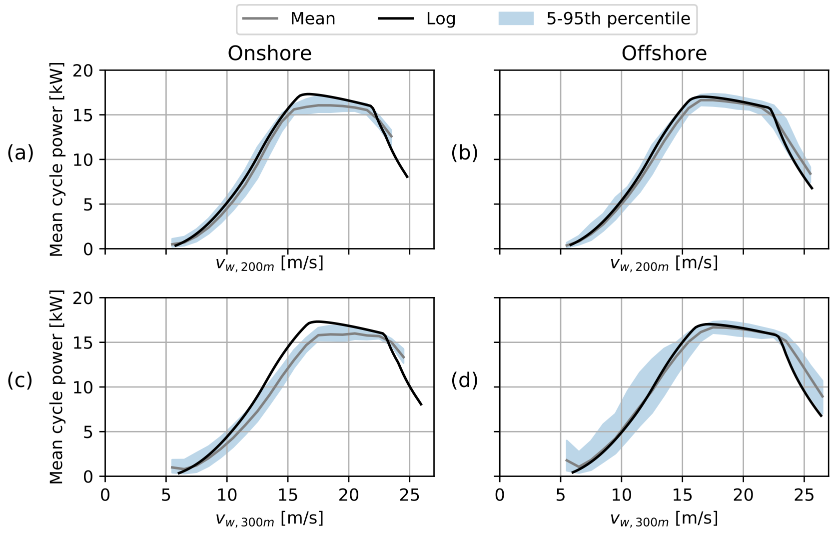

Figure 8.4: The mean and variation of the power output as a function of the wind speed inferred from the brute-force optimisation results for the onshore and offshore locations. These quantities are compared against the baseline optimisation results and expressed with respect to the wind speeds at 200 m (**a**, **b**) and 300 m height (**c**, **d**).

$v_{\mathrm{w,300m}}$ opposed to $v_{\mathrm{w,200m}}$). The variation around the mean curve at the lowest wind speeds has slightly increased as the wind profiles with similar power output are more frequently placed in different wind speed bins.

The same analysis is repeated for the offshore location with the results shown in Figure 8.4b and d. The mean curve shows a good agreement with the power curve based on the neutral logarithmic wind profile with roughness length z_0=0.0002 m for all wind speeds. This suggests a better match of the mean wind profile shape with that of the logarithmic wind profile compared to the onshore analysis. The variation around the mean curve increases substantially when expressed against $v_{\mathrm{w,300m}}$. This indicates that the variation of wind profile shapes within the bins is larger for the offshore location.

The brute-force approach does not rely on a modelled wind profile shape but quickly becomes computationally expensive, especially when using a more detailed pumping flight operation model. The results show that it may be sensible to characterise the power output as a function of the wind speed at a carefully selected height. However, this does not mean that it suffices only to consider a single wind profile shape. Alternatively, the power output can be characterised using a set of power curves for fixed wind profile shapes.

8.5. Wind profile climate description

As part of the proposed AEP estimation for AWE systems, a small group of dominant wind profile shapes is used to characterise the wind climate instead of considering the full spectrum of shapes. The prevalence of the wind profile magnitude for each of these shapes is described with probability distributions. Combined, the wind profile shapes and magnitude distributions yield a description of the wind profile climate.

8.5.1. Representative wind profile shapes

The wind profile shapes considered in the climate description are chosen carefully. The data-driven approach presented in Chapter 4 is slightly modified to identify the dominant wind profile shapes. To align the wind climate description with the input of the pumping flight operation model, only the vertical profile of the wind speed is considered and not the wind direction profile. Although the wind direction profile needs to be considered when planning the flight trajectory, the effect on the power output is assumed to be negligible.

The wind profile climate is site-specific and strongly depends on the orography and atmospheric conditions. Consequently, employing the data-driven approach for choosing the wind profile shapes leads to a site-specific power output characterisation. Having a site-specific characterisation is not very desirable. However, the same approach could in principle be used to generate a more general set of wind profile shapes representative of a larger region. For the purpose of this chapter, the analysis is confined to the on- and offshore reference locations.

The variation of wind profile shapes is investigated with a 2D histogram in the space spanned by the first two principal components (PCs), see Figure 8.6. The following operations are carried out on the wind atlas data to obtain the data in the new coordinate system:

1. The wind profiles in the dataset that have a mean wind speed smaller than 5 m/s are left out;
2. The wind profiles are normalised with their maximum wind speed;
3. The PCs (Figure 8.5b and c) are computed, and the data is transformed to the PC coordinates;
4. The neutral logarithmic wind profile (Figure 8.5a) is transformed to the PC coordinates, and then the data is translated such that the neutral logarithmic wind profile coincides with the origin.

89.5% of the variance in the data is contained in the first two PCs for the onshore dataset and 86.1% for the offshore dataset. Therefore, the difference between possible modes in the distribution of wind profile shapes can primarily be observed in the space spanned by the first two PCs. A denser data concentration in the 2D histogram indicates a dominant wind profile shape.

The orange lines in Figure 8.6 give an indication of where the logarithmic wind profile for different stability conditions is positioned in the PC-projection. To obtain the log-line, the stability-dependent logarithmic wind profiles calculated with Monin-Obukhov similarity theory (Equation 2.3) are normalised with respect to the wind speed at 600 m height and transformed to the new coordinate system. The line shows the resulting coordinates ranging from very unstable (left), to neutral (origin), to very stable (right) stratification. Along the line, the Obukhov length boundaries of the stability classes of Table 2.1 (left to right: -200, -500, 500, and 200 m) are marked with vertical bars. The neutral region covers a large part of the line.

The log lines show some resemblance with the data spread, e.g., departing from the lower left corner the trend flattens for both the data cloud and the log line. Also, with the log line as a reference, it remains problematic to identify the wind profiles that adhere to the logarithmic wind profile in the PC-projection. Note that for the comparison with the logarithmic wind profiles, it would be more sensible to confine the analysis to the surface layer in which the wind profile relationship is strictly valid. The surface layer has a maximum depth of ~100 m. Currently, wind profiles with the same shape up to 100 m may be positioned differently on the projection depending on the upper part of the wind profile (100–600 m). Therefore, only wind profiles that adhere to the logarithmic wind profile over the full 600 m depth will be positioned in the neighbourhood of the log line in the projection plot.

A set of wind profile shapes representative of the dominant shapes is obtained by manually clustering the wind profile shape data. A manual approach is opted for instead of using an algorithm to deal with unevenly sized clusters and consider the connectivity of the data. Simple clustering algorithms such as k-means clustering used in Chapter 4 are less suited for this. The data is divided into six clusters: three for the lower data strip (i.e. data points located approximately below the x-axis for the onshore location) and another three for the data cloud located above. The cluster boundaries are chosen such that the area for each of the two data substructures is approximately evenly partitioned. The cluster boundaries are set with the quadrilaterals in Figure 8.6. The mean-cluster wind profile shapes are shown in Figure 8.7.

The first and second mean-cluster wind profile shapes can be classified as (very) unstable and neutral logarithmic wind profiles, respectively, based on their shear. For the offshore location, the third cluster resembles the logarithmic wind profile with an Obukhov length of approximately 500 m, whereas for the onshore location, such logarithmic wind profile is positioned directly between the third and fourth clusters. The remaining clusters exhibit unalike features, such as jet-like shapes of the sixth onshore cluster and the remaining offshore clusters. Note that although the second and third onshore mean-cluster wind profile shapes and the logarithmic wind profile shapes appear different in Figure 8.7a, up to 200 m height they are actually very similar as shown in Figure 8.7c.

8.5.2. Wind profile magnitude probability

This section does not directly contribute to the targeted climate description but provides relevant insight into the relationship between the shape and magnitude of the wind profiles. The previous section evaluated the wind profile shape in a 2D space. To visualise the full wind profile distribution, a third dimension is introduced.

Each wind profile in the wind atlas data $\hat{\mathbf{v}}_{\mathrm{w}}$ is decomposed using three basis profiles, i.e., one profile for every axis. The three basis profiles are the normalised neutral logarithmic wind profile $\tilde{\mathbf{v}}_{\log}$ and the PCs, $\tilde{\mathbf{v}}_{\mathrm{PC1}}$ and $\tilde{\mathbf{v}}_{\mathrm{PC2}}$. The best fitting linear combination of the three basis profiles is sought using the following least squares fitting problem:

$$\min_{\mathbf{k}} \quad \left\| \hat{\mathbf{v}}_{\mathrm{w}} - \left[\tilde{\mathbf{v}}_{\log}\ \tilde{\mathbf{v}}_{\mathrm{PC1}}\ \tilde{\mathbf{v}}_{\mathrm{PC2}} \right] \mathbf{k} \right\|^2, \tag{8.2}$$

in which the fitting parameters $\mathbf{k} = [v_{\log,200\mathrm{m}}\ k_{\mathrm{PC1}}\ k_{\mathrm{PC2}}]^T$ are the weights of each basis

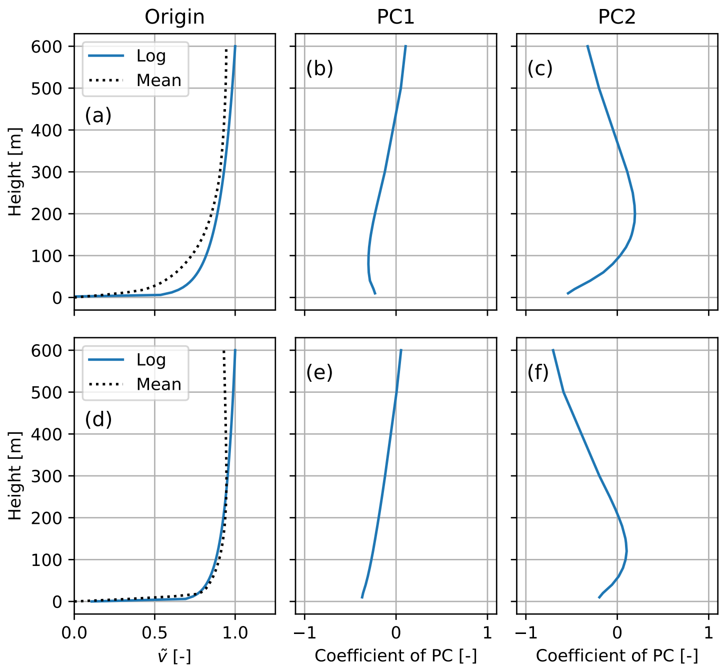

Figure 8.5: Mean (**a**) and first two principal components (**b, c**) in the onshore wind profile shape data. (**d–f**) The same plots for the offshore wind profile shape data. The principal components are used to express the data with just two variables. The coordinate system is translated such that the origin corresponds to the neutral logarithmic wind profile instead of the mean shape.

profile.

The resulting 3D density distributions are depicted in Figure 8.8a and Figure 8.9a for the onshore and offshore locations, respectively. To enable a visualisation, each distribution is dissected using 2D slices along the x-axis (Figure 8.8b–g and Figure 8.9b–g). The elevation and azimuth angle determine the shape of the wind profiles, i.e., two wind profiles that lie on a line passing through the origin have the same shape. The radial coordinate of each data point signifies the magnitude of the wind profile. The wind profile shape distribution of Figure 8.6 can approximately be obtained by projecting the 3D distribution onto a sphere centred around the origin and intersecting with the x-axis at $\tilde{v}_{\mathrm{log,200m}}$=0.87.

In the wind profile distributions, the normalised neutral logarithmic wind profile and the cluster-mean wind profile shapes are represented by lines. For each slice of the distribution, the intersections of these lines are depicted with asterisks. They have approximately the same relative position as the clusters in Figure 8.6 with some discrepancies resulting from only considering the first two PCs in the wind profile decomposition (Equation 8.2).

Inspecting the series of distribution slices enable identifying trends of the wind profile shape with wind speed. At calm wind conditions (Figure 8.8b and Figure 8.9b), the spread of the mean-cluster wind profile shapes (asterisks) is small compared to the

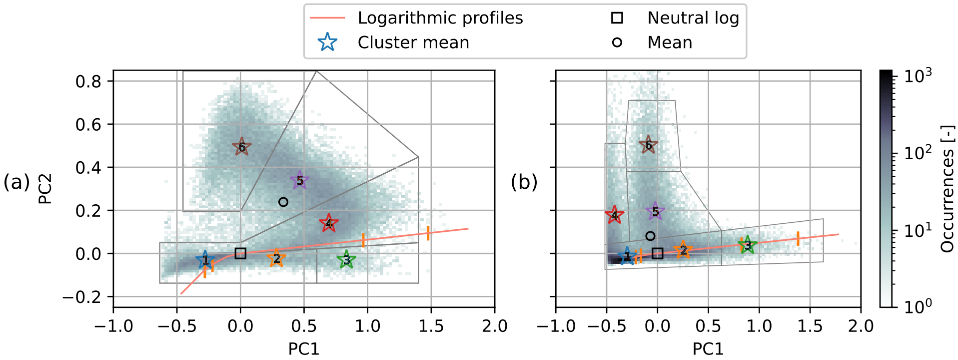

Figure 8.6: The distribution of the shapes of the wind profiles in (**a**) the onshore dataset and (**b**) the offshore dataset expressed with respect to the first two principal components. For reference, the orange line indicates the shape of the logarithmic wind profile for very unstable (left), to neutral (origin), to very stable (right) stability conditions. Along this log line, the Obukhov length boundaries of the stability classes of Table 2.1 (left to right: -200, -500, 500, and 200 m) are marked with vertical bars. The quadrilaterals depict the manually chosen cluster boundaries. The asterisks mark the mean shapes of each cluster.

spread of the distribution, indicating a large variety of wind profile shapes. As the wind speed increases, the wind profile shape variation decreases. The onshore distribution becomes bi-modal for light wind conditions (Figure 8.8c) with peaks close to well-mixed (blue asterisk/cluster 1) and jet-shaped wind profiles (brown asterisk/cluster 6). For strong winds (Figure 8.8f), a peak is found between clusters 3 and 4 (green and red asterisks, respectively) indicating that the wind profiles exhibit less mixing. Similar trends can be observed in the offshore distribution (Figure 8.9b–g)

8.5.3. Discretisation of wind profile shape

The wind profile distribution is discretised by wind profile shape to obtain the targeted climate description referred to as the cluster representation. The considered wind profile shapes are the normalised neutral logarithmic wind profile and the six mean-cluster wind profile shapes. Each wind profile $\hat{\mathbf{v}}_{\mathrm{w}}$ is represented by the best fitting shape i given by:

$$\min_{\mathrm{i}} \left\{ \epsilon_{\mathrm{log}}, \epsilon_1, \ldots, \epsilon_{\mathrm{N_{clusters}}} \right\}, \tag{8.3}$$

where ϵ_i is the sum of weighted squared errors determined for every normalised wind speed profile $\tilde{\mathbf{v}}_{\mathrm{i}}$ by solving a least squares problem:

$$\epsilon_{\mathrm{i}} = \min_{v_{\mathrm{w,200m}}} \left[\mathbf{w} \left(\hat{\mathbf{v}}_{\mathrm{w}} - v_{\mathrm{w,200m}} \tilde{\mathbf{v}}_{\mathrm{i}} \right) \right]^2. \tag{8.4}$$

Weights $\mathbf{w}$ are assigned to each height in the unevenly spaced vertical grid of the wind atlas data (Figure 4.5d) such that the fit is uniformly weighted over the vertical operating range of the system.

The discretisation is analogous to projecting each data point in Figure 8.8a and Figure 8.9a onto the closest line representing one of the considered wind profile shapes.

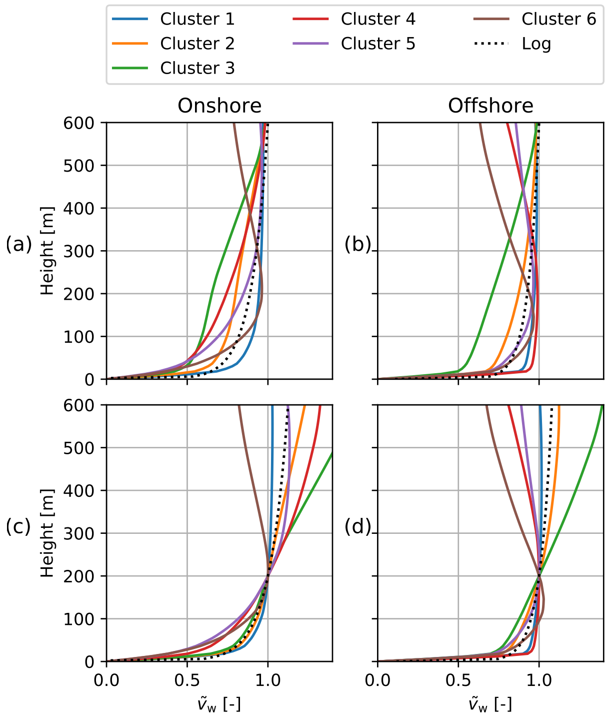

Figure 8.7: Onshore and offshore cluster-mean wind profile shapes normalised with respect to the maximum wind speed (**a**, **b**) and wind speed at 200 m height (**c**, **d**).

Accordingly, the multivariate probability distribution can be substituted with seven univariate probability distributions, i.e., one for every considered wind profile shape. The resulting distributions for the onshore location are shown in Figure 8.10b-g.

Figure 8.10a shows that the wind speed distribution at 200 m height in the wind atlas data agrees well with the wind profile magnitude distribution found for the cluster representation. For comparison, also the distribution is shown that results from an approach in which every wind profile is represented with a neutral logarithmic wind profile. The good agreement between the wind atlas and cluster representation distribution verifies that the wind profile shape discretisation with a small set of shapes did not substantially degrade the wind speed statistics in contrast to the logarithmic representation. For the latter approach, the resulting distribution shows a significant overestimate of the probability of high wind speeds at 200 m.

8.6. Efficient annual energy production estimation

Now that the sets of representative wind profile shapes have been identified, the power curves needed to carry out the proposed AEP estimations can be generated. The procedure described in Section 8.3 is used to derive a power curve for each mean-cluster wind profile shape. The resulting power curves are depicted in Figure 8.11 alongside the power curve for the neutral logarithmic wind profile.

The following observations are based on the onshore power curves (Figure 8.11a), however, similar observations can be made for the offshore power curves. For $v_{\mathrm{w,200m}}$=15.5–22 m/s, the logarithmic-based power curve and the power curve of the second cluster exhibit the highest power output. For a large range $v_{\mathrm{w,200m}}$=6.5–22 m/s, the logarithmic-based power curve is higher than the mean power curve of the hourly brute-force optimisation.

The power curve of the third cluster, which closely resembles a stable logarithmic wind profile, exhibits the highest power output at low values of $v_{\mathrm{w,200m}}$. The higher power output is realised by tapping into the relatively strong winds higher up compared to the other wind profile shapes (for equal $v_{\mathrm{w,200m}}$). Just as for the fourth and fifth clusters, the optimal pumping cycle at cut-in of the third cluster exhibits a higher operating altitude, whereas the cut-in operating altitude for the other clusters is just above the 100 m height limit. At high values of $v_{\mathrm{w,200m}}$, the relatively strong winds higher up turn into a disadvantage as it requires depowering the kite earlier.

In particular, for the offshore location, the cut-out wind speed of the sixth cluster appears to be relatively high. This can be explained by the choice of the height at which the wind speed is expressed. The maximum of the jet-like mean-cluster profile lies close to the 200 m reference height. Although the kite would be overpowered when operating in the jet, the system can still operate safely above the jet. Note that the sixth cluster only occurs for $v_{\mathrm{w,200m}} < 17$ m/s (see Figure 8.10h). This implies that employing this depowering approach will not be relevant in reality.

Unexpectedly, the power curves fall in the upper range of the power outputs determined with the brute-force approach for low wind speeds. Only around $v_{\mathrm{w,200m}}$=16.5 m/s, some power curves intersect with the mean power curve and drop below the 5$^{\mathrm{th}}$ percentile around $v_{\mathrm{w,200m}}$=19 m/s. The relatively high values of the power curves may be explained by the averaging used to obtain the mean-cluster wind profile shapes. Detailed features of the wind profiles that reduce the power output such as fluctuations are likely to be averaged out by the clustering.

The power curves of the onshore location are also displayed in Figure 8.10b–h alongside the wind speed distributions. The part of the wind speed distribution for which operation is infeasible is depicted in orange. A substantial part of the distribution lies below the cut-in of the investigated AWE system and, therefore, will not contribute to the energy production. On the contrary, there are virtually no occurrences of wind conditions above cut-out.

The computed power curves are used to calculate the average generated power of the AWE system:

$$\bar{P} = \sum_{i=1}^{n_\mathrm{c}} \int_0^\infty p_i(v_{\mathrm{w,200m}}) \cdot P_i(v_{\mathrm{w,200m}})\,\mathrm{d}v_{\mathrm{w,200m}} \approx \sum_{i=1}^{n_\mathrm{c}} \sum_{j=1}^{n_\mathrm{b}} \frac{f_{ij}}{n_\mathrm{s}} \cdot P_i(v_{\mathrm{w,200m}}) \quad , \tag{8.5}$$

in which p_i is the wind speed probability and P_i is the power curve of the i^{th} cluster. n_c is the number of clusters. The integral in the expression is solved numerically using $n_\mathrm{b} = 30$ wind speed bins of equal width between cut-in and cut-out. In the resulting right-hand side expression, the number of samples n_s is used to normalise the frequency f_{ij} of the i^{th} cluster and j^{th} wind speed bin.

Table 8.1: Combinations of wind profile shapes used to evaluate the sensitivity of the AEP calculation.

Set name	Considered wind profile shapes Onshore	 Offshore
Log (benchmark)	Neutral logarithmic wind profile	Neutral logarithmic wind profile
Set #1	Log & cluster 1–3	Log & cluster 1–3
Set #2	Log & cluster 1–4	Log & cluster 1–3, 5
Set #3	Log & cluster 1–4, 6	Log & cluster 1–3, 5, 6

To obtain a preliminary AEP estimate, the average generated power is multiplied by the hours in a year. More realistic AEP estimates should also consider the availability of the machine to account for maintenance, take-off limitations, poor visibility, icing, and lightning [47].

To investigate how sensitive the AEP estimation is to which wind profile shapes are considered, the calculation is repeated with reduced sets of wind profile shapes. First, the AEP is estimated based only on the neutral logarithmic wind profile as conventionally used in basic AEP estimations. In the subsequent estimation, the first three mean-cluster shapes are added, which is approximately equivalent to considering logarithmic wind profiles with different atmospheric stabilities. In the next steps, the remaining mean-cluster shapes are added one by one in descending order of their cluster size. This scheme enables evaluating how sensitive the calculation is to the jet-like wind profile shapes. The sets of wind profile shapes used in the series of AEP estimations are listed in Table 8.1. For every calculation over again, the wind profile instances in the wind atlas data are represented with the best fitting wind profile shape contained in the evaluated (reduced) set using the approach described in Section 8.5.3.

Figure 8.12 shows that the AEP estimate for 2008 is converging to the brute-force result as the number of considered wind profile shapes increases. This suggests that including more wind profile shapes in the calculation increases the precision of the AEP estimate. Thereby, the number of wind profile shapes to consider in the calculation can be chosen such that accuracy requirements are met. How much accuracy is required depends on the application of the AEP estimation. This conclusion is reinforced by the observed consistency in how the AEP converges for each year. This consistency underlines that the clustering approach yields a structural decomposition of the wind climate.

Including all seven wind profile shapes shown in Figure 8.7a yields an average AEP of 39.8 MWh over the years 2008–2017 for the onshore location (rightmost point of the blue line in Figure 8.12a). In between the years, the variation in AEP is substantial. The AEP for the year 2010 exhibits the highest difference and is 18% lower than the 10-year-average. The result for 2008 with all wind profile shapes is compared against the results of the brute-force calculations (dashed line), which is only carried out for 2008 in view of the computational cost. The observed AEP underestimation is only 0.5%. This difference is small given that the power curves poorly cover the 5$^{\text{th}}$–95$^{\text{th}}$ percentile power band that results from the brute-force calculations. The result based on the neutral logarithmic wind profile (Log) shows a relatively large AEP overestimation of almost 5%. This indicates that a substantial improvement in accuracy is obtained with the cluster profile based AEP estimation.

The effect of the jet-like wind profile shapes is studied by evaluating the AEP difference from one estimation to the next. The pronounced low-level jet profile (cluster 6) is introduced in the second-to-last step in Set #3. In this step, the AEP estimate increases by 1.18 MWh, which is almost 3%. Such a substantial difference could suggest that low-level jets make a significant contribution to the AEP. However, studying the fluctuations with respect to the brute-force result in Figure 8.12a shows that such a difference has the same order magnitude as the estimation error and, thus, could be attributed to imprecision.

For the offshore location, including all seven wind profile shapes shown in Figure 8.7b yields an average AEP of 58.4 MWh over the years 2008–2017. The variation in AEP between the years is similar to that for the onshore location. Comparing the result for 2008 with all wind profile shapes against the results of the brute-force calculations shows an AEP underestimation of 2%, which is relatively large compared to the onshore location. The result based on the neutral logarithmic wind profile shows a large AEP underestimation of more than 10%.

Three jet-like wind profile shapes resulted from the clustering for the offshore location. The effect of these jet-like shapes is evaluated based on the difference in AEP due to the introduction of cluster 4–6 (between Set #1 and All). In this step, the AEP estimate increases with 5.3 MWh. This 10% increase suggests that considering low-level jets is more important for the evaluated offshore location than for the onshore location.

8.7. Conclusion

The wind profile variability has a substantial effect on annual energy production (AEP). This effect is evaluated using an efficient AEP estimation methodology based on multiple power curves for different wind profile shapes. The power curves are generated with the use of a basic pumping flight operation model and a set of prevalent wind profile shapes identified using clustering. The AEP is calculated with these power curves and a statistical summary of the wind profile climate.

The quasi-steady model is modified to make it more suitable for optimisation. A fixed number of points is used to represent the flight trajectory, and a constraint is added to ensure the periodicity of the pumping cycle. Numerical optimisation is used to find the optimal operation approach of the system at given wind conditions. The operation approach controls the pumping cycle flight trajectory and is set by the duration of the phases, the reel-out elevation angle, the minimum tether length, and the tether forces at every point along the path.

The basic pumping flight model is used to derive a power curve using a neutral logarithmic wind profile. The investigated system, similar to that of Kitepower, reaches the rated mean cycle power of 16.7 kW at $v_{\mathrm{w,200m}}$=16.6 m/s. Slightly below this wind speed, the reel-out power reaches the 40 kW maximum limit imposed by the generator capacity. Consequently, the kite is depowered by increasing the elevation angle more rapidly at higher wind speeds. In this region, the power curve does not exhibit a plateau but shows a modest decline, unlike the power curve of modern utility-scale tower-based wind turbines. Once the maximum elevation angle is reached, the power curve rapidly declines before cut-out.

Instead of deriving a power curve for a fixed wind profile shape, a more general power curve is inferred based on the full wind atlas data. This brute-force approach obtains the optimal power output for every hourly wind profile. The power curve is given by the mean power output as a function of the wind speed. The resulting power curve shows some differences with respect to the power curve for a logarithmic wind profile, suggesting a mismatch between the mean shape of the wind atlas profiles and the neutral logarithmic wind profile shape. Expressing the power curve against the wind speed at 200 m height yields a modest variation around the curve. This suggests that for the given system, $v_{\mathrm{w,200m}}$ is a good predictor for the power output.

The wind profiles in the wind atlas data are manually clustered according to shape. The small set of mean-cluster wind profile shapes is employed to describe the wind profile climate in the AEP estimation. The set includes logarithmic wind profiles for stability conditions ranging from stable to unstable conditions. Moreover, one of the identified shapes shows a pronounced low-level jet. The wind profile probability distribution is discretised by wind profile shape, yielding a univariate distribution for each shape.

For each of the wind profile shapes, a power curve is computed. For most wind speeds, the power curves only partly cover the range of power outputs determined with the brute-force approach. Summing the integrals of the products of each power curve with its respective probability distribution yields the average power output. Including the jet-like wind profile shapes changes the estimate by 3% and 10% for the onshore and offshore locations, respectively. This suggests that considering low-level jets is important for estimating the AEP.

For the onshore location of the met mast Cabauw, only considering the logarithmic wind profile yields an AEP overestimation of 5%. Using seven wind profile shapes reduces the error to a 0.5% underestimation. For the offshore location of the met mast IJmuiden, only considering the logarithmic wind profile gives an AEP underestimation of more than 10%, which is reduced to 2% when using seven wind profile shapes. The associated convergence trends show that including more wind profile shapes in the calculation increases the precision of the AEP estimate. Thereby, the number of wind profile shapes to consider in the calculation can be chosen such that accuracy requirements are met. How much accuracy is required depends on the application of the AEP estimation.

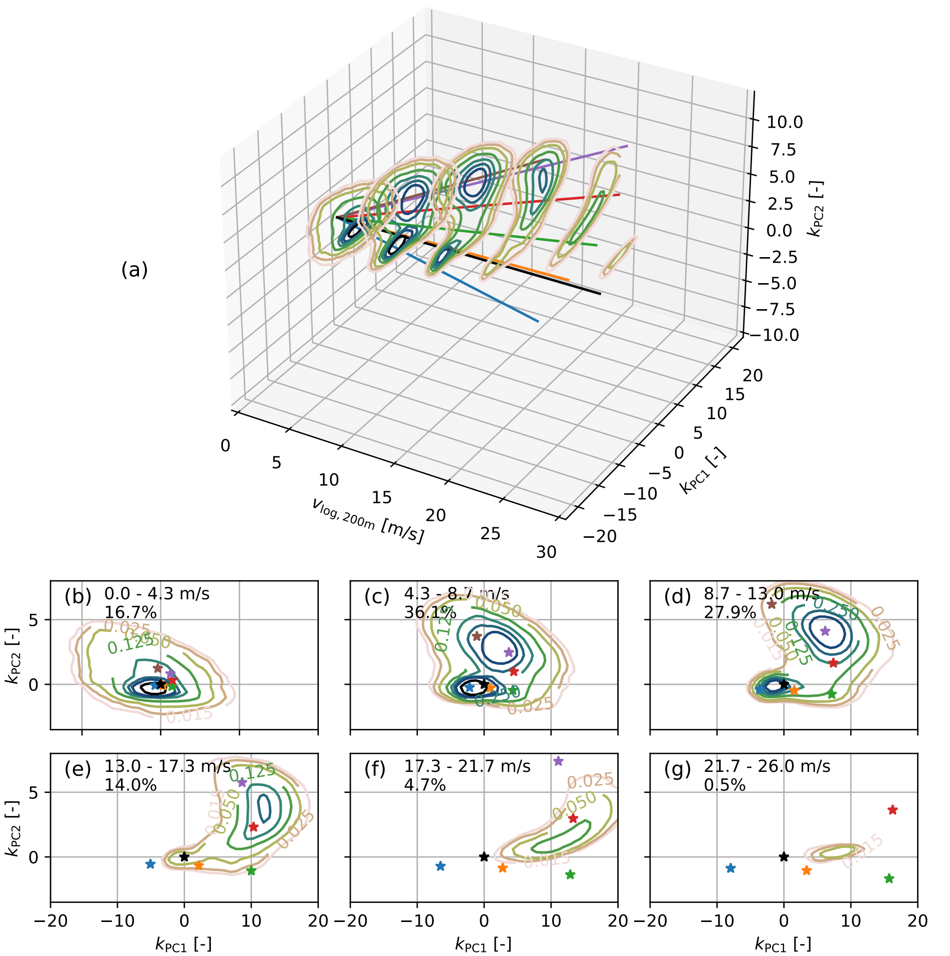

Figure 8.8: (**a**) Full wind profile distribution of the onshore dataset depicting the relationship between shape and magnitude. Each wind profile is decomposed based on the neutral logarithmic wind profile and principal components with the fitted weights $v_{\mathrm{log,200m}}$, k_{PC1}, and k_{PC2}. Wind profiles with equal shapes but different magnitudes lie on the same line passing through the origin. The lines corresponding to the cluster-mean shapes (see Figure 8.6a) are depicted with coloured lines. (**b**-**g**) The 3d density distribution is broken down into six slices along the x-axis.

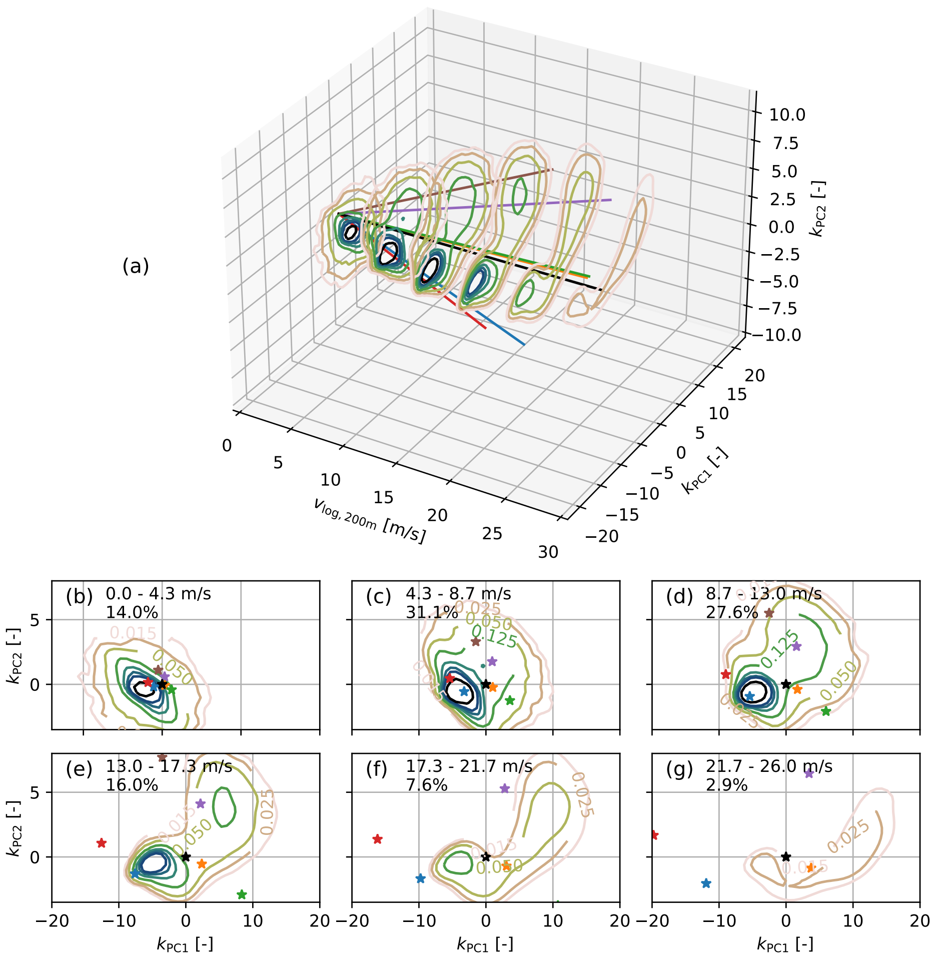

Figure 8.9: (**a**) Full wind profile distribution of the offshore dataset depicting the relationship between shape and magnitude. Each wind profile is decomposed based on the neutral logarithmic wind profile and principal components with the fitted weights $v_{\mathrm{log,200m}}$, k_{PC1}, and k_{PC2}. Wind profiles with equal shapes but different magnitudes lie on the same line passing through the origin. The lines corresponding to the cluster-mean shapes (see Figure 8.6b) are depicted with coloured lines. (**b**-**g**) The 3d density distribution is broken down into six slices along the x-axis.

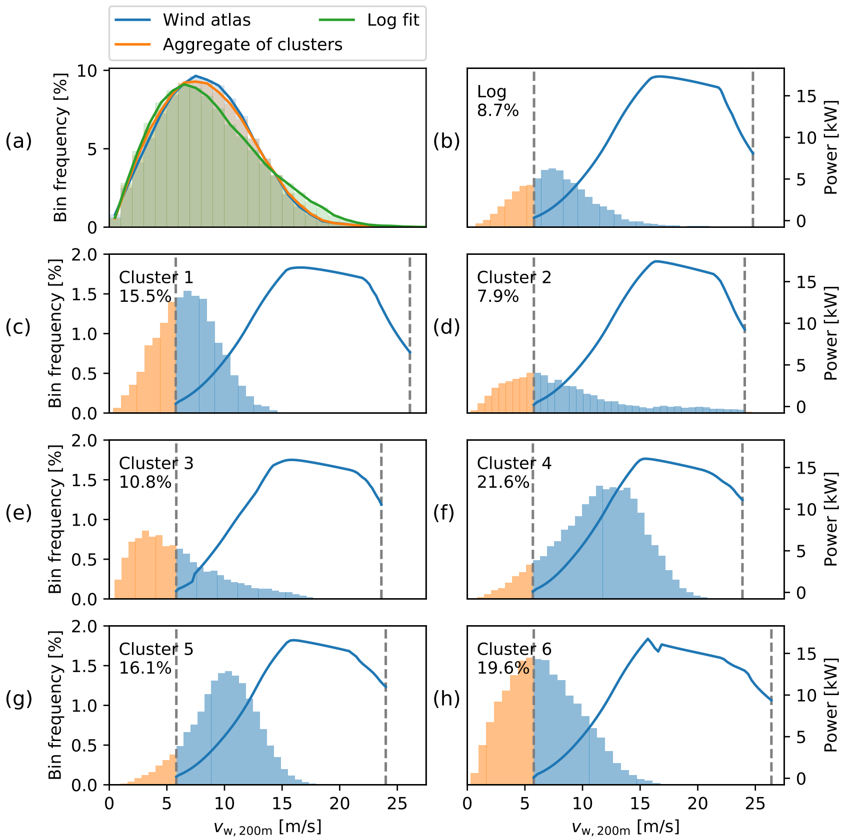

Figure 8.10: Wind profile magnitude distributions of all clusters combined (**a**) and each separate cluster (**b**-**h**) in the proposed wind climate description for the onshore location. The aggregated distribution is compared against the distributions in the wind atlas data and in the benchmark wind climate description. Also depicted are the power curves of the clusters from which the cut-in and cut-out wind speeds are inferred.

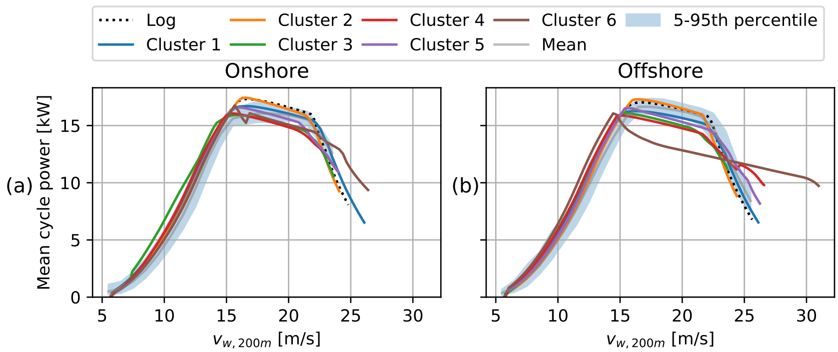

Figure 8.11: The power curves for the cluster-mean and neutral logarithmic wind profile shapes compared against the brute-force optimisation results.

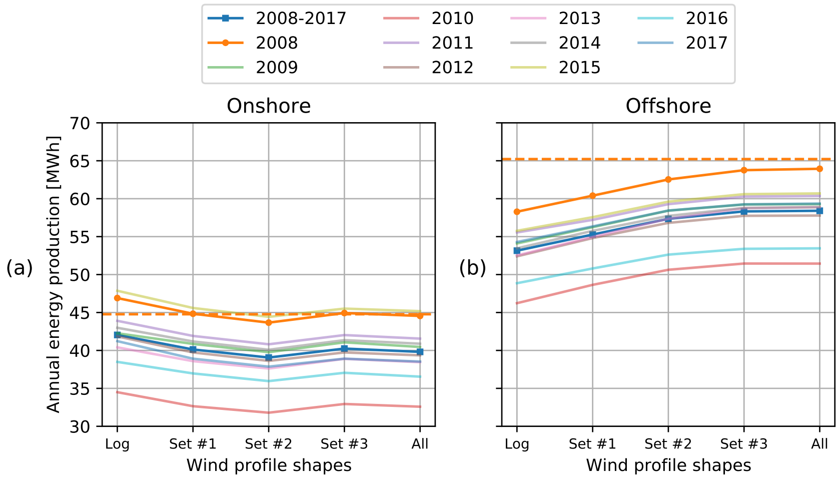

Figure 8.12: The AEP estimated with the use of increasing numbers of mean-cluster wind profile shapes (see Table 8.1). For the onshore location, jet-like wind profiles are only introduced in the last step, while, for the offshore location The dashed line shows the average AEP of 2008 calculated with the brute-force optimisation results.

9

Conclusion

This last chapter starts directly by answering the main research question, followed by answering the underlying research questions. Hereafter follows the discussion of the findings of this thesis, including the recommendations for future work and the practical implications of this dissertation.

Main conclusions

How significant is the wind profile variability to the annual energy production estimation? (main research question)

Not considering the wind profile variability in calculating the annual energy production (AEP) suggests the potential for a significant error exceeding 10%, based on AEP estimations performed for an on- and offshore location in the Netherlands. The error is quantified for a 17 kW airborne wind energy (AWE) system by benchmarking AEP estimations with a wind resource representation assuming a neutral logarithmic wind profile against AEP calculations based on hourly wind profiles of wind atlas data. The latter computational estimate lies closer to the 'ground truth' as it does not require statistically summarising the wind climate. For the offshore location, the AEP benchmark is underestimated by more than 10%, while for the onshore location the benchmark is overestimated by 5%.

The 10% AEP estimation error for the offshore location is unacceptable for most applications and underlines the demand for a refined wind profile climate representation. On the other hand, the benchmark AEP calculation is very refined but may not meet requirements on computational cost. This thesis proposes using a wind resource representation based on clustered wind profile shapes, which allows trading off the accuracy of the AEP estimation against the computational cost by tailoring the level of detail of the wind resource to the application of the AEP estimation.

Employing seven prevalent wind profile shapes to represent the wind climate reduces the AEP estimation error to 0.5% and 2% for the onshore and offshore locations, respectively. To quantify the effect of the wind profile variability on the AEP estimation, a sequence of annual energy production estimations is conducted in which progressively more wind profile shapes are considered. These mean-cluster wind profile shapes are obtained by clustering wind atlas data of wind profiles and represent the prevalent wind profiles. Whether or not including the jet-like mean-cluster wind profile shapes changes the estimate by 3% and 10% for the onshore and offshore locations, respectively. Although, for the onshore location, the difference has the same order of magnitude as the error relative to the AEP benchmark, the difference found for the offshore location is substantial and cannot be attributed to the imprecision of the estimation alone. As such, these results suggest that considering low-level jets is important for estimating the AEP.

The conducted AEP estimation has not been validated, as long-term operational data from AWE systems in continuous operation are currently unavailable. The AEP estimations assume that optimal pumping cycles are flown by the evaluated AWE system, which is not the case in currently available operational data. The AEP estimation assumes that the AWE system continuously adapts its flight operation to the instantaneous wind profile. However, in practice, often only limited information about the wind profile is available, impeding real-time optimisation of the flight operation.

How does flexible-height operation change the accessible wind resource? (Chapter 3)

The variable-height wind speed distributions of locations across Europe show improved characteristics compared to the wind speed distribution at the hub height of tower-based turbines. These distributions are compiled from the instantaneous maximum wind speeds within the operational height range using reanalysis wind data. Conventionally, wind speed distributions characterise the wind climate at a fixed altitude for a given site. Such a fixed-height distribution is commonly used with tower-based wind turbines to assess the energy production. By continuously adjusting the harvesting height, AWE systems can tailor the effective wind speed distribution experienced by the kite to increase energy production. Therefore, the energy production of AWE systems may be more closely related to the variable-height wind speed distribution.

Variable-height harvesting of wind energy increases the availability of wind speeds exceeding either the approximate cut-in or rated conditions for most regions over land, though exceptions may occur. The availability indicates the percentage of time for which a wind power density is exceeded. To provide a tangible context, the selected threshold values for evaluating availability align with typical cut-in and rated wind speeds of conventional wind turbines. At coastal areas of the North Sea, the potential to reduce the percentage of time below the approximate cut-in condition is relatively small but still significant compared to other regions, while the percentage of time above the approximate rated condition is relatively large. Mediterranean coastal areas show an opposite trend. Over the sea, the increase in availability is substantially smaller as low altitudes already exhibit good wind conditions.

How can the wind profile variability be integrated into a statistical description of the wind climate? (Chapter 4)

A compelling approach involves using a data-driven methodology to distil the full spectrum of wind profiles into a concise set of prevailing wind profile shapes sufficient to capture the wind profile variability at a given deployment site. This approach does not confine the considered wind profile shapes to forms described by well-known relationships such as Monin-Obukhov Similarity, which are only valid in the surface layer. The methodology adopted in this thesis heavily relies on decomposing the wind profile into its shape and magnitude components. Clustering techniques are used to identify the dominant wind profile shapes from wind atlas data. Additionally, the distribution of the wind profile magnitudes is inferred and expressed as a function of the reference wind speed at 200 m height. The set of wind profile shapes and corresponding frequency distributions together form a compact wind climate representation.

Although the choice for the number of clusters highly depends on the application, the coherent structures observed in the wind profile data of the investigated locations indicate that using a minimum of six clusters could provide an adequate wind climate representation. The identified clusters will highly depend on the chosen clustering algorithm and the parameters controlling the learning process, such as the number of requested clusters. It is crucial to maintain a critical perspective on the resulting clusters and assess whether they are associated with coherent structures in the wind profile data. Specifically, an incoherent cluster could potentially lead to a mean-cluster shape that does not accurately reflect a realistic wind profile. The share of outliers and differences in cluster size are determining factors for selecting the clustering algorithm.

The adopted approach is not limited to identifying wind profile shape clusters for a single site but has the capability to generate a unified set of wind profile shapes valid for a larger area. This is demonstrated by generating eight clusters for the whole of the Netherlands. A clear distinction is observed between clusters linked to on- and offshore wind conditions. As expected, the sharply defined patterns in the spatial variability of the frequency of each cluster coincide with orographic features.

A wind climate description based on clustered wind profile shapes can adequately summarise the wind conditions for an average year when sufficient years of wind atlas data are considered. This claim is supported by the agreement between the clusters and local weather patterns regarding their temporal occurrence and relationship to atmospheric stability. A unique aspect of the proposed wind climate representation is that it incorporates jet-like wind profiles, which cannot be achieved using conventional wind profile relationships. However, it is inevitable that certain characteristic properties of the wind profiles, such as the height of the atmospheric boundary layer, may not be fully preserved.

How can the pumping flight operation be modelled efficiently with due regard to the influence of the wind profile variability? (Chapter 5–7)

The trade-off between using a quasi-steady or dynamic framework to resolve the pumping cycle operation mostly concerns simplicity and accuracy. Quasi-steady frameworks link independent states by solving the motion along an idealised flight trajectory for which the motion is prescribed to some extent. The simplicity of the input and their robustness make them a powerful tool for the preliminary calculation of the performance of AWE systems. On the other hand, dynamic frameworks require carefully compiled control input to obtain a realistic pumping cycle. This task could also be achieved by employing the model in an optimal control problem. Unlike quasi-steady frameworks, dynamic frameworks account for transient effects, and as a result, the dominant factor determining the accuracy is not the solver but rather the model description of the system. The exact magnitude of the difference in accuracy between the two frameworks has not been determined.

Within these frameworks, a rigidly-linked, two-point model of the kite is an efficient configuration to model a flexible kite with suspended control unit. Such a model realistically resolves the characteristic swinging motion of the kite during turns and the increased pitch during depowered flight, as demonstrated with a revised steady-rotation-state and dynamic model. This important aspect of the turning mechanism cannot be modelled with single-point kite models. Moreover, the two-point kite model complements the aerodynamic model as it allows computing the angle of attack of the wing by resolving the pitch. These characteristics improve the generalisation of the kite model with little additional computational effort. Moreover, the two-point model of the kite may conveniently extend a lumped-mass tether model. Alternatively, rigid kites are typically modelled as a rigid body. However, this is excessive for flexible kites, as their yawing motion is frequently described using a kinematic relationship in performance models.

The model development for flexible-kite AWE systems is hindered by the increased complexity of the state estimation and system identification compared to rigid-kite systems. A significant factor is the continuous deformation of the kite during flight, making it challenging to determine the quantities required for the system identification. Together with the initially overlooked off-centre flow measurement, this has hindered the aerodynamic model identification carried out in this thesis. Since no clear relationships with the angle of attack or steering input could be identified, a model structure with solely a dependency on the depower signal appears to be the most sensible option. Concerning the state estimation, the insufficient quality of the flight data impeded confidently reconstructing the flight trajectory, as required for studying the swinging motion of the kite.

Due to the inability to accurately estimate the instantaneous wind profile, the available experimental data proves inadequate for validating the quasi-steady model (QSM). While the power curve of the figure-of-eight manoeuvre as a function of the apparent wind speed computed with the QSM exhibits reasonable agreement with the flight data, the agreement significantly diminishes when the power curve and flight data are expressed against the reconstructed wind speed. This discrepancy is explained by a faulty conversion of the measured apparent wind speed to the reconstructed wind speed, resulting from deficiencies in the wind speed reconstruction. The mismatch of the calculated and observed power curve of the figures of eight also persists in the mean cycle power curve. The absence of an accurate wind profile estimate prohibits validating the QSM and its modules, as well as assessing the effect of incorporating a newly developed realistic figure-of-eight parameterisation. Direct wind profile measurements would be highly beneficial, if not mandatory, to facilitate validation.

Applying the QSM in performance optimisations proves effective in exploring the operational decision space to find the pumping cycle that maximises power output. To increase efficiency, the QSM is integrated with the optimiser to take over finding the kinematic ratio of the underlying steady states and ensure system limits are adhered to.

Caution is required when interpreting the optimal power output because the optimisation may venture outside the established regions of the operational decision space, particularly with the QSM, as it lacks validation.

How sensitive is the power output to the wind profile? (Chapter 8)

For each of the identified mean-cluster wind profile shapes, one power curve is generated for a range of wind profiles with increasing magnitude expressed by the wind speed at 200 m height. The curves are obtained by optimising the power output of a 17 kW AWE system. The resulting mean-cluster power curves for the on- and offshore reference locations are similar in shape but show slightly different mean cycle power maxima and relationships with the wind speed. Below the rated wind speed, the power curves exhibit a similar slope but are vertically shifted by a maximum of 1.2 kW. Above the rated wind speed, larger variations between the curves are observed, especially towards cut-out.

Alternatively, the power production of the AWE system is characterised by a 5^{th}–95^{th} percentile power band and the mean power curve as a function of the wind speed at 200 m height. These power output statistics are obtained by optimising the pumping cycle operation for every hourly wind profile in one year of wind atlas data of a given location. The results of the on- and offshore reference locations show that deviations due to wind profile variability of $\pm5\%$ relative to the mean power curve are frequent. The power statistics exhibit a relatively low deviation when expressed against the wind speed at 200 m, which can be attributed to the operation of the kite being centred around at this height during the reel-out phase.

Comparing the two power characterisations shows that the mean-cluster power curves indicate a higher power output in the region below rated wind speed, where the investigated system most frequently operates. Despite the statistical power curve being derived from a year with relatively high winds, the mean-cluster power curves indicate optimistic power characteristics. The mismatch between the two characterisations may be explained by the loss of detailed wind profile features due to the clustering, which otherwise would penalise the power output.

An important aspect of the mean-cluster power curves is their applicability in estimating energy production for sites other than those used to identify the wind profile clusters. A requirement for the alternative site is that it experiences similar variations in wind profile shape, albeit with different probabilities. On the contrary, the statistically obtained mean power curve is site-dependent as it conceals the site-dependent wind profile shape probability.

Discussion

A novel framework is developed to estimate energy production by directly parameterising power output as a function of wind profile shape and magnitude. Similar to the approach proposed by Ranneberg et al. [47], multiple power curves are employed for this purpose. However, the statistical approach used in this thesis considers the interdependency of instantaneous wind speeds at different heights. By characterising the annual probability of the wind profile in terms of its shape function and magnitude component, the wind profile probability can be coupled to the proposed power characterisation. Due to yearly reoccurring patterns observed in the probability of wind conditions, this characterisation may accurately represent the wind statistics of an average year and enables an efficient AEP estimation. On the contrary, the approach of Ranneberg et al. does not derive wind climate characteristics but depends more directly on long-term wind profile data.

The findings regarding the AEP error resulting from assuming a neutral logarithmic wind profile are consistent with those of Ranneberg et al. [47]. Their study reports a 10% AEP error when comparing a simple AEP estimation for a 100 kW pumping system based on a neutral logarithmic wind profile (613 MWh) against an estimation based on forecasted wind data (686 MWh) for an onshore location in Germany. This AEP error is similar to the error found in the present thesis for the offshore reference location. Unlike this thesis, Ranneberg et al. do not investigate the cause of the reported error.

Limitations and recommendations

This study is unable to provide a thorough validation of the AEP estimation framework, the underlying power curves, and the employed pumping flight operation model. Consequently, it is difficult to draw definitive conclusions regarding the practical implications of this thesis. The validation is hindered by the unavailability of long-term operational data from AWE systems, the absence of direct wind profile measurements, and inadequate measurements of the state of the kite. Moreover, system identification is hindered by inadequate airflow measurements relative to the kite. Conducting experiments, including system identification, should be integral to developing the pumping flight operation model. Therefore, an essential step towards model refinement is improving the collaboration between academia and AWE companies in designing experiments. Concrete recommendations for improving measurements include measuring the wind profile directly with a lidar system, equipping the suspended control unit with motion sensors, and positioning the flow sensors, including a wind vane to measure side slip, in the symmetry plane of the kite. The improved measurements may contribute to enhancing state estimation and wind profile reconstruction.

Another limitation of this study pertains to the investigated locations. The results are primarily based on analysing two locations in the Netherlands: the onshore location of the met mast Cabauw and the offshore location of the met mast IJmuiden. These locations are expected to have representative wind climates for the Netherlands. However, more locations need to be analysed, covering a larger area and having a larger variety of wind climates, to enable generalisation of the results, for example, to provide a generic conclusion on how many clusters are needed to perform sufficiently accurate AEP estimations.

This work does not address local effects on the wind profile and turbulence. A specific but relevant example is the impact of wakes from upstream systems when operating in a wind farm configuration. The mean-cluster wind profile shapes, deduced from mesoscale wind data, do not accommodate incorporating local effects. To increase the accuracy of the AEP estimation, future research is needed to investigate how local effects can be imposed on the wind profile after the clustering and how to incorporate turbulence in performance models. The work of Haas et al. [125] is a good reference for this purpose.

A strong focus is placed on a single type and size of flexible-kite pumping AWE system similar to Kitepower's system, operated in a similar way. The availability of experimental data of a system with a small kite with a rated power of approximately 17 kW motivated the choice for the investigated system. Moreover, the investigated operational approach for the reel-out phase is limited to flying figure-of-eight manoeuvres along an elevated axis. Consequently, the shapes of the derived power curves and specifics of the AEP estimation may not be generalised to different types and modern AWE systems. Specifically, the power output of systems employing a more horizontally aligned reel-out phase may be significantly less sensitive to the wind profile shape.

Although this thesis touches upon the trade-off between using a quasi-steady or dynamic framework to resolve the pumping cycle, it does not compare these frameworks quantitatively. This could be achieved by implementing a dynamic model for a flexible-kite pumping AWE system and applying it in an optimal control problem to perform performance optimisations without needing to address the specifics of the control system. It is highly recommended to compare the performance of a cross-wind manoeuvre before zooming out to the full pumping cycle. This comparison could contribute to identifying the limits of the quasi-steady framework and shed light on its suitability for performance optimisation.

Practical implications

A substantial computational cost reduction is achieved with the cluster-based AEP estimation relative to the brute-force approach. The cluster-based AEP estimation in this thesis requires approximately 350 performance optimisations: 50 performance optimisations to construct each of the seven power curves. This is significantly less than a brute-force approach, which requires one optimisation for each hourly wind profile in wind atlas data, summing up to 8760 optimisations for each year. Particularly considering that time windows of 30 years are typically analysed to evaluate the wind climate. An often-used workaround, also used in this thesis, is reducing the amount of brute-force calculations by narrowing the time window. However, the resulting AEP estimate may not be valid for an average year.

Characterising the power production of AWE systems as a function of wind profile shape and magnitude offers a new perspective to the ongoing debate on how to characterise system performance uniformly [13]. In contrast to expressing the power curve against the average pattern trajectory height [55], the proposed characterisation employs an unambiguous set of parameters and enables a direct coupling to the wind profile probability. Publishing sets of wind profile shapes characteristic of specific areas may promote the standardisation of wind conditions for which AWE systems are rated in terms of power production. The wind profile shapes identified in this work for the onshore and offshore reference locations in the Netherlands are published together with the corresponding magnitude probability distributions. Other researchers may contribute by using these to quantify the effect of the wind profile variability on different types and sizes of AWE systems or by employing different flight operation models.

This thesis underlines that wind resource assessments for AWE systems need to move away from using the methodologies borrowed from conventional wind energy. These methodologies, developed for tower-based wind turbines, neglect the wind profile variability and assume a neutral logarithmic wind profile. Although this may be acceptable as a first estimate of the AEP of an AWE system, it is not sufficiently accurate for most wind resource assessments. The separated wind profile climate and power production characterisation of the presented AEP estimation framework offer a practical solution to efficiently account for the wind profile variability. Adopting this framework requires characterising the performance of the investigated system in the form of a set of power curves corresponding to a set of prevalent wind profile shapes. This characterisation facilitates incorporating the advantages of the operational flexibility of AWE systems.

The significant influence of wind profile variability on AEP underlines the potential of real-time, informed planning of the operations of AWE systems, requiring knowledge about the instantaneous wind profile. The AEP contribution attributed to accounting for jet-like wind profiles suggests that a substantial gain in AEP can be achieved by tailoring the flight operation to the specifics of the wind profile. A prerequisite for this optimisation is the measurement or estimation of the instantaneous wind profile in the vicinity of the operating AWE system. Employing the kite as a flying probe to estimate the wind profile is an economical option, while wind profile measurements with lidar entail high costs. Additionally, real-time planning necessitates online optimisation with efficient flight operation models to determine the optimal operational approach.

This thesis makes a small but important contribution to reducing the uncertainty of long-term performance estimation for AWE systems. The presented AEP estimation framework enables addressing the operational benefits of AWE systems compared to tower-based wind turbines and may be used to further establish the potential of AWE technology. Nonetheless, further research is needed to evaluate the cost-competitiveness and complementarity of AWE with respect to tower-based wind turbines and improve the understanding of the viability of large-scale deployment of AWE systems. Assuming positive outcomes, these studies will help to secure investments, accelerate technology development, and ultimately create a role for AWE in the future energy mix. As such, this dissertation is titled ‘Power to the Airborne Wind Energy Performance Model’.

A

Geometric bridle model

A simple geometrical model of the bridle system is used to approximate the depower angle α_d as a function of the power setting based on the cosine rule. As illustrated in Figure 2.13, this angle varies depending on the power setting of the kite:

$$\sin\alpha_\mathrm{d} = \frac{(l_0 + \Delta l)^2 - d^2 - c_\mathrm{ref}^2}{2\,d\,c_\mathrm{ref}} \quad , \tag{A.1}$$

where $d = 11.4\,\mathrm{m}$ is the length of the front bridle between the bridle point and the chord line, $c_\mathrm{ref} = 1.8\,\mathrm{m}$ is the length along the chord line between the front and rear bride line connections, l_0 is the length along the depower tape between the bridle point and the trailing edge for $\alpha_\mathrm{d} = 0°$ at the power setting $u_\mathrm{p,ref} = 0.82$, and Δl is the difference in this length due to a power setting u_p other than the reference value. The latter length difference is given by:

$$\Delta l = \frac{\Delta l_\mathrm{d}}{2} = -\frac{u_\mathrm{p} - u_\mathrm{p,ref}}{2} l_\mathrm{d} \quad , \tag{A.2}$$

in which $l_\mathrm{d} = 5\,\mathrm{m}$ is the depower tape length difference between the fully powered and de-powered setting: $u_\mathrm{p} = 1$ and $u_\mathrm{p} = 0$, respectively. Note that Δl is half the de-power tape length difference due to the pulley connection between the de-power tape and the rear bridles.

The power settings during the reel-out and reel-in phases are 0.78 and 0.7, respectively. This corresponds to a depower angle of 3.2° and 9.8°.

B

Flight trajectory reconstruction

The kinematics of the wing recorded in the flight data show inconsistencies in the measured tether reel-out speed and are reconstructed in a preprocessing step to remove anomalies. The dynamic simulation relies on the recorded wing kinematics and tether reel-out speed for its input. Directly using these recorded quantities as input leads to faulty simulations, and a workaround is needed to obtain coherent input. The reconstruction is carried out for the full 65$^{\text{th}}$ pumping cycle.

A preliminary evaluation of the wing kinematics in the flight data shows that the vertical speed does not fully agree with the derivative of the vertical position of the wing, even though it does for the horizontal components. The largest mismatch occurs during the turns, where the recorded vertical speed is more negative than the derivative of the vertical position. The recorded vertical position is GPS data enhanced with barometer measurements. However, it is expected that the vertical speed was not updated accordingly.

The inconsistent vertical speed leads to a discrepancy between the derivative of the measured radial position $\dot{\hat{r}}_{\mathrm{k}}$ and the measured radial component of the wing velocity $\hat{v}_{\mathrm{k,r}}$, while in theory, they should be the same. These quantities are depicted with the blue and red lines, respectively, in Fig. B.1c. The radial component of the wing velocity is calculated with:

$$v_{\mathrm{k,r}} = \frac{\mathbf{r}_{\mathrm{k}} \cdot \mathbf{v}_{\mathrm{k}}}{\|\mathbf{r}_{\mathrm{k}}\|}, \tag{B.1}$$

in which $\mathbf{r}_{\mathrm{k}}$ and $\mathbf{v}_{\mathrm{k}}$ are the position and velocity of the wing, respectively. An objective of the intended flight trajectory reconstruction is to ensure that the updated radial component of the wing velocity and the derivative of the radial position agree.

As an additional check, the derivative of the measured radial position of the wing $\dot{\hat{r}}_{\mathrm{k}}$ is compared to the measured tether reel-out speed $\dot{\hat{l}}_{\mathrm{t}}$ (dotted black line in Fig. B.1c). The derivative of the radial position shows large fluctuations around the tether reel-out speed in the reel-out phase. The magnitude of the fluctuations conflicts with our expectation that the changes in tether slack (difference between the tether length and radial position of the kite) and stretch are small in this phase. Towards the end of the right turns (at the end of the blue intervals), the derivative of the radial position even tends to become shortly negative.

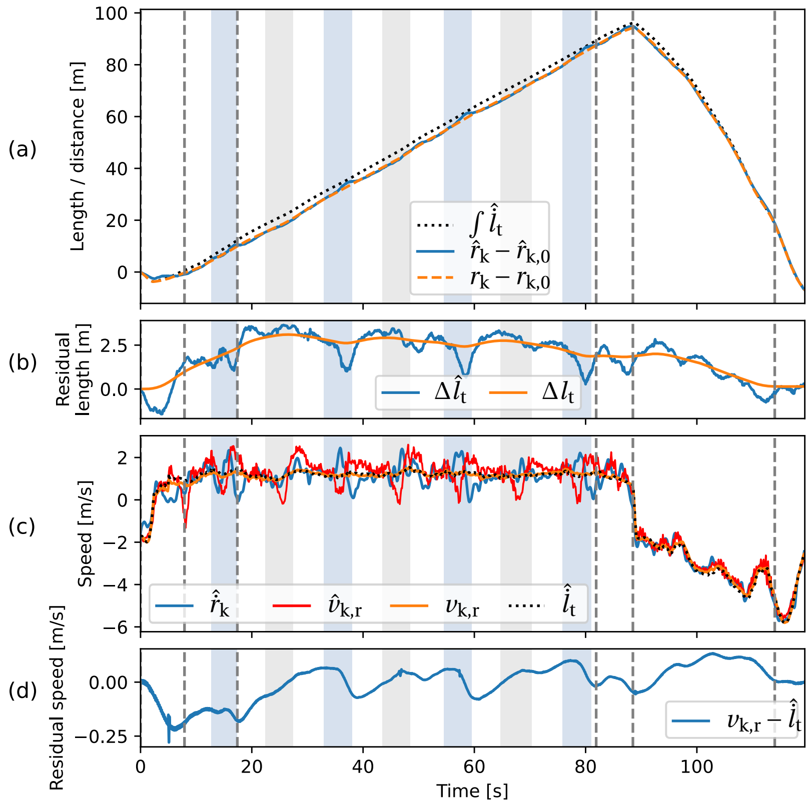

Figure B.1: (**a**) Evolution of unstrained tether length $\int \hat{\dot{l}}_t$ and the measured and reconstructed radial distances of the wing, $\hat{r}_k$ and r_k, all with their initial values subtracted. (**b**) Difference between the tether length and the measured radial distance of the wing $\Delta\hat{l}_t$ and its equivalent after the reconstruction Δl_t. (**c**) Time-derivative of measured radial position of the wing $\hat{\dot{r}}_k$, measured and reconstructed radial speeds of the wing, $\hat{v}_{k,r}$ and $v_{k,r}$, and measured tether reel-out speed $\hat{\dot{l}}_t$. (**d**) Residual between the tether reel-out speed and reconstructed radial speed. The intervals shaded grey and blue indicate left and right turns, respectively, from a downwind perspective.

Figure B.1a shows how the integrated measured reel-out speed (dotted black line) evolves with respect to the measured radial position of the wing $\hat{r}_k$ (blue line). During the right turns, the inferred tether length increases approximately linearly, while the radial position exhibits subtle local maxima. These local maxima coincide with the large discrepancies between the derivative of the radial position and the tether reel-out speed observed in Fig. B.1c. Note that the tether length lines depict the relative lengths with respect to the start of the pumping cycle. The lines need to be shifted up with their initial values to obtain their respective absolute values. Unfortunately, the absolute tether length is unknown as it is not measured directly.

The residual between the inferred tether length and measured radial position $\Delta\hat{l}_t$ is shown in Fig. B.1b. During the right turns, the residual changes roughly 2 m (depth of the valley) within a couple of seconds. The corresponding relatively large increase in radial position can partly be attributed to decreased tether slack and increased tether stretch. However, the magnitude of the change is deemed to be too large to be attributed only to changes in these quantities. Note that also here, the line may shift vertically depending on the initial values. As such, no conclusions can be drawn based on the magnitude of the residual but merely on how it changes with time. The given residual length has an unknown offset with respect to the tether slack. Note that the tether slack cannot be negative.

The maxima in the recorded radial position do not need to be purely physical. Another possible cause is GPS inaccuracy during manoeuvres, which has previously been reported in the literature. Borobia et al. [126] reported measured radial position exceeding varying more than 3 m while none was expected. Considering the imprecision of the recorded position, the wing kinematics is adapted by letting the radial wing speed follow the measured reel-out speed as closely as possible.

The flight trajectory reconstruction is obtained using a discrete-time optimisation problem that minimises the error between the modelled radial wing speed and recorded tether reel-out speed while limiting the bias between the modelled and recorded wing position

$$\begin{aligned} \min_{\mathbf{r}_\mathrm{k}(\cdot),\mathbf{v}_\mathrm{k}(\cdot),\mathbf{a}_\mathrm{k}(\cdot)} \quad & \sum_{i=0}^{\mathrm{N}} \left[\mathrm{w}\left(v_{\mathrm{k,r}} - \hat{\dot{l}}_\mathrm{t} \right)^2 + \left(\mathbf{r}_\mathrm{k} - \hat{\mathbf{r}}_\mathrm{k}\right)^\top \left(\mathbf{r}_\mathrm{k} - \hat{\mathbf{r}}_\mathrm{k}\right) \right]_{t=\frac{i}{10}} \\ \text{s.t.} \quad & \mathbf{a}_\mathrm{k} = \dot{\mathbf{v}}_\mathrm{k} = \ddot{\mathbf{r}}_\mathrm{k}\,. \end{aligned} \tag{B.2}$$

Quantities marked with a hat indicate measured quantities, whereas the absence of a hat indicates modelled quantities. A discrete function is used for the acceleration of the wing, and continuous trajectories are used for the velocity and position of the wing. The decision variables consist of the wing accelerations during the control intervals $\mathbf{a}_\mathrm{k}(\cdot)$ and the velocities $\mathbf{v}_\mathrm{k}(\cdot)$ and positions $\mathbf{r}_\mathrm{k}(\cdot)$ at the control interval boundaries. N is the number of time steps, and the weighing factor w = 25 is chosen as it leads to a good balance between the two objectives. Note that having matching reel-out and radial wing speeds does not necessarily mean that also the tether length is the same as the radial position. However, it does mean that the tether slack stays constant.

In line with the dynamic simulation, the fitting problem uses discrete control input trajectories. It assumes a constant acceleration within each simulation time step of 0.1 s. Between the corresponding control intervals, the values may vary. Due to the step function form of the acceleration, the velocity and position are linear and quadratic functions, respectively, within the control intervals. These low-order forms allow for sufficient detail due to the small time step. The fitting problem is solved in CasADi using a multiple-shooting approach. This approach is not hindered by integration drift causing an accumulating error with time.

The flight trajectory reconstruction results are shown with the orange lines in Fig. B.1. The reconstruction shaves off the local maxima in the recorded radial position, as can be observed in Fig. B.1a. Figure B.1c shows that the reconstructed radial wing speed follows the measured reel-out speed more closely. The residual speed, which is penalised by the first term of the objective function, is illustrated in Fig. B.1d. The optimiser reduces the position bias, which is penalised by the second term of the objective function, by allowing small changes to the radial wing speed with respect to the measured reel-out speed. As a consequence, the reconstruction does not lower the residual length substantially but keeps it close to the original residual length, as can be seen in Fig. B.1b.

The reconstructed radial wing acceleration $a_{k,r}$ is used as tether reel-out acceleration input $\ddot{l}_t$ for the simulation. Thus, not only the flight trajectory is reconstructed but also the tether reel-out speed is modified with respect to the measurements. As a result, the tether slack remains constant in the simulation and is set by the choice for the initial tether length. In reality, changes in slack length will occur, especially during the transition phases. Therefore, this approach might be sub-optimal for simulating the entire pumping cycle. Nonetheless, it is suitable for simulating intervals where only small tether slack and stretch changes are expected, such as the reel-out phase.

It is acknowledged that the flight trajectory reconstruction might not be strictly valid. However, it serves the main objective of this study by enabling the simulation of a short interval that encompasses a figure-of-eight manoeuvre during reel-out. A more educated reconstruction would require a lot more resources and probably more testing and is recommended as a possible future improvement.

C

Pitch and roll angle definitions

Expressing the attitude of the kite using pitch and roll angles with respect to the wind reference frame gives large variations of these angles along the flight trajectory. Consequently, the kite attitude is difficult to interpret from these angles. Variations are smaller when the pitch and roll angles are expressed with respect to the tangential plane, which is perpendicular to the position vector of the kite and shown with the black rectangle in Fig. C.1. The variations are smaller since the up-direction (positive z-axis) of the kite and the direction of the position vector in the wind reference frame are not far apart, especially during the reel-out phase, where the tether is relatively straight due to the high pulling force of the kite.

Measured attitude of the kite

The rotation matrix for the transformation from the earth to the tangential reference frame is calculated by:

$$\mathbb{T}_{\tau\mathrm{e}} = \begin{bmatrix} \sin\hat{\beta} & 0 & -\cos\hat{\beta} \\ 0 & 1 & 0 \\ \cos\hat{\beta} & 0 & \sin\hat{\beta} \end{bmatrix} \begin{bmatrix} \cos\left(\hat{\varphi}+\hat{\varphi}_{\mathrm{we}}\right) & \sin\left(\hat{\varphi}+\hat{\varphi}_{\mathrm{we}}\right) & 0 \\ -\sin\left(\hat{\varphi}+\hat{\varphi}_{\mathrm{we}}\right) & \cos\left(\hat{\varphi}+\hat{\varphi}_{\mathrm{we}}\right) & 0 \\ 0 & 0 & 1 \end{bmatrix}, \tag{C.1}$$

in which subscripts τ, w, and e refer to the tangential, wind, and earth reference frames, respectively, the hat denotes a measured quantity, β is the elevation angle, and φ is the azimuth angle.

The measured pitch, roll, and yaw of the wing of the kite are expressed using 3-2-1 Euler angles. The corresponding rotation matrix for the transformation from the earth to the top wing surface reference frame is calculated by:

$$\mathbb{T}_{\text{tws-e}} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\hat{\phi} & \sin\hat{\phi} \\ 0 & -\sin\hat{\phi} & \cos\hat{\phi} \end{bmatrix} \begin{bmatrix} \cos\hat{\theta} & 0 & -\sin\hat{\theta} \\ 0 & 1 & 0 \\ \sin\hat{\theta} & 0 & \cos\hat{\theta} \end{bmatrix} \begin{bmatrix} \cos\hat{\psi} & \sin\hat{\psi} & 0 \\ -\sin\hat{\psi} & \cos\hat{\psi} & 0 \\ 0 & 0 & 1 \end{bmatrix}, \tag{C.2}$$

in which subscripts tws and e refer to the top wing surface and earth reference frames, respectively, ϕ is the roll angle, θ is the pitch angle, and ψ is the yaw angle.

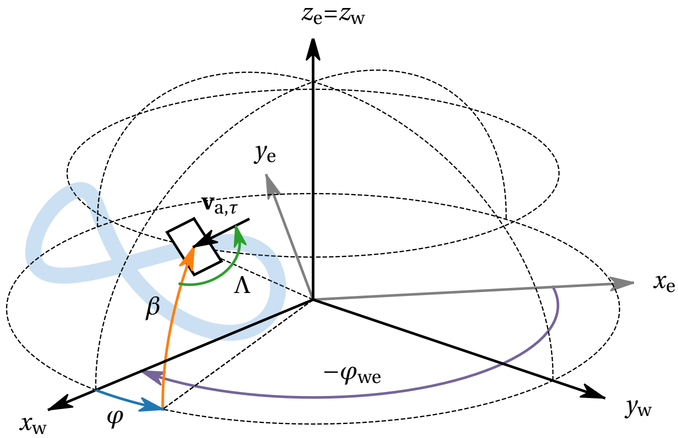

Figure C.1: Earth reference frame x_e, y_e, z_e and wind reference frame x_w, y_w, z_w together with the yawed tangential plane lying on the projection of a figure-of-eight flight path. This plane is yawed such that it heads into the apparent wind velocity and serves as a departure point for expressing the kite attitude, illustrated in Fig. C.2. The corresponding yaw angle Λ is equal to the kite heading in case of zero side slip.

The attitude of the kite is not affected by the depower signal and can be approximated by pitching the wing reference frame with the negative of the depower angle α_d depicted in Fig. 2.13

$$\mathbb{T}_{\text{b-tws}} = \begin{bmatrix} \cos\alpha_d & 0 & \sin\alpha_d \\ 0 & 1 & 0 \\ -\sin\alpha_d & 0 & \cos\alpha_d \end{bmatrix}, \tag{C.3}$$

in which subscript b denotes the bridle reference frame. The depower angle is calculated using a geometrical model from the power setting [127] and yields a nose-down pitch angle of roughly 6.6° during the reel-in phase.

The rotation matrix for the transformation from the tangential to the bridle reference frame is derived from the previously presented matrices:

$$\mathbb{T}_{b\tau} = \mathbb{T}_{\text{b-tws}}\,\mathbb{T}_{\text{tws-e}}\,\mathbb{T}_{\tau e}^{\top}. \tag{C.4}$$

A rotation matrix can be represented with a set of 3-2-1 Euler angles. The yaw, pitch, and roll corresponding to these three angles can be calculated using the lower expressions:

$$\psi = \operatorname{arctan2}(\mathbb{T}_{12}, \mathbb{T}_{11}), \tag{C.5}$$

$$\theta = -\operatorname{arctan2}\left(\mathbb{T}_{13}, \sqrt{\mathbb{T}_{23}^2 + \mathbb{T}_{33}^2}\right), \tag{C.6}$$

$$\phi = \operatorname{arctan2}(\mathbb{T}_{23}, \mathbb{T}_{33}), \tag{C.7}$$

in which $\mathbb{T}_{ij}$ denotes the transformation matrix element at the i^{th} row and j^{th} column. The Euler angles corresponding to $\mathbb{T}_{b\tau}$ are denoted without a subscript. The definitions of the pitch and roll angles are illustrated in Fig. C.2, taking the yawed tangential plane as the point of departure.

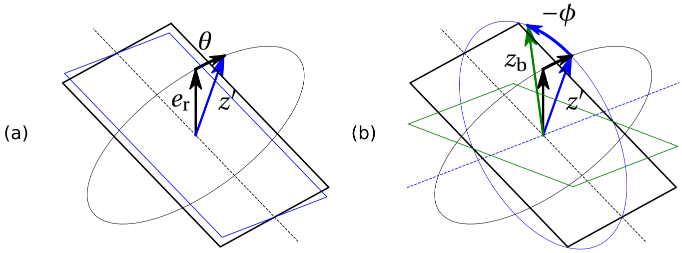

Figure C.2: Last two rotations in the 3-2-1 sequence (Euler angles) to get from the tangential to the bridle reference frame: (**a**) a positive pitch rotation and (**b**) a negative roll rotation. The black rectangle illustrates the yawed tangential plane, introduced in Fig. C.1.

Λ in Fig. C.1 describes the orientation of the tangential projection of the modelled apparent wind velocity, also shown in Fig. 6.1. In case of no side slip, Λ equals the heading angle. The heading angle inferred from measurements and Λ has a small periodic misalignment (not plotted), which may indicate a side slip. However, the constant wind assumption and measurement errors introduce too much uncertainty to confirm this. Also, the side slip angle was not measured in the studied test flight and thus can not be validated. Nevertheless, some side slip can be expected, as previously shown in the experiments by Oehler and Schmehl [66].

Modelled attitude of the kite

Expressing the Euler angles of the kite element of the model requires assigning a local reference frame to the element. The model does not specify a full reference frame but only specifies the axial direction of the element. This axial direction is used as the z-axis for the local reference frame. To differentiate between the roll and pitch, also the x-axis and y-axis need to be specified. The x-axis is chosen such that it lies in the plane spanned by the position vector and the vertical direction z_e. The y-axis then follows from the other two axes and is oriented horizontally.

Other than for securing the alignment between the roll and pitch definitions of the measured and modelled kite attitude, the yaw of the tether is not of interest to this study. It does not affect the kite attitude itself, and therefore, the resulting yaw angles are left out of Fig. 6.8. The modelled yaw of the kite is similar to that inferred from the wing attitude measurements and, thereby, facilitates comparing the measured and modelled roll and pitch.

References

1. Masson-Delmotte, V., Pörtner, H.-O., Skea, J., Zhai, P., Roberts, D., Shukla, P. R.: Global warming of 1.5°C. xiii, 616 p. : (2019). http://digitallibrary.un.org/record/3893415
2. United Nations Framework Convention on Climate Change (UNFCCC) Conference of the Parties (COP): Adoption of the Paris Agreement. Proposal by the President, Decision 1/CP.21, Dec 12, 2015. http://undocs.org/FCCC/CP/2015/L.9/Rev.1
3. Ritchie, H., Roser, M., Rosado, P.: CO2 and Greenhouse Gas Emissions. Our World in Data. https://ourworldindata.org/emissions-by-sector (2020)
4. Tracking Power 2021. Tracking report, IEA, Paris, France, 2021. https://www.iea.org/reports/tracking-power-2021
5. World Energy Outlook 2021. Flagship report, IEA, Paris, France, 2021. https://www.iea.org/reports/world-energy-outlook-2021
6. Veers, P., Dykes, K., Lantz, E., Barth, S., Bottasso, C. L., Carlson, O., Clifton, A., Green, J., Green, P., Holttinen, H., Laird, D., Lehtomäki, V., Lundquist, J. K., Manwell, J., Marquis, M., Meneveau, C., Moriarty, P., Munduate, X., Muskulus, M., Naughton, J., Pao, L., Paquette, J., Peinke, J., Robertson, A., Rodrigo, J. S., Sempreviva, A. M., Smith, J. C., Tuohy, A., Wiser, R.: Grand challenges in the science of wind energy. Science **366**(6464), eaau2027 (2019). doi: 10.1126/science.aau2027
7. Watson, S., Moro, A., Reis, V., Baniotopoulos, C., Barth, S., Bartoli, G., Bauer, F., Boelman, E., Bosse, D., Cherubini, A., Croce, A., Fagiano, L., Fontana, M., Gambier, A., Gkoumas, K., Golightly, C., Latour, M. I., Jamieson, P., Kaldellis, J., Macdonald, A., Murphy, J., Muskulus, M., Petrini, F., Pigolotti, L., Rasmussen, F., Schild, P., Schmehl, R., Stavridou, N., Tande, J., Taylor, N., Telsnig, T., Wiser, R.: Future emerging technologies in the wind power sector: A European perspective. Renewable and Sustainable Energy Reviews **113**, 109270 (2019). doi: 10.1016/j.rser.2019.109270
8. Schmidt, H., Vries, G. de, Renes, R. J., Schmehl, R.: The Social Acceptance of Airborne Wind Energy: A Literature Review. Energies **15**(4) (2022). doi: 10.3390/en15041384
9. Vermillion, C., Cobb, M., Fagiano, L., Leuthold, R., Diehl, M., Smith, R. S., Wood, T. A., Rapp, S., Schmehl, R., Olinger, D., Demetriou, M.: Electricity in the air: Insights from two decades of advanced control research and experimental flight testing of airborne wind energy systems. Annual Reviews in Control **52**, 330–357 (2021). doi: https://doi.org/10.1016/j.arcontrol.2021.03.002
10. Fagiano, L., Quack, M., Bauer, F., Carnel, L., Oland, E.: Autonomous Airborne Wind Energy Systems: Accomplishments and Challenges. Annual Review of Control, Robotics, and Autonomous Systems **5**(1), 603–631 (2022). doi: 10.1146/annurev-control-042820-124658
11. Wrage, S.: Making AWE a Reality. In: Schmehl, R., Fagiano, L., Croce, A., Thoms, S. (eds.). Book of Abstracts of the International Airborne Wind Energy Conference 2021, p. 77, Delft, Netherlands (2022). http://resolver.tudelft.nl/uuid:5e7ba020-de8d-4ed1-ac04-5466f8750a55

12. Peschel, J., Breuer, J.: Kitepower's Journey to the Islands and Beyond. In: Schmehl, R., Fagiano, L., Croce, A., Thoms, S. (eds.). Book of Abstracts of the International Airborne Wind Energy Conference 2021, p. 175, Delft, Netherlands (2022). http://resolver.tudelft.nl/uuid:c1759c23-1d89-4648-8301-d9410a0001a2

13. Van Hussen, K., Dietrich, E., Smeltink, J., Berentsen, K., Sleen, M. van der, Haffner, R., Fagiano, L.: Study on challenges in the commercialisation of airborne wind energy systems. ECORYS Report PP-05081-2016, European Commission, Brussels, Sept 2018. doi: 10.2777/87591

14. Commission, E.: Resource Efficient Automatic Conversion of High-Altitude Wind. https://cordis.europa.eu/project/id/691173. Accessed Aug 17, 2023

15. Commission, E.: Duurzame energie opwekken met vlieger. https://netherlands.representation.ec.europa.eu/strategie-en-prioriteiten/succesverhalen/duurzame-energie-opwekken-met-vlieger_nl. Accessed Aug 17, 2023

16. Breuer, J.: REACH: A H2020 FTI Project to Develop a 100 kW AWE System. In: Schmehl, R., Tulloch, O. (eds.). Book of Abstracts of the International Airborne Wind Energy Conference 2019, p. 41, Delft, Netherlands (2019). http://resolver.tudelft.nl/uuid:307feb11-1073-4c7f-87bb-df3c1529c78c

17. Diehl, M.: Airborne Wind Energy: Basic Concepts and Physical Foundations. In: Ahrens, U., Diehl, M., Schmehl, R. (eds.) Airborne Wind Energy, Green Energy and Technology, Chap. 1, pp. 3–22. Springer, Berlin Heidelberg (2013). doi: 10.1007/978-3-642-39965-7_1

18. Loyd, M. L.: Crosswind kite power. Journal of Energy **4**(3), 106–111 (1980). doi: 10.2514/3.48021

19. Vander Lind, D.: Analysis and Flight Test Validation of High Performance Airborne Wind Turbines. In: Ahrens, U., Diehl, M., Schmehl, R. (eds.) Airborne Wind Energy, Green Energy and Technology, Chap. 28, pp. 473–490. Springer, Berlin Heidelberg (2013). doi: 10.1007/978-3-642-39965-7_28

20. Salma, V., Friedl, F., Schmehl, R.: Reliability and Safety of Airborne Wind Energy Systems. Wind Energy **23**(2), 340–356 (2019). doi: 10.1002/we.2433

21. Fechner, U.: A Methodology for the Design of Kite-Power Control Systems. Ph.D. Thesis, Delft University of Technology, Nov 23, 2016. doi: 10.4233/uuid:85efaf4c-9dce-4111-bc91-7171b9da4b77

22. Kitemill: Great performance statistics. https://kitemill.webflow.io/news/energy-produced-during-testing. Accessed Nov 18, 2022

23. Oland, E., Brink, A. van der: Technical Progress in Kitemill. In: Schmehl, R., Fagiano, L., Croce, A., Thoms, S. (eds.). Book of Abstracts of the International Airborne Wind Energy Conference 2021, p. 61, Delft, Netherlands (2022). http://resolver.tudelft.nl/uuid:39e797a0-7959-44bb-9ccf-efb5351123d4

24. Malz, E., Koenemann, J., Sieberling, S., Gros, S.: A reference model for airborne wind energy systems for optimization and control. Renew. Energ. **140**, 1004–1011 (2019). doi: https://doi.org/10.1016/j.renene.2019.03.111

25. Stull, R.: Practical Meteorology: An Algebra-based Survey of Atmospheric Science. University of British Columbia (2017). https://www.eoas.ubc.ca/books/Practical_Meteorology/

26. Stull, R. B.: An introduction to boundary layer meteorology, vol. 13. Atmospheric and Oceanographic Sciences Library. Springer Netherlands (1988). doi: 10.1007/978-94-009-3027-8

27. Monin, A., Obukhov, A.: Basic laws of turbulent mixing in the ground layer of the atmosphere (in Russian). Trudy Geofiz. Inst. Akad. Nauk SSSR **151**, 163–187 (1954)

28. Peterson, E. W., Hennessey, J. P.: On the Use of Power Laws for Estimates of Wind Power Potential. J. Appl. Meteorol. **17**(3), 390–394 (1978). doi: 10.1175/1520-0450(1978)017<0390:OTUOPL>2.0.CO;2

29. Park, J., Basu, S., Manuel, L.: Large-eddy simulation of stable boundary layer turbulence and estimation of associated wind turbine loads. Wind Energy **17**(3), 359–384 (2014). doi: 10.1002/we.1580

30. Optis, M., Monahan, A., Bosveld, F. C.: Moving Beyond Monin–Obukhov Similarity Theory in Modelling Wind-Speed Profiles in the Lower Atmospheric Boundary Layer under Stable Stratification. Bound.-Lay. Meteorol. **153**(3), 497–514 (2014). doi: 10.1007/s10546-014-9953-z

31. Holtslag, M. C., Bierbooms, W. A. A. M., Bussel, G. J. W. van: Estimating atmospheric stability from observations and correcting wind shear models accordingly. J. Phys. Conf. Ser. **555**, 012052 (2014). doi: 10.1088/1742-6596/555/1/012052

32. Gryning, S.-E., Batchvarova, E., Brümmer, B., Jørgensen, H., Larsen, S.: On the extension of the wind profile over homogeneous terrain beyond the surface boundary layer. Bound.-Lay. Meteorol. **124**(2), 251–268 (2007). doi: 10.1007/s10546-007-9166-9

33. Holtslag, M., Bierbooms, W., Bussel, G. van: Extending the diabatic surface layer wind shear profile for offshore wind energy. Renew. Energ. **101**, 96–110 (2017). doi: https://doi.org/10.1016/j.renene.2016.08.031

34. ECMWF: IFS Documentation CY41R2 - Part IV: Physical Processes. eng. In: 4. doi: 10.21957/tr5rv27xu

35. Olauson, J.: ERA5: The new champion of wind power modelling? Renewable Energy **126**, 322–331 (2018). doi: 10.1016/j.renene.2018.03.056

36. Petersen, E. L., Mortensen, N. G., Landberg, L., Højstrup, J., Frank, H. P.: Wind power meteorology. Part II: siting and models. Wind Energy **1**(2), 55–72 (1998). doi: https://doi.org/10.1002/(SICI)1099-1824(199812)1:2<55::AID-WE5>3.0.CO;2-R

37. European Centre for Medium Range Weather Forecasts (ECMWF): ERA5 data documentation. Accessed 12 September 2022. https://confluence.ecmwf.int/display/CKB/ERA5%3A+compute+pressure+and+geopotential+on+model+levels%2C+geopotential+height+and+geometric+height

38. Wijnant, I. L., Ulft, B. van, Stratum, B. van, Barkmeijer, J., Onvlee, J., Valk, C. de, Knoop, S., Kok, S., Marseille, G. J., Klein Baltink, H., Stepek, A.: The Dutch Offshore Wind Atlas (DOWA): description of the dataset. TR-380, KNMI, 2019, p. 60. http://bibliotheek.knmi.nl/knmipubTR/TR380.pdf Accessed July 26, 2019

39. Bengtsson, L., Andrae, U., Aspelien, T., Batrak, Y., Calvo, J., Rooy, W. de, Gleeson, E., Hansen-Sass, B., Homleid, M., Hortal, M., Ivarsson, K.-I., Lenderink, G., Niemelä, S., Nielsen, K. P., Onvlee, J., Rontu, L., Samuelsson, P., Muñoz, D. S., Subias, A., Tijm, S., Toll, V., Yang, X., Køltzow, M. Ø.: The HARMONIE–AROME Model Configuration in the ALADIN–HIRLAM NWP System. Mon. Weather Rev. **145**(5), 1919–1935 (2017). doi: 10.1175/MWR-D-16-0417.1

40. Kalverla, P. C., Holtslag, A. A. M., Ronda, R. J., Steeneveld, G.-J.: Quality of wind characteristics in recent wind atlases over the North Sea. Quarterly Journal of the Royal Meteorological Society **146**(728), 1498–1515 (2020). doi: https://doi.org/10.1002/qj.3748

41. Miller, L. M., Gans, F., Kleidon, A.: Estimating maximum global land surface wind power extractability and associated climatic consequences. Earth System Dynamics **2**(1), 1–12 (2011). doi: 10.5194/esd-2-1-2011

42. Marvel, K., Kravitz, B., Caldeira, K.: Geophysical limits to global wind power. Nature Climate Change **3**, 118–121 (2013). doi: 10.1038/nclimate1683

43. Miller, L. M., Brunsell, N. A., Mechem, D. B., Gans, F., Monaghan, A. J., Vautard, R., Keith, D. W., Kleidon, A.: Two methods for estimating limits to large-scale wind power generation. Proceedings of the National Academy of Sciences (2015). doi: 10.1073/pnas.1408251112

44. Jacobson, M. Z., Archer, C. L.: Saturation wind power potential and its implications for wind energy. Proceedings of the National Academy of Sciences (PNAS) **109**(39), 15679–15684 (2012). doi: 10.1073/pnas.1208993109

45. Archer, C. L., Caldeira, K.: Global Assessment of High-Altitude Wind Power. Energies **2**(2), 307–319 (2009). doi: 10.3390/en20200307

46. van der Vlugt, R., Bley, A., Noom, M., Schmehl, R.: Quasi-Steady Model of a Pumping Kite Power System. Renewable Energy **131**, 83–99 (2019). doi: 10.1016/j.renene.2018.07.023

47. Ranneberg, M., Wölfle, D., Bormann, A., Rohde, P., Breipohl, F., Bastigkeit, I.: Fast Power Curve and Yield Estimation of Pumping Airborne Wind Energy Systems. In: Schmehl, R. (ed.) Airborne Wind Energy: Advances in Technology Development and Research, pp. 623–641. Springer, Singapore (2018). doi: 10.1007/978-981-10-1947-0_25

48. Williams, P.: Cable Modeling Approximations for Rapid Simulation. Journal of Guidance, Control, and Dynamics **40**(7), 1779–1788 (2017). doi: 10.2514/1.G002354

49. Fechner, U., Vlugt, R. van der, Schreuder, E., Schmehl, R.: Dynamic Model of a Pumping Kite Power System. Renewable Energy **83**, 705–716 (2015). doi: 10.1016/j.renene.2015.04.028

50. Houska, B., Diehl, M.: Optimal control for power generating kites. In: 2007 European Control Conference (ECC), pp. 3560–3567, (2007). doi: 10.23919/ECC.2007.7068861

51. Gros, S., Diehl, M.: Modeling of Airborne Wind Energy Systems in Natural Coordinates. In: Ahrens, U., Diehl, M., Schmehl, R. (eds.) Airborne Wind Energy, Green Energy and Technology, Chap. 10, pp. 181–203. Springer, Berlin Heidelberg (2013). doi: 10.1007/978-3-642-39965-7_10

52. Licitra, G., Bürger, A., Williams, P., Ruiterkamp, R., Diehl, M.: Aerodynamic model identification of an autonomous aircraft for airborne wind energy. Optimal Control Applications and Methods **40**(3), 422–447. doi: https://doi.org/10.1002/oca.2485

53. Zanon, M., Gros, S., Andersson, J., Diehl, M.: Airborne Wind Energy Based on Dual Airfoils. IEEE Transactions on Control Systems Technology **21**(4), 1215–1222 (2013). doi: 10.1109/TCST.2013.2257781

54. Antoniou, I., Pedersen, S. M., Enevoldsen, P. B.: Wind Shear and Uncertainties in Power Curve Measurement and Wind Resources. Wind Engineering **33**(5), 449–468 (2009). doi: 10.1260/030952409790291208

55. Europe, A. W.: Glossary – definitions of AWE-specific terms. https://airbornewindeurope.org/resources/glossary-2/. Accessed Nov 22, 2022

56. Licitra, G., Koenemann, J., Horn, G., Williams, P., Ruiterkamp, R., Diehl, M.: Viability assessment of a rigid wing airborne wind energy pumping system. In: 2017 21st International Conference on Process Control (PC), pp. 452–458, (2017). doi: 10.1109/PC.2017.7976256

57. Sommerfeld, M., Dörenkämper, M., De Schutter, J., Crawford, C.: Offshore and onshore ground-generation airborne wind energy power curve characterization. Wind Energy Science Discussions **2020**, 1–39 (2020). doi: 10.5194/wes-2020-120

58. Malz, E., Verendel, V., Gros, S.: Computing the power profiles for an Airborne Wind Energy system based on large-scale wind data. Renew. Energ. (2020). doi: https://doi.org/10.1016/j.renene.2020.06.056

59. Malz, E., Hedenus, F., Göransson, L., Verendel, V., Gros, S.: Drag-mode airborne wind energy vs. wind turbines: An analysis of power production, variability and geography. Energy **193**, 116765 (2020). doi: https://doi.org/10.1016/j.energy.2019.116765

60. Diezel, J. M., Reuder, J.: Using the Lidar-Validated Hindcast Model NORA3 for Resource Estimates of Airborne Wind Energy Systems. In: Schmehl, R., Fagiano, L., Croce, A., Thoms, S. (eds.). Book of Abstracts of the International Airborne Wind Energy Conference 2021, p. 89, Delft, Netherlands (2022). http://resolver.tudelft.nl/uuid:6917a009-8bd8-4062-ade8-596dcc3b30bb

61. Vlugt, R. van der, Peschel, J., Schmehl, R.: Design and Experimental Characterization of a Pumping Kite Power System. In: Ahrens, U., Diehl, M., Schmehl, R. (eds.) Airborne Wind Energy, Green Energy and Technology, Chap. 23, pp. 403–425. Springer, Berlin Heidelberg (2013). doi: 10.1007/978-3-642-39965-7_23

62. Fechner, U., Schmehl, R.: Model-Based Efficiency Analysis of Wind Power Conversion by a Pumping Kite Power System. In: Ahrens, U., Diehl, M., Schmehl, R. (eds.) Airborne Wind Energy, Green Energy and Technology, Chap. 14, pp. 249–269. Springer, Berlin Heidelberg (2013). doi: 10.1007/978-3-642-39965-7_14

63. Weber, J., Marquis, M., Lemke, A., Cooperman, A., Draxl, C., Lopez, A., Roberts, O., Shields, M.: Proceedings of the 2021 Airborne Wind Energy Workshop. (2021). https://www.osti.gov/biblio/1810067

64. Breuer, J., Peschel, J., Ghilardi, M.: First Airborne Wind Energy Operation on a Tropical Island. In: Schmehl, R., Fagiano, L., Croce, A., Thoms, S. (eds.). Book of Abstracts of the International Airborne Wind Energy Conference 2021, p. 51, Delft, Netherlands (2022). http://resolver.tudelft.nl/uuid:1ef9d8c7-c575-446b-86e6-c35fd9ec9ffb

65. Kitepower B.V.: Onshore Containerised AWES-100 Kitepower Falcon. https://thekitepower.com/product/. Accessed Mar 1, 2023

66. Oehler, J., Schmehl, R.: Aerodynamic characterization of a soft kite by in situ flow measurement. Wind Energy Science **4**(1), 1–21 (2019). doi: 10.5194/wes-4-1-2019

67. Peschel, J., Breuer, J., Schmehl, R.: Kitepower – Commercializing a 100 kW Mobile Wind Energy System. In: Diehl, M., Leuthold, R., Schmehl, R. (eds.). Book of Abstracts of the International Airborne Wind Energy Conference (AWEC 2017), pp. 47–51, Freiburg, Germany, Oct 5–6, 2017. http://resolver.tudelft.nl/uuid:9e9a6bdb-f5f4-4868-bed6-f25dc5ee42f0

68. Archer, C. L., Jacobson, M. Z.: Evaluation of global wind power. Journal of Geophysical Research: Atmospheres **110**(D12), 1–20 (2005). doi: 10.1029/2004JD005462

69. Adams, A. S., Keith, D. W.: Are global wind power resource estimates overstated? Environmental Research Letters **8**(1), 015021 (2013). doi: 10.1088/1748-9326/8/1/015021

70. Emeis, S. (ed.): Wind Energy Meteorology – Atmospheric Physics for Wind Power Generation. Green Energy and Technology. Springer, Cham (2018). doi: 10.1007/978-3-319-72859-9

71. Dupont, E., Koppelaar, R., Jeanmart, H.: Global available wind energy with physical and energy return on investment constraints. Applied Energy **209**, 322–338 (2018). doi: https://doi.org/10.1016/j.apenergy.2017.09.085

72. Archer, C. L., Caldeira, K.: Atlas of high altitude wind power, Department of Global Ecology, Carnegie Institute for Science, 260 Panama Street, Stanford, CA 94305, June 16, 2008. http://www.mdpi.com/1996-1073/2/2/307/s1

73. Miller, L. M., Gans, F., Kleidon, A.: Jet stream wind power as a renewable energy resource: little power, big impacts. Earth System Dynamics **2**(2), 201–212 (2011). doi: 10.5194/esd-2-201-2011

74. Archer, C. L., Monache, L. D., Rife, D. L.: Airborne wind energy: Optimal locations and variability. Renewable Energy **64**, 180–186 (2014). doi: 10.1016/j.renene.2013.10.044

75. Gambier, A., Bastigkeit, I., Nippold, E.: Projekt OnKites II : Untersuchung zu den Potentialen von Flugwindenergieanlagen (FWEA) Phase II. Final Project Report, Fraunhofer Institute for Wind Energy and Energy System Technology IWES, Bremerhaven, Germany, June 2017. 105 pp. doi: 10.2314/GBV:1009915452

76. Gambier, A.: Projekt OnKites : Untersuchung zu den Potentialen von Flugwindenergieanlagen (FWEA). Final Project Report, Fraunhofer Institute for Wind Energy and Energy System Technology IWES, Bremerhaven, Germany, 2014. 155 pp. doi: 10.2314/GBV:81573428X

77. Lunney, E., Ban, M., Duic, N., Foley, A.: A state-of-the-art review and feasibility analysis of high altitude wind power in Northern Ireland. Renewable and Sustainable Energy Reviews **68**, 899–911 (2017). doi: https://doi.org/10.1016/j.rser.2016.08.014

78. Yip, C. M. A., Gunturu, U. B., Stenchikov, G. L.: High-altitude wind resources in the Middle East. Scientific Reports **7**(9885), 1–10 (2017). doi: 10.1038/s41598-017-10130-6

79. Mann, S., Gunn, K., Harrison, G., Beare, B., Lazakis, I.: Wind Yield Assessment for Airborne Wind Energy. Poster presented at the EWEA Offshore 2015 conference, Copenhagen, Mar 2015. http://www.ewea.org/offshore2015/conference/allposters/PO090.pdf Accessed Jan 18, 2016

80. Ruiterkamp, R., Sieberling, S.: Description and Preliminary Test Results of a Six Degrees of Freedom Rigid Wing Pumping System. In: Ahrens, U., Diehl, M., Schmehl, R. (eds.) Airborne Wind Energy, Green Energy and Technology, Chap. 26, pp. 443–458. Springer, Berlin Heidelberg (2013). doi: 10.1007/978-3-642-39965-7_26

81. Salma, V., Ruiterkamp, R., Kruijff, M., Van Paassen, M. M. R., Schmehl, R.: Current and Expected Airspace Regulations for Airborne Wind Energy Systems. In: Schmehl, R. (ed.) Airborne Wind Energy – Advances in Technology Development and Research, Green Energy and Technology, Chap. 29, pp. 703–725. Springer, Singapore (2018). doi: 10.1007/978-981-10-1947-0_29

82. Bechtle, P., Schelbergen, M., Schmehl, R., Zillmann, U., Watson, S.: Python code for airborne wind energy resource analysis based on the ERA5 dataset, 2018. https://github.com/rschmehl/awe-era5/

83. Bechtle, P., Schelbergen, M., Schmehl, R., Zillmann, U., Watson, S.: ERA5 reanalysis data for Airborne Wind Energy Resource Analysis. Dataset, 2018. doi: 10.4121/uuid:646eaf3f-c90b-4f22-89bf-8986804def3c

84. Archer, C. L.: An Introduction to Meteorology for Airborne Wind Energy. In: Ahrens, U., Diehl, M., Schmehl, R. (eds.) Airborne Wind Energy, Green Energy and Technology, Chap. 5, pp. 81–94. Springer, Berlin Heidelberg (2013). doi: 10.1007/978-3-642-39965-7_5

85. Brown, A. R., Beljaars, A. C. M., Hersbach, H., Hollingsworth, A., Miller, M., Vasiljevic, D.: Wind turning across the marine atmospheric boundary layer. Q. J. Roy. Meteor. Soc. **131**(607), 1233–1250 (2005). doi: 10.1256/qj.04.163

86. Floors, R., Peña, A., Gryning, S.-E.: The effect of baroclinicity on the wind in the planetary boundary layer. Q. J. Roy. Meteor. Soc. **141**(687), 619–630 (2015). doi: 10.1002/qj.2386

87. Ranjha, R., Svensson, G., Tjernström, M., Semedo, A.: Global distribution and seasonal variability of coastal low-level jets derived from ERA-Interim reanalysis. Tellus A **65**(1), 20412 (2013). doi: 10.3402/tellusa.v65i0.20412

88. Kalverla, P. C., Duncan Jr., J. B., Steeneveld, G.-J., Holtslag, A. A. M.: Low-level jets over the North Sea based on ERA5 and observations: together they do better. Wind Energ. Sci. **4**(2), 193–209 (2019). doi: 10.5194/wes-4-193-2019

89. Basu, S.: A simple recipe for estimating atmospheric stability solely based on surface-layer wind speed profile. Wind Energy **21**(10), 937–941 (2018). doi: 10.1002/we.2203

90. Sommerfeld, M., Crawford, C., Monahan, A., Bastigkeit, I.: LiDAR-based characterization of mid-altitude wind conditions for airborne wind energy systems. Wind Energy **22**(8), 1101–1120 (2019). doi: 10.1002/we.2343

91. Durán, P., Basu, S., Meißner, C., Adaramola, M. S.: Automated classification of simulated wind field patterns from multiphysics ensemble forecasts. Wind Energy **23**(4), 898–914 (2020). doi: 10.1002/we.2462

92. Amante, C., Eakins, B. W.: ETOPO1 1 Arc-Minute Global Relief Model: Procedures, Data Sources and Analysis. NOAA Technical Memorandum NESDIS NGDC-24. National Geophysical Data Center, NOAA. (2009). Accessed Apr 1, 2020

93. Kalverla, P. C., Steeneveld, G.-J., Ronda, R. J., Holtslag, A. A.: An observational climatology of anomalous wind events at offshore meteomast IJmuiden (North Sea). J. Wind Eng. Ind. Aerod. **165**, 86–99 (2017). doi: https://doi.org/10.1016/j.jweia.2017.03.008

94. Wijnant, I. L., Brink, H. W. van den, Stepek, A.: North Sea wind climatology part 1: a review of existing wind atlases. TR-342, KNMI, May 2015, p. 62. http://bibliotheek.knmi.nl/knmipubTR/TR342.pdf Accessed June 1, 2020

95. Kelly, M., Gryning, S.-E.: Long-Term Mean Wind Profiles Based on Similarity Theory. Bound.-Lay. Meteorol. **136**(3), 377–490 (2010). doi: 10.1007/s10546-010-9509-9

96. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine Learning in Python. J. Mach. Learn. Res. **12**, 2825–2830 (2011)

97. Burk, S. D., Thompson, W. T.: The Summertime Low-Level Jet and Marine Boundary Layer Structure along the California Coast. Mon. Weather Rev. **124**(4), 668–686 (1996). doi: 10.1175/1520-0493(1996)124<0668:TSLLJA>2.0.CO;2

98. Parish, T. R.: Forcing of the Summertime Low-Level Jet along the California Coast. J. Appl. Meteorol. **39**(12), 2421–2433 (2000). doi: 10.1175/1520-0450(2000)039<2421:FOTSLL>2.0.CO;2

99. Mahrt, L., Vickers, D., Andreas, E. L.: Low-Level Wind Maxima and Structure of the Stably Stratified Boundary Layer in the Coastal Zone. J. Appl. Meteorol. Clim. **53**(2), 363–376 (2014). doi: 10.1175/JAMC-D-13-0170.1

100. Shapiro, A., Fedorovich, E., Rahimi, S.: A Unified Theory for the Great Plains Nocturnal Low-Level Jet. J. Atmos. Sci. **73**(8), 3037–3057 (2016). doi: 10.1175/JAS-D-15-0307.1

101. Verkaik, J. W.: On Wind and Roughness over Land. Ph.D. Thesis, Wageningen University, 2006. https://edepot.wur.nl/121786 Accessed June 1, 2020

102. Schmehl, R., Noom, M., Vlugt, R. van der: Traction Power Generation with Tethered Wings. In: Ahrens, U., Diehl, M., Schmehl, R. (eds.) Airborne Wind Energy, Green Energy and Technology, Chap. 2, pp. 23–45. Springer, Berlin Heidelberg (2013). doi: 10.1007/978-3-642-39965-7_2

103. Houska, B., Diehl, M.: Optimal control for power generating kites. In: Proceedings of the 9th European Control Conference, pp. 3560–3567, Kos, Greece, July 2–5, 2007

104. Gohl, F., Luchsinger, R. H.: Simulation Based Wing Design for Kite Power. In: Ahrens, U., Diehl, M., Schmehl, R. (eds.) Airborne Wind Energy, Green Energy and Technology, Chap. 18, pp. 325–338. Springer, Berlin Heidelberg (2013). doi: 10.1007/978-3-642-39965-7_18

105. Rapp, S., Schmehl, R., Oland, E., Haas, T.: Cascaded Pumping Cycle Control for Rigid Wing Airborne Wind Energy Systems. Journal of Guidance, Control, and Dynamics **42**(11), 2456–2473 (2019). doi: 10.2514/1.G004246

106. Williams, P., Lansdorp, B., Ockels, W. J.: Modeling and Control of a Kite on a Variable Length Flexible Inelastic Tether. AIAA Paper 2007-6705. In: Proceedings of the AIAA Modeling and Simulation Technologies Conference and Exhibit, Hilton Head, SC, USA, Aug 20–23, 2007. doi: 10.2514/6.2007-6705

107. Sánchez-Arriaga, G., Pastor-Rodríguez, A., Sanjurjo-Rivo, M., Schmehl, R.: A Lagrangian Flight Simulator for Ground and Fly Generation Airborne Wind Energy Systems. Applied Mathematical Modelling **69**, 665–684 (2019). doi: 10.1016/j.apm.2018.12.016

108. Schmehl, R. (ed.): Airborne Wind Energy – Advances in Technology Development and Research. Green Energy and Technology. Springer, Singapore (2018). 754 pp. doi: 10.1007/978-981-10-1947-0

109. Breukels, J., Schmehl, R., Ockels, W.: Aeroelastic Simulation of Flexible Membrane Wings based on Multibody System Dynamics. In: Ahrens, U., Diehl, M., Schmehl, R. (eds.) Airborne Wind Energy, Green Energy and Technology, Chap. 16, pp. 287–305. Springer, Berlin Heidelberg (2013). doi: 10.1007/978-3-642-39965-7_16

110. Bosch, A., Schmehl, R., Tiso, P., Rixen, D.: Nonlinear Aeroelasticity, Flight Dynamics and Control of a Flexible Membrane Traction Kite. In: Ahrens, U., Diehl, M., Schmehl, R. (eds.) Airborne Wind Energy, Green Energy and Technology, Chap. 17, pp. 307–323. Springer, Berlin Heidelberg (2013). doi: 10.1007/978-3-642-39965-7_17

111. Geschiere, N. H.: Dynamic modelling of a flexible kite for power generation: Coupling a fluid-structure solver to a dynamic particle system, May 1, 2014. http://resolver.tudelft.nl/uuid:6478003a-3c77-40ce-862e-24579dcd1eab

112. Duport, C.: Modeling with consideration of the fluid-structure interaction of the behavior under load of a kite for auxiliary traction of ships. PhD thesis, ENSTA Bretagne, 2018

113. Oehler, J., Reijen, M. van, Schmehl, R.: Experimental investigation of soft kite performance during turning maneuvers. Journal of Physics: Conference Series **1037**(5), 052004 (2018). doi: 10.1088/1742-6596/1037/5/052004

114. Thedens, P.: An Integrated Aero-Structural Model for Ram-Air Kite Simulations With Application to Airborne Wind Energy. Ph.D. Thesis, Delft University of Technology, Delft, Apr 25, 2022. doi: 10.4233/uuid:16e90401-62fc-4bc3-bf04-7a8c7bb0e2ee

115. Folkersma, M. A. M.: Aeroelasticity of Membrane Kites: Airborne Wind Energy Applications. Ph.D. Thesis, Delft University of Technology, Delft, Sept 7, 2022. doi: 10.4233/uuid:eae39f5a-49bc-438b-948f-b6ab51208068

116. Poland, J.: Modelling aeroelastic deformation of soft wing membrane kites. http://resolver.tudelft.nl/uuid:39d67249-53c9-47b4-84c0-ddac948413a5. M.Sc.Thesis, Delft University of Technology, 2022

117. Cayon, O.: Fast aeroelastic model of a leading-edge inflatable kite for the design phase of airborne wind energy systems. http://resolver.tudelft.nl/uuid:aede2a25-4776-473a-8a75-fb6b17b1a690. M.Sc.Thesis, Delft University of Technology, 2022

118. Roullier, A.: Experimental analysis of a kite system's dynamics. M.Sc.Thesis, EPFL, 2020

119. Andersson, J. A. E., Gillis, J., Horn, G., Rawlings, J. B., Diehl, M.: CasADi – A software framework for nonlinear optimization and optimal control. Mathematical Programming Computation **11**, 1–36 (2019). doi: 10.1007/s12532-018-0139-4

120. Luchsinger, R. H.: Pumping Cycle Kite Power. In: Ahrens, U., Diehl, M., Schmehl, R. (eds.) Airborne Wind Energy, Green Energy and Technology, Chap. 3, pp. 47–64. Springer, Berlin Heidelberg (2013). doi: 10.1007/978-3-642-39965-7_3

121. Kulfan, B. M.: Universal Parametric Geometry Representation Method. Journal of Aircraft **45**(1), 142–158 (2008). doi: 10.2514/1.29958

122. Licitra, G., Koenemann, J., Bürger, A., Williams, P., Ruiterkamp, R., Diehl, M.: Performance assessment of a rigid wing Airborne Wind Energy pumping system. Energy **173**, 569–585 (2019). doi: https://doi.org/10.1016/j.energy.2019.02.064

123. Erhard, M., Strauch, H.: Flight control of tethered kites in autonomous pumping cycles for airborne wind energy. Control Engineering Practice **40**, 13–26 (2015). doi: 10.1016/j.conengprac.2015.03.001

124. Perez, R. E., Jansen, P. W., Martins, J. R. R. A.: pyOpt: A Python-Based Object-Oriented Framework for Nonlinear Constrained Optimization. Struct. Multidiscip. O. **45**(1), 101–118 (2012). doi: 10.1007/s00158-011-0666-3

125. Haas, T., De Schutter, J., Diehl, M., Meyers, J.: Large-eddy simulation of airborne wind energy farms. Wind Energy Science **7**(3), 1093–1135 (2022). doi: 10.5194/wes-7-1093-2022

126. Borobia, R., Sanchez-Arriaga, G., Serino, A., Schmehl, R.: Flight Path Reconstruction and Flight Test of Four-line Power Kites. 1–11 (2018). doi: 10.2514/1.G003581

127. Schelbergen, M., Kalverla, P. C., Schmehl, R., Watson, S. J.: Clustering Wind Profile Shapes to Estimate Airborne Wind Energy Production. Wind Energy Science **5**(3), 1097–1120 (2020). doi: 10.5194/wes-5-1097-2020

List of Publications

Journal and conference papers

1. **Schelbergen, M.**, Schmehl, R.: Swinging Motion of a Kite with Suspended Control Unit Flying Turning Manoeuvres. Wind Energy Science Discussions (2023). doi: 10.5194/wes-2023-121

2. **Schelbergen, M.**, Kalverla, P. C., Schmehl, R., Watson, S. J.: Clustering Wind Profile Shapes to Estimate Airborne Wind Energy Production. Wind Energy Science 5(3), 1097–1120 (2020). doi: 10.5194/wes-5-1097-2020

3. **Schelbergen, M.**, Schmehl, R.: Validation of the Quasi-steady Performance Model for Pumping Airborne Wind Energy Systems. Journal of Physics: Conference Series 1618(3) (2020). doi: 10.1088/1742-6596/1618/3/032003

4. Bechtle, P., **Schelbergen, M.**, Schmehl, R., Zillmann, U., Watson, S.: Airborne Wind Energy Resource Analysis. Renewable Energy 141, 1103–1116 (2019). doi: 10.1016/j.renene.2019.03.118

Data and software repositories

1. Bechtle, P., **Schelbergen, M.**, Schmehl, R., Zillmann, U., Watson, S.: ERA5 reanalysis data for Airborne Wind Energy Resource Analysis. 4TU.ResearchData. Dataset. doi: 10.4121/uuid:646eaf3f-c90b-4f22-89bf-8986804def3c

2. **Schelbergen, M.**, Schmehl, R., Buchholz, B., Breuer, J., Peschel, J.: Kitepower Flight Data Acquired on 8 October 2019. 4TU.ResearchData. Dataset. doi: 10.4121/19376174

3. **Schelbergen, M.**, Schmehl, R.: Prevalent Wind Profile Shapes and Magnitude Probability Distributions for an On- and Offshore Location in the Netherlands. 4TU.ResearchData. Dataset. doi: 10.4121/c3957db1-55b9-448a-95d6-8ad5b6c7494d

Software repositories can be found on:
github.com/awegroup and github.com/markschelbergen