AWESPA: a modular toolchain for assessing airborne wind energy systems and farms

J.A.A. Bredael

Delft University of Technology

TU Delft

AWESPA: a modular toolchain for assessing airborne wind energy systems and farms

by

J.A.A. Bredael

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Wednesday July 29, 2026 at 13:00.

Student number:	5583810	
Project duration:	November 3, 2025 – July 29, 2026	
Thesis committee:	Dr. D. Ragni,	TU Delft, Chair
	Dr. Ir. D. A. M. Tavernier,	TU Delft, Examiner
	Dr.-Ing. R. Schmehl,	TU Delft, Responsible supervisor
	Ir. O. Cayon,	TU Delft, Daily supervisor
	Ir. H. Yousef,	TU Delft, Daily supervisor

Cover: Kitepower V9.60 kite ready for launch in Bangor Erris

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

TU Delft

Preface

I started my studies in Aerospace Engineering because of my interest in spaceflight. During the first year, however, I realized that my main interest was not in spacecraft, but in renewable energy. This interest eventually led me to work on airborne wind energy during my bachelor end project. That project was my first real introduction to the field and showed me how technically challenging and exciting airborne wind energy is. This interest was strengthened further when I followed the airborne wind energy course taught by Dr.-Ing. Roland Schmehl. In the summer of 2025, I had the opportunity to do an internship at Kitepower, where I worked in the field on flight tests. Being involved in practical test operations gave me a much better understanding of the amount of work required to bring airborne wind energy technology closer to application. With this background, starting a thesis in this field felt like a natural next step.

This thesis marks the end of my master's degree in Wind Energy at TU Delft. I would like to thank my supervisors, Dr.-Ing. Roland Schmehl, PhD candidate ir. Oriol Cayon, and PhD candidate ir. Hanna Yousef, for their supervision, expertise, enthusiasm, and feedback throughout this project. Through our discussions, they helped me to clarify the direction of the research and improve the thesis at each stage of the project. I would especially like to thank Dr.-Ing. Roland Schmehl for all the opportunities to present and share my work, and for introducing me to the airborne wind energy community.

On a personal note, I would like to thank my parents for their support throughout my studies. I am also very grateful to my girlfriend and friends, who have endured listening to me talk about kites over the past few years. Finally, I would like to give a special thanks to Alyssa Van Meijvis for designing the awesIO logo.

This thesis has been a challenging but rewarding final step in my studies. It allowed me to contribute to a field that I have become genuinely interested in, and I am grateful to everyone who supported me along the way.

J.A.A. Bredael
Delft, July 2026

Abstract

Airborne wind energy systems can access wind resources at higher altitudes while using less material than conventional wind turbines. Their performance assessment, however, requires several connected modeling steps, including wind resource analysis, power estimation, farm operation, and annual energy production calculation. These steps are often implemented in separate tools with different input and output formats, which makes results difficult to compare and analyses difficult to repeat. This thesis develops AWESPA, a modular Python toolchain for airborne wind energy system and farm performance assessment, together with awesIO, a standardized and validated input and output framework for airborne wind energy data.

The toolchain is built around exchangeable modules for each step in the performance assessment. The awesIO files provide common data definitions for the main information exchanged between modules and can be checked against predefined schemas to detect missing parameters, inconsistent file structures, and incorrect data definitions. The wind resource module converts ERA5 wind data into representative vertical wind profiles using principal component analysis and k-means clustering. A comparison with lidar measurements at the Bangor Erris test site shows that ERA5 captures the first-order wind climate, including the mean profile shape, Weibull scale parameter, and dominant wind direction. However, ERA5 underrepresents local wind-speed variability and produces a narrower range of profile shapes than the lidar data, especially for high-shear, low-shear, and low-level-jet-like profiles.

Three power estimation models are implemented and compared: the original Luchsinger model, the soft-kite Luchsinger model, and an inertia-free quasi-steady model. The Luchsinger models are useful for testing the modular structure of AWESPA, but give more optimistic power estimates than the inertia-free quasi-steady model. The difference is mainly caused by the treatment of gravitational loading and the reel-in phase. The quasi-steady model is therefore used for the case-study power estimates and is extended with variable reel-out elevation angles, a reel-out to reel-in transition, and constrained power-curve optimization.

At the farm level, a geometric maximum phase offset is derived for closely spaced wind-aligned systems with overlapping reel-in and reel-out envelopes. The results show that wind-aligned systems do not always need to operate synchronously. Limited asynchronous operation can reduce the peak-to-peak variation of the aggregated farm power output while maintaining geometric separation.

Two case studies are considered: a $100\ \mathrm{kW}$ isolated system at Bangor Erris and a $1\ \mathrm{MW}$ farm at Maasvlakte 2. The first case study compares system design choices and evaluates the selected system using site-specific wind profiles, giving an annual energy production of $488.6\ \mathrm{MWh}$ with the clustered wind-profile representation. The second case study applies the same system in a 10-system farm and gives an annual energy production of $4361\ \mathrm{MWh}$ with the clustered wind-profile representation. Together, the case studies show how AWESPA can be used to assess different sites, system configurations, and farm operating strategies in a consistent and reusable way.

Contents

List of Figures

List of Tables

Nomenclature

Abbreviations

Abbreviation	Definition
AEP	Annual energy production
AWE	Airborne wind energy
AWES	Airborne wind energy system
AWESPA	Airborne Wind Energy Systems Performance Assessment
DOWA	Dutch Offshore Wind Atlas
ECMWF	European Centre for Medium-Range Weather Forecasts
ERA5	Fifth-generation ECMWF atmospheric reanalysis
IQR	Interquartile range
ISA	International Standard Atmosphere
KCU	Kite control unit
LES	Large-eddy simulation
PCA	Principal component analysis
QSM	Quasi-steady model
SI	International System of Units
SLSQP	Sequential Least Squares Programming
WCSS	Within-cluster sum of squares
YAML	YAML Ain't Markup Language

Symbols

Symbol	Definition	Unit
A	Projected kite surface area	$[\mathrm{m}^2]$
A_{flat}	Flat kite surface area	$[\mathrm{m}^2]$
A_k	ERA5 hybrid-coordinate coefficient at model half-level k	$[\mathrm{Pa}]$
B_k	ERA5 hybrid-coordinate coefficient at model half-level k	$[-]$
c	Weibull scale parameter	$[\mathrm{m\,s^{-1}}]$
C_{D}	Drag coefficient	$[-]$
$C_{\mathrm{D,i}}$	Depowered drag coefficient during reel-in	$[-]$
$C_{\mathrm{D,o}}$	Powered drag coefficient during reel-out	$[-]$
$C_{\mathrm{D,t}}$	Tether drag coefficient	$[-]$
C_{L}	Lift coefficient	$[-]$
$C_{\mathrm{L,i}}$	Depowered lift coefficient during reel-in	$[-]$
$C_{\mathrm{L,o}}$	Powered lift coefficient during reel-out	$[-]$
d	Ground-station spacing	$[\mathrm{m}]$
d_{t}	Tether diameter	$[\mathrm{m}]$
E	Lift-to-drag ratio	$[-]$
E_{c}	Net mechanical energy per cycle	$[\mathrm{J}]$
E_{i}	Lift-to-drag ratio during reel-in	$[-]$
E_{o}	Lift-to-drag ratio during reel-out	$[-]$
f	Dimensionless reeling factor	$[-]$
f_{i}	Reel-in reeling factor	$[-]$

Symbol	Definition	Unit
f_o	Reel-out reeling factor	$[-]$
F_a	Aerodynamic force	$[\mathrm{N}]$
F_t	Tether force	$[\mathrm{N}]$
$F_\mathrm{t,i}$	Reel-in tether force	$[\mathrm{N}]$
$F_\mathrm{t,min}$	Minimum tether force	$[\mathrm{N}]$
$F_\mathrm{t,n}$	Nominal tether-force limit	$[\mathrm{N}]$
$F_\mathrm{t,o}$	Reel-out tether force	$[\mathrm{N}]$
g	Gravitational acceleration	$[\mathrm{m\,s^{-2}}]$
g_ref	Reference gravitational acceleration used for height conversion	$[\mathrm{m\,s^{-2}}]$
G	Prescribed lift-to-drag ratio in the QSM solver	$[-]$
J	Peak-to-peak farm-power objective function	$[\mathrm{W}]$
k	Weibull shape parameter or number of clusters	$[-]$
l	Tether length	$[\mathrm{m}]$
l_max	Maximum tether length	$[\mathrm{m}]$
l_min	Minimum tether length	$[\mathrm{m}]$
l_reel	Reeled tether length during one pumping phase	$[\mathrm{m}]$
m	Mass	$[\mathrm{kg}]$
m_t	Tether mass	$[\mathrm{kg}]$
p_c	Normalized mechanical cycle-power factor	$[-]$
p_k	Pressure at ERA5 model half-level k	$[\mathrm{Pa}]$
p_s	Surface pressure	$[\mathrm{Pa}]$
P	Mechanical power	$[\mathrm{W}]$
P_farm	Aggregated farm power output	$[\mathrm{W}]$
P_n	Nominal mechanical power limit	$[\mathrm{W}]$
P_w	Wind power density	$[\mathrm{W\,m^{-2}}]$
q	Specific humidity	$[\mathrm{kg\,kg^{-1}}]$
r_max	Maximum tether length	$[\mathrm{m}]$
$r_\mathrm{m,o}$	Mean tether length	$[\mathrm{m}]$
r_min	Minimum tether length	$[\mathrm{m}]$
R_d	Specific gas constant for dry air	$[\mathrm{J\,kg^{-1}\,K^{-1}}]$
s	Reel-out distance before envelope overlap starts	$[\mathrm{m}]$
t	Time	$[\mathrm{s}]$
t_i	Reel-in time	$[\mathrm{s}]$
t_o	Reel-out time	$[\mathrm{s}]$
t_s	Geometric time offset between neighboring systems	$[\mathrm{s}]$
t_safe	Safety-time margin in the farm phase constraint	$[\mathrm{s}]$
T	Temperature	$[\mathrm{K}]$
T_v	Virtual temperature	$[\mathrm{K}]$
u_norm	Normalized wind speed	$[-]$
u^*	Friction velocity	$[\mathrm{m\,s^{-1}}]$
v	Wind velocity	$[\mathrm{m\,s^{-1}}]$
v_a	Apparent wind speed	$[\mathrm{m\,s^{-1}}]$
v_i	Reel-in speed	$[\mathrm{m\,s^{-1}}]$
$v_\mathrm{i,max}$	Maximum reel-in speed	$[\mathrm{m\,s^{-1}}]$
v_o	Reel-out speed	$[\mathrm{m\,s^{-1}}]$
$v_\mathrm{o,max}$	Maximum reel-out speed	$[\mathrm{m\,s^{-1}}]$
v_r	Reeling speed	$[\mathrm{m\,s^{-1}}]$
$v_\mathrm{r,max}$	Maximum drum reeling speed	$[\mathrm{m\,s^{-1}}]$
v_ref	Reference wind speed	$[\mathrm{m\,s^{-1}}]$
$\bar{v}_\mathrm{w}$	Average wind speed over the operating altitude interval	$[\mathrm{m\,s^{-1}}]$
w_{ij}	Bilinear interpolation weight for grid point (i,j)	$[-]$
z	Height above ground level	$[\mathrm{m}]$
z_0	Surface roughness length	$[\mathrm{m}]$

Symbol	Definition	Unit
z_{max}	Maximum altitude of the operating interval	$[\mathrm{m}]$
z_{min}	Minimum altitude of the operating interval	$[\mathrm{m}]$
z_{ref}	Reference height	$[\mathrm{m}]$
α	Power-law shear exponent	$[-]$
β	Elevation angle	$[\mathrm{deg}]$ or $[-]$
β_{i}	Reel-in elevation angle	$[\mathrm{deg}]$
β_{o}	Reel-out elevation angle	$[\mathrm{deg}]$
γ	Angle in the wind-aligned farm-envelope geometry	$[\mathrm{deg}]$
γ_{i}	Aerodynamic force factor during reel-in	$[-]$
γ_{o}	Aerodynamic force factor during reel-out	$[-]$
$\Delta\beta_{\mathrm{i}}$	Reel-in half-cone width	$[\mathrm{deg}]$
$\Delta\beta_{\mathrm{o}}$	Reel-out half-cone width	$[\mathrm{deg}]$
$\Delta\Phi_k$	Geopotential difference across layer k	$[\mathrm{m}^2\,\mathrm{s}^{-2}]$
$\Delta\psi_{\mathrm{geom}}$	Geometric phase difference	$[-]$
$\Delta\psi_{\mathrm{max}}$	Maximum allowed phase difference	$[-]$
θ	Elevation angle	$[\mathrm{deg}]$
$\theta_{\mathrm{RORI,end}}$	End elevation angle of the reel-out to reel-in transition	$[\mathrm{deg}]$
κ	Kinematic ratio in the QSM steady-state solver	$[-]$
λ	Longitude	$[\mathrm{deg}]$
μ_{F}	Wind-speed ratio used in the force-limited operating region	$[-]$
μ_{P}	Wind-speed ratio used in the power-limited operating region	$[-]$
ρ	Air density	$[\mathrm{kg}\,\mathrm{m}^{-3}]$
ρ_{t}	Tether material density	$[\mathrm{kg}\,\mathrm{m}^{-3}]$
ϕ	Azimuth angle or latitude	$[\mathrm{deg}]$
Φ	Geopotential	$[\mathrm{m}^2\,\mathrm{s}^{-2}]$
Φ_{s}	Surface geopotential	$[\mathrm{m}^2\,\mathrm{s}^{-2}]$

1 Introduction

Climate change is caused by human activities, mainly through greenhouse gas emissions, and the global energy system is the largest source of CO_2 emissions [1]. Limiting global warming therefore requires large reductions in energy-sector emissions while maintaining a reliable energy supply. Wind energy is already used widely for this purpose, but conventional wind turbines are limited by tower height, foundation requirements, transport constraints, and the amount of material needed for large rotor diameters. Airborne wind energy (AWE) is being developed as an alternative wind energy technology that uses tethered flying devices to harvest wind energy at higher altitudes than conventional wind turbines [2].

An AWE system flies the main aerodynamic component in the air instead of supporting it with a tower and is connected to a ground station via a tether. This allows the system to operate at altitudes above conventional wind turbine hub heights and reduces the need for large tower structures. Most AWE concepts use crosswind flight, where the wing flies mainly perpendicular to the wind direction. This increases the apparent wind speed at the wing and therefore increases the aerodynamic load on the flying device [2].

AWE systems can be divided into systems with on-board generation and systems with ground-based generation. In on-board generation, electricity is produced on the flying device and transmitted to the ground through the tether. In ground-based generation, the tether transmits mechanical power to a generator on the ground. An example of ground-based generation is pumping AWE systems, which is the focus of this thesis. In these systems, the generator is connected to a winch or drum, and power is produced by reeling out the tether under high tension [2], [3].

The operating principle of a ground-generation pumping system is shown in Figure 1.1. The cycle consists of a traction phase, a retraction phase, and two transition phases. During the traction phase, the wing flies crosswind patterns, often figure-eight or circular paths, while the tether is reeled out. The high apparent wind speed during this phase creates a large tether force, which drives the ground-based generator. Once the maximum tether length is reached, the system enters the retraction phase through a transition phase. The wing is then flown in a low-force state and the tether is reeled back in. While this consumes energy, the required energy is smaller than the energy generated during traction. The difference between the generated and consumed energy gives the net energy output of the cycle [2].

The main control objective of a pumping system is to maintain a large force during reel-out and a small force during reel-in. During traction, the flight path is chosen to keep the wing in a part of the wind window where it can generate a high tether force. During retraction, the wing is depowered or guided along a flight path that reduces the aerodynamic load. This alternating operation makes the power output periodic. The average power of the system depends on the phase durations, the tether force level in each phase, and the reel-out and reel-in speeds [3], [5].

The performance assessment of AWE systems requires several linked models. The wind resource model must describe the vertical wind profile over the altitude range used by the system, because the

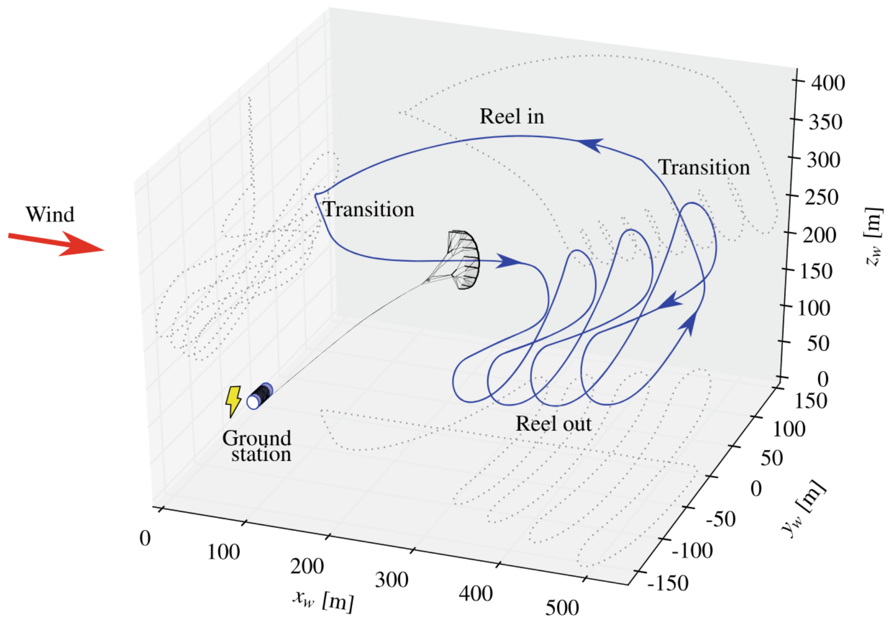

Figure 1.1: The flight path of the flexible-kite pumping AWE system, with kite and drum not to scale, from [4].

profile shape can affect both the power output and the operating strategy. For pumping AWE systems, the power model must model the cyclic operation because the system alternates between a power-producing reel-out phase and a power-consuming reel-in phase. At the farm level, this cyclic behavior affects the aggregated power output and creates a relation between layout, spacing, and phase timing.

These modeling steps are often treated separately. Wind resource tools, power models, and farm models can use different input files, output files, assumptions, and naming conventions. This makes it difficult to combine models in one workflow, repeat analyses, or compare results from different modeling choices. It also makes case-study analysis less reusable, because changes to one part of the workflow can require manual changes in other parts.

This issue is relevant for the case studies considered in this thesis. The first case study considers a $100\ \mathrm{kW}$ isolated AWE system at the Kitepower test site in Bangor Erris, Ireland. The second case study considers a $1\ \mathrm{MW}$ AWE farm at the Tweede Maasvlakte in the Port of Rotterdam. Both case studies require wind resource processing, power estimation, and farm output modeling. A consistent workflow is therefore needed to apply these steps to different sites and system configurations.

Earlier work in the open-source AWERA tool combined wind profile clustering with power estimation for AWE resource assessment [6], [7]. However, the present work requires a workflow in which the different modeling steps can be exchanged more easily and in which the main input and output files follow a common structure. This leads to the research objective defined in Chapter 3:

> To enable consistent, reusable, and comparable performance assessment of airborne wind energy systems by developing a modular toolchain with standardized input and output definitions and interchangeable modeling modules.

This objective is addressed through the development of the Airborne Wind Energy Systems Performance Assessment toolchain, abbreviated as AWESPA. AWESPA is a modular Python toolchain that connects the main steps of AWE performance assessment. The toolchain is structured so that individual modules can be replaced without changing the full workflow. This makes it possible to compare different modeling choices using the same system description, wind resource input, and output structure.

During the development of AWESPA [8], the need for a common input and output structure led to the development of awesIO [9]. awesIO defines standardized files for AWE system configurations, wind resources, and power curves. These files are used to exchange data between the modules in AWESPA and to reduce tool-specific data conversion. The structure of awesIO and AWESPA is described in Chapter 4.

Chapter 2 reviews the literature needed for AWE performance assessment, including ERA5 reanalysis data, wind profile clustering, power estimation models, and farm modeling. Chapter 3 defines the objective and scope of the research and gives the research questions. Chapter 4 presents awesIO and AWESPA. Chapter 5 describes the wind profile clustering module and compares ERA5 data with Lidar measurements from the Bangor Erris test site. Chapter 6 presents the implemented power models, including the Luchsinger model and the inertia-free quasi-steady model based on the work of Van der Vlugt et al. and Schelbergen [5], [10], [11], [12]. Chapter 7 investigates asynchronous operation of wind-aligned AWE farms and derives a maximum phase offset for farm power smoothing while maintaining geometric separation between systems. Chapter 8 applies the workflow to the two case studies, and Chapter 9 gives the main conclusions.

2

Literature review

This chapter reviews the literature needed for the performance assessment of airborne wind energy (AWE) systems and farms. The use of ERA5 reanalysis data for AWE resource and performance assessment is first discussed in Section 2.1, with attention to the need for vertical wind-profile information. Wind-profile clustering is then reviewed in Section 2.2 as a method to reduce large atmospheric datasets to a limited number of representative profiles. Relevant power estimation models for pumping-cycle AWE systems are summarized in Section 2.3. Finally, Section 2.4 reviews AWE farm modeling, focusing on layout constraints, synchronization, phase shifting, and farm-level power smoothing.

2.1. ERA5 reanalysis data used for AWE performance assessment

ERA5 is the fifth-generation global atmospheric reanalysis product developed by the European Centre for Medium-Range Weather Forecasts (ECMWF) within the Copernicus Climate Change Service. A reanalysis combines a numerical weather prediction model with historical observations using data assimilation, producing a physically consistent reconstruction of the atmospheric state over long time periods. Compared to its predecessor ERA-Interim, ERA5 provides higher spatial and temporal resolution, with hourly output on a global grid of approximately $31\ \mathrm{km}$ and 137 vertical model levels extending up to $1\ \mathrm{Pa}$ [13].

The ERA5 data used in this thesis are obtained on pressure levels rather than directly in terms of geometric altitude. For AWES applications, these pressure levels must therefore be converted to corresponding altitudes, as described in Section 5.1.

ERA5 reanalysis data are widely used in wind resource assessment because they provide long-term, hourly, global atmospheric data over many vertical levels. For AWES applications, this is useful because the operating altitudes extend above conventional wind turbine hub heights and require a description of the vertical wind structure rather than only near-surface conditions [6], [14], [15]. Bechtle et al. [14] used ERA5 to compare the wind resource available to conventional wind turbines with the variable-height resource accessible to AWES. Their analysis showed that access to higher altitudes and variable harvesting height can increase the available wind resource and reduce its temporal variability compared with a fixed-height reference case [14].

Bechtle et al. [14] limited their study to wind resource availability and did not model the power harvesting process of an AWES. The optimal operating height was therefore approximated as the height with the maximum wind speed within the admissible range. This assumption is suitable for a first-order resource assessment, but it does not account for the fact that operation at higher altitude may reduce energy conversion efficiency and may not coincide with the power-optimal trajectory of a real AWES [14]. Thimm extended this resource-oriented framework within the AWERA toolchain [7] by coupling wind resource parametrization to power production estimates through the quasi-steady model (QSM) originally developed by van der Vlugt et al. [16]. This integration changed the approach from a wind resource description to a formulation that can be used for AWES performance assessment [6].

Although ERA5 has broad coverage and is easy to access, its limitations must also be considered. Gualtieri [17] showed that ERA5 can provide acceptable results for offshore and flat sites, whereas its accuracy decreases in mountainous, forested or otherwise complex terrain because of its limited spatial resolution. Caution is also required in coastal regions, where the sea–land discontinuity leads to strong contrasts in surface roughness and atmospheric stability, which can reduce the reliability of reanalysis data. ERA5 is therefore well suited for large-scale screening studies and pre-feasibility assessments, while local topography, land-use variability and coastal effects require more detailed data. [17]. This is relevant for AWES studies because vertical gradients and profile shape can affect the estimated system performance.

Recent work has also shown that the wind profiles relevant to AWES are often more complex than the standard monotonic profiles assumed in simplified engineering models. He et al. [15] investigated wind profiles up to $3000\ \mathrm{m}$ using ERA5 data together with lidar measurements and found that negative shear occurs frequently in instantaneous profiles, especially between $500\ \mathrm{m}$ and $1500\ \mathrm{m}$.

2.2. Wind profile clustering

Since AWES performance depends on the full vertical wind profile, long-term ERA5 datasets contain a large number of relevant atmospheric states. Directly evaluating every individual profile in an optimization or power-production model is computationally inefficient. Wind-profile clustering is therefore used to reduce the full dataset to a limited number of representative profile shapes and associated occurrence frequencies. Among the available clustering techniques, k-means clustering has become a common choice because it is computationally efficient and yields a discrete number of clusters that can be used directly in subsequent performance analyses [6], [18], [19]. In the AWERA framework, clustering is used to group similar normalized wind profiles so that the full wind resource can be represented by a compact set of profile shapes and occurrence frequencies [6]. This parametrization retains the main differences in vertical profile structure and can capture profile shapes that are not well described by standard logarithmic laws.

2.2.1. Principal component analysis

As part of the clustering procedure, principal component analysis (PCA) can first be used to project the wind-profile dataset onto a lower-dimensional space while retaining the dominant variability in the original data. In this transformation, the coordinate system is shifted to the mean of the dataset, after which the first principal component is aligned with the direction of maximum variance. Each following component is orthogonal to the preceding components and is chosen to capture the largest remaining share of the variance. Consequently, the later components represent progressively less of the variability in the dataset. The principal components themselves form unit vectors along the positive axes of the transformed coordinate system [19].

The amount of variance preserved by PCA depends on the number of principal components used. Schelbergen et al. found that, for the offshore wind-profile dataset considered in their study, the first four principal components already captured more than $90\ \%$ of the total variance. This strong dimensionality reduction is possible because wind speeds at neighboring vertical levels are highly correlated. Although retaining four components was already sufficient to preserve most of the variance, the use of five principal components was selected because the additional component still provided a small increase in retained variance. The normalized wind profiles were therefore mapped to the space spanned by the first five principal components before clustering was performed [19].

2.2.2. k-means clustering

K-means clustering is then used to identify a finite set of representative wind-profile shapes [20]. In this method, each cluster is described by a centroid, and every sample is assigned to the cluster whose centroid is closest in terms of Euclidean distance. The algorithm repeatedly updates the centroid locations in order to minimize the sum of the squared distances between samples and their corresponding centroids. This objective function is commonly referred to as the within-cluster sum of squares (WCSS). After convergence, the centroids in principal-component space can be transformed back to the original physical space, where they represent the mean wind-profile shapes of the clusters [19].

Although k-means is simple and always returns a clustering solution, its results should be interpreted with care. The method favors clusters that are approximately spherical and of comparable size, which may not correspond to the actual structure of wind-profile datasets. For this reason, the number of clusters, k, must be chosen carefully rather than assumed a priori. Schelbergen et al. evaluated the appropriate value of k using both the elbow method and the silhouette method. The elbow method examines how the WCSS decreases as the number of clusters increases. In addition, the silhouette score assesses cluster quality by comparing the similarity of each sample to other samples within the same cluster with its similarity to samples in the nearest neighboring cluster. The highest mean silhouette score in their study was obtained for $k = 2$, indicating that two clusters produced the most cohesive separation of the dataset. The decreasing silhouette score for larger values of k suggested that using too many clusters reduces cluster cohesiveness, even though it improves the WCSS [19].

By associating each cluster centroid with a frequency of occurrence, clustered wind profiles can be used as representative operating conditions in optimization or power-curve calculations [6], [18]. For this thesis, wind-profile clustering provides a way to translate large ERA5 datasets into a limited number of characteristic wind conditions that remain representative of the vertical variability relevant to AWES operation.

2.3. Power estimation

AWE power models reported in the literature span a broad range of complexity, from highly idealised cycle-averaged formulations to more detailed quasi-steady and dynamic simulation models.

In this section, the definitions of steady, quasi-steady, and dynamic models are briefly discussed, followed by an overview of the most relevant power models for pumping-cycle AWE systems.

2.3.1. Definitions of steady, quasi-steady, inertia-free, and dynamic motion

The terms *steady*, *quasi-steady*, *inertia-free*, and *dynamic* are used throughout the airborne wind energy literature to describe different levels of modelling fidelity, but their meaning is not always stated consistently. Following the terminology discussed by Van Deursen et al. [21], these concepts are defined explicitly here.

A *steady* flight state refers to a situation in which the kite motion, wind conditions, and gravity are constant in an appropriate body-fixed frame. In such a state, the aerodynamic and tether forces remain constant, and no change in the motion variables occurs with time. For pumping AWE systems, this is a highly restrictive condition and is generally not representative of crosswind operation. [21]

An *inertia-free* model assumes that inertial forces are negligible compared with the external forces acting on the system. Under this assumption, the kite is taken to adapt instantaneously to changes in the applied forces, so that the equations of motion reduce to a direct force equilibrium. The resulting behavior is therefore governed only by the balance between aerodynamic, tether, and gravitational forces, without accounting for the effects of mass and acceleration. [21]

A *quasi-steady* model represents the motion as a sequence of equilibrium states, assuming that perturbations from the quasi-steady equilibrium decay sufficiently fast for the instantaneous motion to remain approximated by equilibrium relations. In this formulation, the equilibrium state is governed by constant tangential and radial speed components, and the kite is assumed to follow a curved trajectory at constant speed subject to a constant reel-out speed. Unlike an inertia-free description, the quasi-steady assumption retains selected inertial effects associated with curved flight, including centrifugal and Coriolis forces, as well as part of the Euler forces. As a result, the dynamic equations of motion are reduced to a three-degree-of-freedom system of first-order differential equations, in which the inertial force becomes time-invariant. [21]

A *dynamic* model, by contrast, retains the full equations of motion and explicitly accounts for inertial effects arising from acceleration. These include relative, centrifugal, Coriolis, and Euler accelerations, depending on the chosen reference frame. [21]

2.3.2. Overview of relevant power models

In this work, only inertia-free quasi-static models are used because of their simplicity and low computational cost. Figure 2.1 shows the development of the main theories, equations, and models for pumping-cycle crosswind airborne wind energy systems relevant to this work. The arrows indicate that a given work uses, extends, or follows from an earlier theory, equation, or model. Although some of these relationships are not always strictly defined, the figure provides a useful overview of the main development paths in the literature.

The starting point is the work of Loyd [22], who introduced the fundamental crosswind power theory for tethered wings. Almost thirty years later, Argatov et al. [23] used this theory to develop an analytical model for estimating the mechanical energy output of a kite pumping cycle.

Fechner and Schmehl [24] introduced a quasi-static system model for a pumping-cycle AWE system. In this model, the reel-out phase is represented by a single operating state, while the reel-in phase is calculated iteratively. Crosswind effects are included through quasi-static force relations, but the detailed reel-out trajectory is not resolved explicitly. The kite is modeled as massless and represented by aerodynamic lift and drag forces. Although the effect of mass was investigated, it was found to be small and was therefore neglected. The model also includes optimization of the reel-out speed, average height, and elevation angle.

Noom and Schmehl [3], [25] continued from the equations of Argatov et al. [23] and derived the full formulation for a dynamic model and developed a quasi-steady model for calculating the traction power of a tethered kite. In contrast to the earlier analytical formulations, this model included the effect of gravity.

Van der Vlugt et al. [16] built further on the work of Noom and Schmehl and developed a more detailed inertia-free quasi-steady model. In this model, all inertial forces are neglected and each operating state is obtained from a force equilibrium. The kite and tether are represented in lumped form and effects such as tether sag and tether drag are included. The retraction and transition phases are modeled through simplified trajectory updates, whereas the traction phase is represented by a sequence of representative crosswind states rather than by explicitly resolving the full figure-of-eight motion. This results in a more detailed description of the pumping cycle than the earlier analytical models, while remaining computationally efficient enough for optimization and scaling studies. Venegas [26] used a related approach, but iteratively calculated the lift-to-drag ratio in a slightly different way than Van der Vlugt et al. [16]. The model of Van der Vlugt et al. [16] was also later used by Johnson [27] to study airborne wind energy farms.

A separate branch in the development is the work of Ranneberg et al. [28], who proposed an inertia-free quasi-steady model aimed at fast optimization and power-curve estimation. In this approach, much of the trajectory is prescribed rather than solved directly. The traction phase is represented by points on a prescribed figure-of-eight surface, while the retraction phase is simplified to a single representative flight state with a fixed retraction slope typical of fixed-wing kites. At each point, optimization is used to determine the control and operating variables that satisfy force equilibrium and maximize the mean cycle power. This makes the model suitable for rapid performance estimation, although the prescribed trajectory limits its use for detailed flight-dynamics studies.

Schelbergen [5], [12] further developed the quasi-steady modeling framework for pumping soft-kite systems, building on the model of Van der Vlugt et al. [16]. His work placed particular emphasis on performance prediction and long-term energy estimation, and the model was validated against data from 87 consecutive pumping cycles [12]. In this framework, the traction phase can be represented in two ways. In the simpler line representation, the crosswind motion is replaced by a representative flight state at fixed angles, similar to the model of Van der Vlugt et al. [16]. In the more advanced surface representation, multiple steady states are evaluated on a prescribed figure-of-eight surface, which enables a more direct comparison with flight data and generally improves the power estimate. This surface representation was introduced by Talmar [29] as part of further enhancements to the quasi-steady model. This model was then used by Thimm [6] in the AWERA toolchain, which combines wind resource parameterisation with AWE power harvesting tools. More recently, Bertozzi et al. [30] used this model with a new optimizer and the surface representation, and validated it against the TU Delft V3 reference kite.

Cayon et al. [31] and Van Deursen et al. [21] further investigated the relation between quasi-steady and dynamic modeling for airborne wind energy systems. Van Deursen et al. derived dynamic equations of motion and compared steady, quasi-steady, and dynamic modeling approaches to identify when

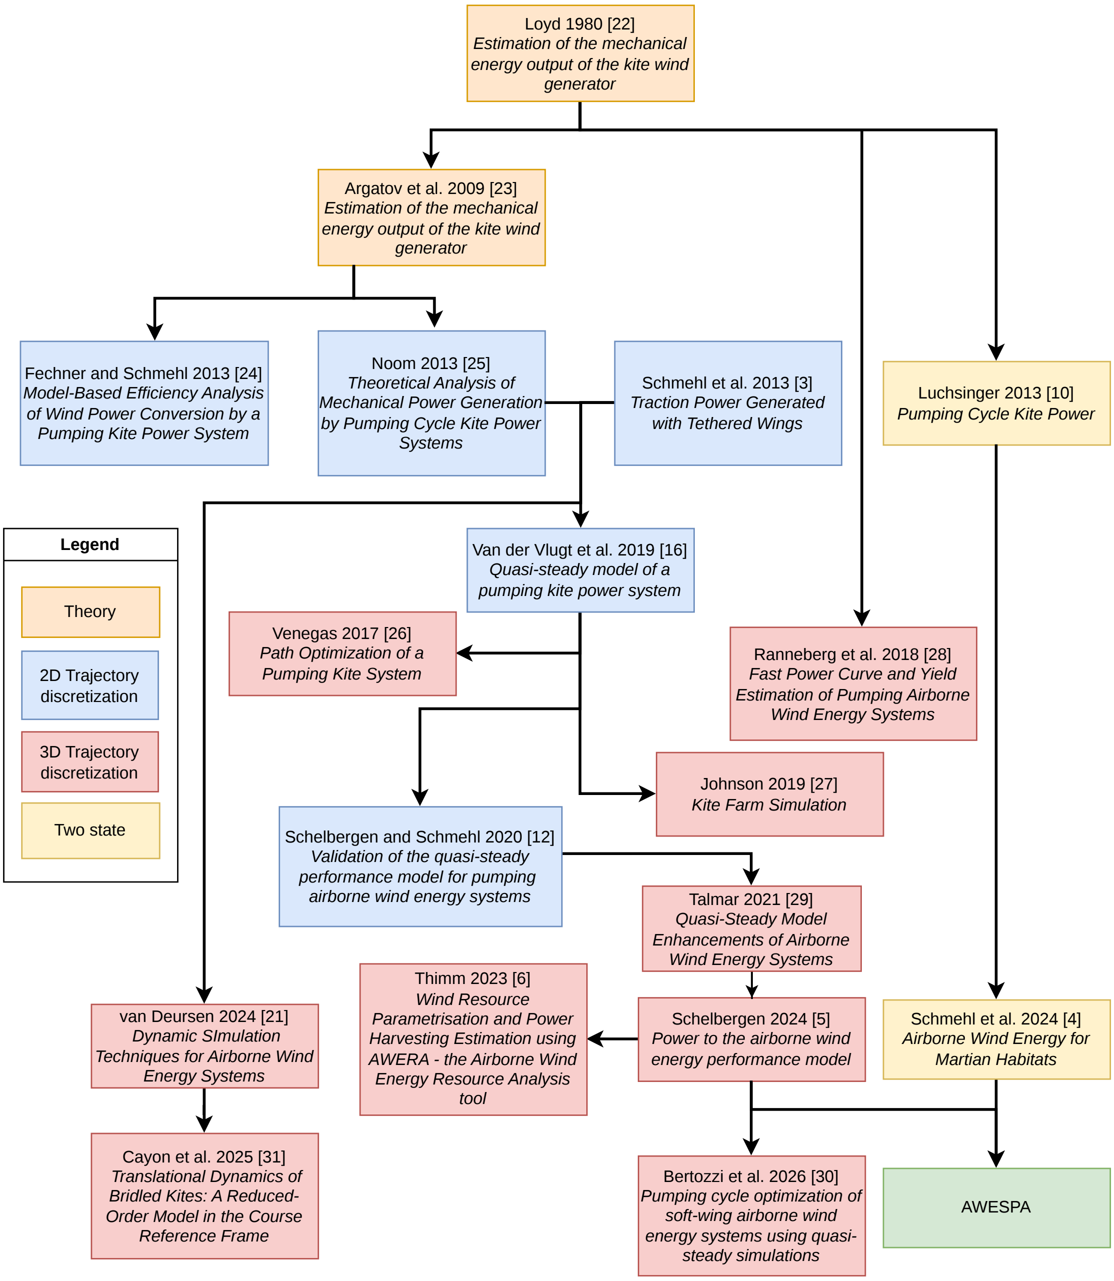

Figure 2.1: Evolution of power estimation models and theory for pumping-cycle airborne wind energy systems. The arrows indicate that a given work uses, extends, or follows from an earlier theory, equation, or model. Orange boxes denote works that mainly describe theory without introducing a full simulation model. Blue boxes indicate models that discretize the kite path in two dimensions, while red boxes indicate models that discretize the kite path in three dimensions. Yellow boxes denote models that use one representative state for each reel-out and reel-in phase. The green box indicates where the present work, AWESPA, fits within this development.

the quasi-steady assumption remains valid. Their work shows that quasi-steady models can predict time-averaged quantities, such as mean traction power, with reasonable accuracy. However, by neglecting inertial effects, quasi-steady models cannot accurately capture phase shifts, amplitudes, and acceleration-dependent behavior [21]. Cayon et al. extended this discussion for bridled soft kites by developing a reduced-order translational model in a course reference frame. Their formulation retains gravitational and inertial terms from the curvilinear reference frame, while applying a quasi-steady condition with zero path-aligned acceleration. The results show that for low wing loadings typical of soft kites, the quasi-steady approximation gives mean reel-out power errors below 1%, whereas for higher wing loadings inertial effects become more important and can lead to power deviations of up to 14% [31]. Together, these studies clarify the validity limits of quasi-steady modeling and show when dynamic effects must be retained.

Returning to the branch starting from Loyd's crosswind theory, Luchsinger [10] developed a highly simplified analytical model for pumping-cycle kite power. This model extends Loyd's crosswind power analysis [22] by including both the traction and retraction phases. The reel-out and reel-in phases are each represented by an idealized quasi-steady flight state, and compact expressions are derived for the average cycle power using dimensionless reeling factors and aerodynamic force factors. The simplicity of the model makes it attractive for first-order analysis and conceptual design. However, the model neglects tether drag and gravity, assumes constant aerodynamic coefficients within each phase, and does not represent the detailed kite trajectory. These assumptions limit its ability to describe realistic flight behavior. The model also introduces a three-phase operating strategy, in which the wind-speed range is divided into unconstrained, tether-force-limited, and power-limited operation.

Schmehl et al. [4] later adapted the three-phase Luchsinger approach for soft-kite systems by having the reel-in phase use a fixed lift over drag ratio instead of a fixed elevation angle used in Luchsinger [10]. The tether drag is included, aerodynamic force factors are different, and the reeling factors are optimized in a similar manner to the original Luchsinger model [10]. This provides a simple analytical model that remains close to the original three-phase strategy, while making the formulation more suitable for soft-kite applications.

2.4. Farm modeling

The farm modeling of airborne wind energy (AWE) systems is different from that of conventional wind farms. In conventional wind farms, the spacing between turbines is mainly governed by wake losses. For AWE farms, the layout problem is mainly governed by the safe separation of kites, tethers, and ground stations during cyclic operation [32]. A compact AWE farm requires small spacing between systems, but this spacing is limited by the flight envelopes and by the retraction trajectories of neighboring systems.

For ground-generation pumping systems, the spatial layout problem is coupled to the time variation of the power output. During the traction phase, the tether is reeled out and power is produced. During the retraction phase, the tether is reeled in and power is consumed. If several systems operate in phase, the farm power output has large peaks during simultaneous traction and low values during simultaneous retraction. A phase shift between systems can reduce this variation because some systems produce power while others consume power. However, for wind-aligned systems with overlapping or closely spaced operational zones, phase shifting can also increase the risk of kite or tether interference. The farm-level problem is therefore to combine high system density with smooth power output and collision-safe operation.

Faggiani and Schmehl [33], [34] studied the layout and operation of pumping kite wind parks using an optimized single-system model as the building block for a farm model. In their approach, the spacing between systems is determined by safe collective operation rather than by wake interaction. The two basic spacing directions are shown in Figure 2.2. For systems placed perpendicular to the wind direction, the minimum lateral spacing can be written as

$$d = 2r_{\max}\sin(\phi), \tag{2.1}$$

where d is the distance between ground stations, $r_{\max}$ is the maximum tether length, and ϕ is half of the operational envelope in the azimuth direction.

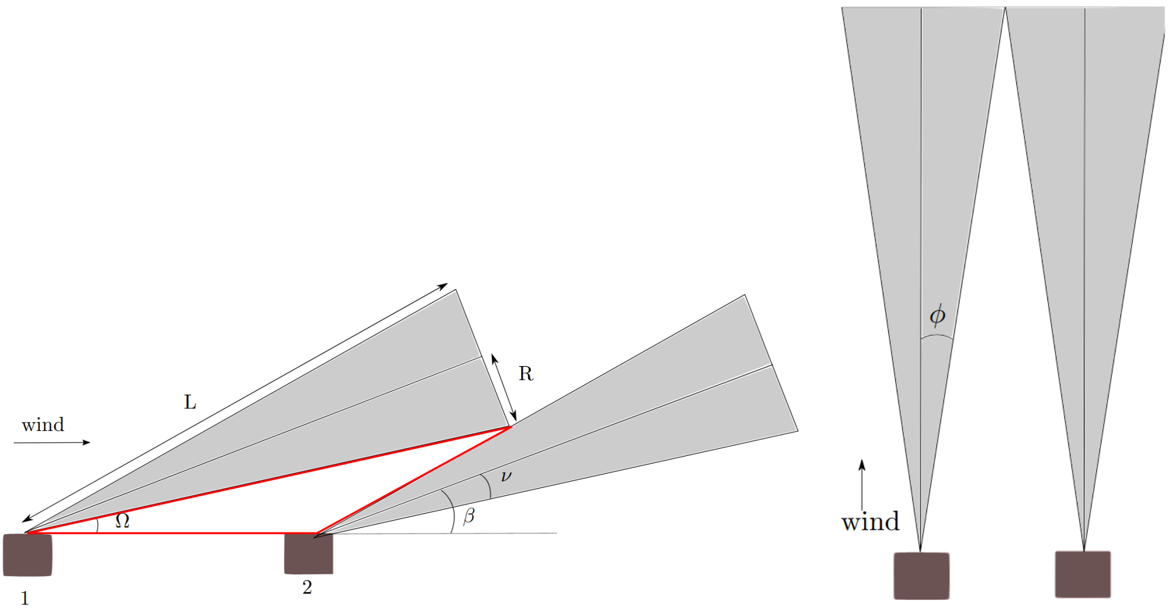

(a) Wind-aligned layout [33], [34]

(b) Wind-perpendicular layout [33]

Figure 2.2: Wind-aligned and wind-perpendicular spacing between two pumping kite systems [33], [34].

The wind-aligned spacing is more restrictive because the downwind system can operate below or behind the upwind system. Faggiani and Schmehl derived the minimum longitudinal spacing from the side-view geometry of the operational envelopes. In the notation used here, this spacing is

$$d = \frac{r_{\max}}{\sin(\beta - \nu_1)\left(\frac{1}{\tan(\beta-\nu_1)} + \frac{1}{\tan(\nu_1+\nu_2)}\right)}, \tag{2.2}$$

where β is the average elevation angle during traction, ν_1 is half of the elevation-angle range of the upwind system, and ν_2 is half of the elevation-angle range of the downwind system. If both systems use the same operational envelope, ν_1 and ν_2 are equal. The spacing in Equation 2.2 is minimized when the wind-aligned systems operate synchronously. For pumping operation, this means that systems in the same wind-aligned column are kept in phase. [33], [34]

This synchronization requirement is also linked to the farm power output. Synchronous operation allows close wind-aligned spacing, but it causes the power cycles of the systems in the same column to overlap. Faggiani and Schmehl [33], [34] therefore proposed that systems within a wind-aligned column operate synchronously, while different columns operate with a phase shift. The column-based grouping depends on the wind direction, as shown in Figure 2.3. When the inflow is aligned with the farm array, the columns have similar size. For diagonal inflow, the number of systems per column becomes uneven, which changes the aggregate power waveform. [33], [34]

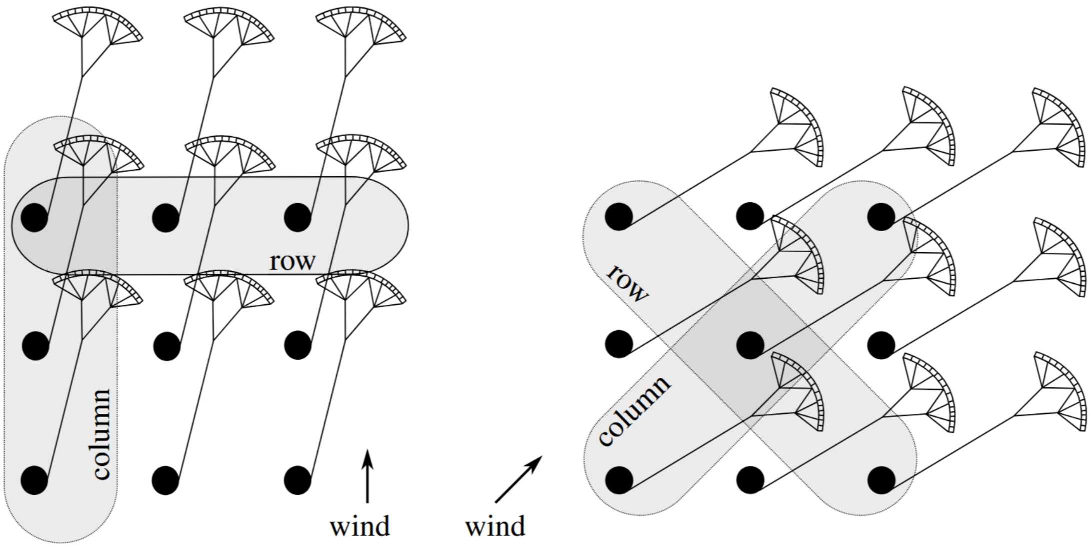

Figure 2.3: Farm layout for inflow aligned with the array and inflow diagonal to the array [33].

The effect of this collective operation on the instantaneous farm power is shown in Figure 2.4. The single-system power signal has a cyclic production and consumption pattern. For a larger farm, the aggregate output is smoother because the systems in different columns can be shifted in time. This reduces the need for temporary storage for retraction power, but it does not remove the synchronization constraint within each wind-aligned column.

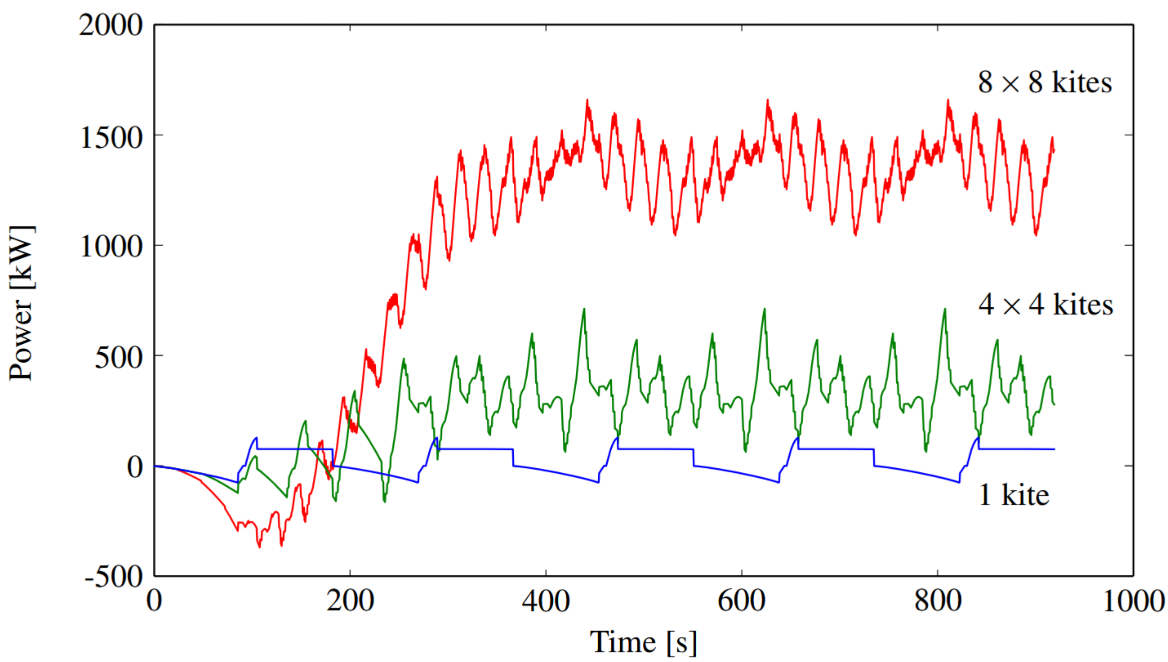

Figure 2.4: Instantaneous electricity generation for one system and for two kite wind farm sizes [34].

Johnson [27] extended this type of farm analysis by replacing the simplified two-dimensional flight path with a three-dimensional path model. The traction phase was modeled using Lissajous curves, so that the figure-eight motion was included in the position of the kite. Johnson [27] used this model to compare wind-aligned and wind-perpendicular layouts, with and without phase shifts. The results show the same trade-off as in Faggiani and Schmehl [33], [34], but with a more detailed geometric description of the flight path. Phase-shifted operation gives a more level farm power output because the systems are not all in traction or retraction at the same time. However, the phase shift can increase the required ground-station spacing, especially in the wind-aligned case. Systems without phase shift can therefore reach higher spatial density, while phase-shifted systems reach more stable farm power output.

Licitra et al. [35] studied trajectory optimization for a rigid-wing pumping system and compared circular and lemniscate trajectories under equal boundary conditions. The average power output was found to be nearly independent of the trajectory topology in their case study. However, the circular trajectory occupied a smaller lateral width than the lemniscate trajectory. This result is relevant for farm modeling because two trajectories with similar single-system power can still lead to different farm densities. The layout comparison is shown in Figure 2.5. In the illustrated case, circular trajectories require less area for the same maximum tether length, which can increase the surface power density of an AWE farm.

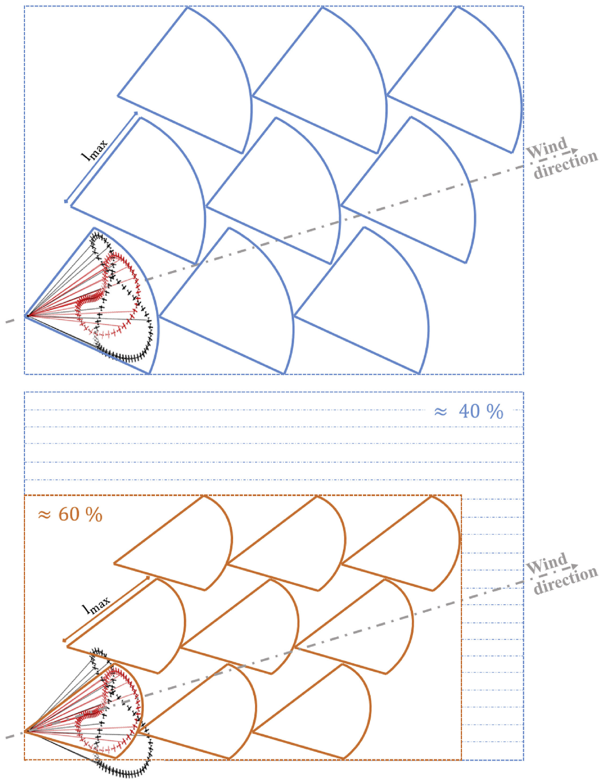

Figure 2.5: AWE farm layout comparison for lemniscate trajectories and circular trajectories [35].

Da Costa et al. [32] studied the combined problem of AWE farm layout and farm power smoothing. Their work addresses the same constraint that appears in the studies by Faggiani and Schmehl and Johnson: when systems are placed close together in the wind-aligned direction, the retraction path of one system can intersect the operating region of another system. With the usual reel-in strategy, these wind-aligned systems therefore have to be synchronized in groups to avoid collision risks. This grouped synchronization limits the possible phase shifts between systems and can lead to large variations in farm power output when many systems enter traction or retraction at the same time.

To avoid this synchronization requirement, da Costa et al. proposed a modified reel-in strategy in which the kite retracts inside the conical-shaped traction envelope, as shown in Figure 2.6. In this approach, the retraction motion remains inside the volume already reserved for the system during traction. As a result, neighboring systems can be placed with disjoint traction volumes and can still operate out of sync. This allows individual phase shifting between systems, which reduces the farm-level power variation without requiring wind-aligned systems to follow the same cycle timing. The drawback is that reel-in takes place at a lower elevation angle than in the usual strategy, which increases the power consumed during retraction. Da Costa et al. found this increase to be small compared with the gain in power smoothing. [32]

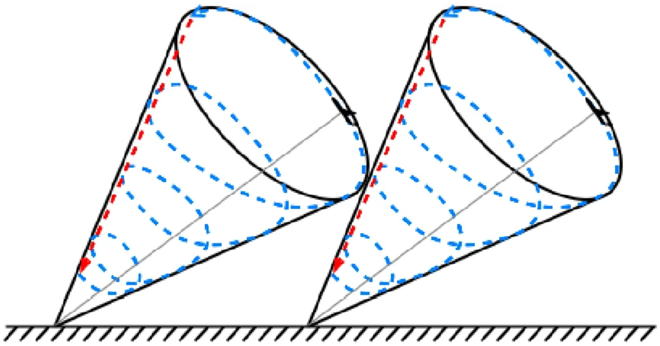

Figure 2.6: Reel-in strategy proposed by da Costa et al. [32], where the kite retracts inside the conical-shaped flight envelope.

These studies all treat synchronization as a constraint for closely spaced wind-aligned systems when the operational zones or retraction paths can interfere. Faggiani and Schmehl require systems in the same wind-aligned column to operate nearly in phase to allow close spacing. Johnson also shows that introducing phase shifts can increase the required spacing when the full three-dimensional flight path is considered. Da Costa et al. avoid this requirement by changing the reel-in strategy so that retraction remains inside the traction envelope. However, this solution changes the operating cycle itself. The remaining question is therefore whether wind-aligned systems with overlapping or adjacent operational zones must always be synchronized, or whether asynchronous operation can still be possible for the original type of operation under specific geometric or timing conditions.

3
Objective and scope of the research

This chapter defines the overall objective of the thesis and formulates the research questions that guide the development and application of the proposed toolchain. The questions are structured according to the main components of the research: the development of the AWESPA and awesIO framework, wind resource analysis, the comparison and extension of power models and the farm module.

3.1. Research objective

The research objective of this thesis is formulated as follows:

> To enable consistent, reusable, and comparable performance assessment of airborne wind energy systems by developing a modular toolchain with standardized input and output definitions and interchangeable modeling modules.

3.2. Research questions

To achieve the research objective, several main research questions are formulated.

3.2.1. AWESPA and awesIO development

The first part of this research focuses on the development of the AWESPA toolchain and the standardized input/output framework awesIO. The corresponding main research question is:

> How can a modular toolchain be developed to assess the performance of airborne wind energy systems in a consistent, reusable, and comparable manner?

To answer this main question, the following sub-questions are addressed:

1. How should the toolchain be structured such that individual modules can be exchanged easily without disrupting the overall workflow?
2. How should input and output data be standardized to ensure consistent communication between the different modules in the toolchain?
3. How can validation procedures be incorporated to improve the reliability and robustness of the toolchain and the standard data framework?

3.2.2. Wind resource

The second part of the research focuses on the processing of wind resource data and the comparison between ERA5 reanalysis data and Lidar measurements at the Kitepower test site in Bangor Erris. The corresponding main research question is:

> How can wind resource data be processed into representative wind profile shapes, and how well does ERA5 represent the local wind conditions compared with Lidar measurements?

To answer this main question, the following sub-questions are addressed:

1. How does ERA5 compare with Lidar measurements at the Bangor Erris test site in terms of mean wind speed, wind-speed variability, Weibull distribution, wind direction, and vertical profile shape?
2. To what extent does ERA5 smooth local wind-speed and profile-shape variability compared with Lidar measurements?

3.2.3. Power models

The third part of the research considers the influence of different power simulation models on the predicted performance of an airborne wind energy system. The main research question is:

> How do different power simulation models for airborne wind energy systems influence the predicted performance of a given system design?

To answer this main question, the following sub-questions are addressed:

1. How does the original Luchsinger model for fixed wing kites compare with the soft kite variant of the Luchsinger model?
2. How does the power estimation of soft kite variant Luchsinger model compare to the inertia-free quasi-steady model?
3. How can the inertia-free quasi-steady model be extended to include a non-constant reel-out elevation angle?

3.2.4. Farm model

The fourth part of the research addresses the farm-level operation of cyclic airborne wind energy systems. The focus is on wind-aligned systems with overlapping reel-in and reel-out envelopes, where previous studies generally require synchronized operation to avoid collision risks. This part investigates whether this synchronization requirement can be relaxed by deriving a geometric maximum phase offset and using it as a constraint in the farm phase optimization. The main research question is:

> To what extent can closely spaced wind-aligned airborne wind energy systems operate asynchronously while avoiding collision, and how does this affect the peak-to-peak variation of the aggregated farm power output?

To answer this main question, the following sub-questions are addressed:

1. How can a maximum allowable phase offset be derived for wind-aligned systems with adjacent reel-in and reel-out envelopes?
2. How does constrained asynchronous operation affect the peak-to-peak farm power variation compared with synchronized operation?
3. How do wind-aligned and wind-perpendicular layouts compare for two-system and 10-system farm cases?

3.2.5. Case studies

The last part of the research applies the wind resource, power, and farm analyses to two site-specific case studies. The first case study considers a $100\ \mathrm{kW}$ isolated airborne wind energy system at the Kitepower test site in Bangor Erris, Ireland, and investigates the influence of kite size and wind-profile representation on the predicted annual energy production. The second case study considers a $1\ \mathrm{MW}$ airborne wind energy farm consisting of 10 systems at the Tweede Maasvlakte in the Port of Rotterdam, with focus on the expected farm performance and aggregated power-output variation. The main research question is:

> How do wind resource representation, system sizing, and farm operation influence the predicted performance of the selected airborne wind energy case studies?

To answer this main question, the following sub-questions are addressed:

1. How do system parameters, such as kite size, maximum tether force, and maximum generator power, affect the power curve, annual energy production, and capacity factor of the (100 kW)

system?

2. How does including multiple representative wind profiles affect the predicted annual energy production?
3. What performance can be expected from the $100\ \mathrm{kW}$ isolated system at Bangor Erris and the $1\ \mathrm{MW}$ farm at the Tweede Maasvlakte?
4. How does farm-level operation affect the power-output variation of the $1\ \mathrm{MW}$ airborne wind energy farm?

4

AWESPA and awesIO: a standardized and modular framework for airborne wind energy assessment

Performance assessment for airborne wind energy systems involves wind resource modeling, power modeling, farm modeling, and economic analysis. This requires several modeling tools. These tools are often developed separately and use custom input and output formats. This makes it difficult to combine tools in one analysis, reproduce results, and compare results between models.

In conventional wind energy, similar issues have been addressed through standardized data interfaces, such as windIO, developed within IEA Wind Task 37 [36]. These interfaces provide a common structure for data exchange between tools and support modular modeling workflows. However, windIO was introduced after many organizations and companies had already developed internal data formats. This has limited the adoption of a single common standard.

The airborne wind energy sector is at an earlier stage of development. This creates an opportunity to define a common data format before many incompatible formats become established. A shared data format can reduce differences between tool interfaces and support data exchange between tools and organizations.

This chapter presents the framework developed in this thesis. The framework consists of two parts: awesIO, a standardized data exchange format, and AWESPA, a modular Python toolchain that connects the different modeling steps. Together, awesIO and AWESPA provide a reproducible approach for airborne wind energy performance assessment, from wind resource characterization to annual energy production (AEP) estimation.

4.1. awesIO as a standardized data framework

awesIO (Airborne Wind Energy Systems Input/Output) [9] is a standardized and validated data framework for airborne wind energy applications. Its main purpose is to provide a tool-independent format for system configurations, environmental conditions, and other simulation inputs and outputs.

The framework uses human-readable YAML files that follow predefined schemas. These schemas define the structure, naming conventions, and required fields for each file type. This allows different tools to interpret the same file in a consistent way. By using a common data format, awesIO allows independently developed models to exchange data without tool-specific manual conversion.

A main feature of awesIO is its built-in validation capability. Before a file is used in a simulation, or after a file has been created, it can be checked for:

- compliance with the defined schema,

- presence of all required parameters,
- consistent naming conventions.

This validation step checks whether files comply with the awesIO standard. It allows researchers to verify whether their configuration files or simulation outputs follow the expected structure. This reduces ambiguity in data exchange and helps prevent errors caused by incomplete or inconsistent input files.
awesIO is hosted in a public GitHub repository: `https://github.com/awegroup/awesIO`. This makes the schemas, documentation, and examples accessible to other researchers and developers. It also allows the framework to be developed collaboratively, for example through issues, discussions, and contributions to the repository. In addition, the validation functionality can be installed as a Python package using `pip`.

Figure 4.1: awesIO logo

The public repository and validation package make awesIO usable outside AWESPA. This is relevant because awesIO is intended as a data standard that can also be used by other tools, research groups, and companies.

4.2. awesIO scope and structure

To improve consistency across tools and institutions, all quantities in awesIO are expressed using standard SI units.

At the time of writing, awesIO contains three schemas. These schemas were selected because they cover the main files required for the analyses performed in this thesis. The system configuration schema is the central schema. It defines the physical properties of the complete airborne wind energy system. This file contains parameters such as mass, wing area, force limits, material densities, and other system-level properties. It does not define the aerodynamic properties of the airborne system.

Aerodynamic properties are excluded from the system configuration file because they are tool-dependent. Some models require only a small number of aerodynamic coefficients for a limited set of operating states. Other tools use higher-fidelity aerodynamic models to determine the aerodynamic behavior of the kite or aircraft. Keeping aerodynamic properties outside the system configuration file avoids imposing one aerodynamic model structure on all tools.

The system configuration file describes the complete AWE system in one file, from the ground station to the tether and airborne system. This structure is shown in Figure 4.2. The components in the system configuration file are the kite, tether, and ground station. These components are defined at the same hierarchy level, which reflects the modular structure of the system.

AWE systems can differ in physical layout and operating principle. For example, a fly-generation system may require a conductive tether, while a ground-generation system uses a non-conductive tether. The airborne component may also differ between systems. It can be a leading-edge inflatable soft kite, a ram-air kite, a single-skin kite, or a fixed-wing aircraft. A single schema therefore needs to represent several system types.

To address this, each component specifies its component type. This type determines which fields are required or applicable for that component. For example, a conductive tether requires different parameters from a non-conductive tether, and a soft kite requires different properties from a fixed-wing aircraft. This type-based structure allows awesIO to represent different AWE system concepts while preserving a common file structure.

Where possible, variables are kept consistent between related component types. For example, different soft-kite types, such as ram-air, single-skin, and leading-edge inflatable kites, share common parameters where appropriate. The same approach is used for different fixed-wing configurations. This reduces unnecessary differences between similar systems and supports comparison between designs.

In addition to the component type, each component includes a name and version identifier. These metadata fields improve traceability by making clear which component definition is used. Equivalent

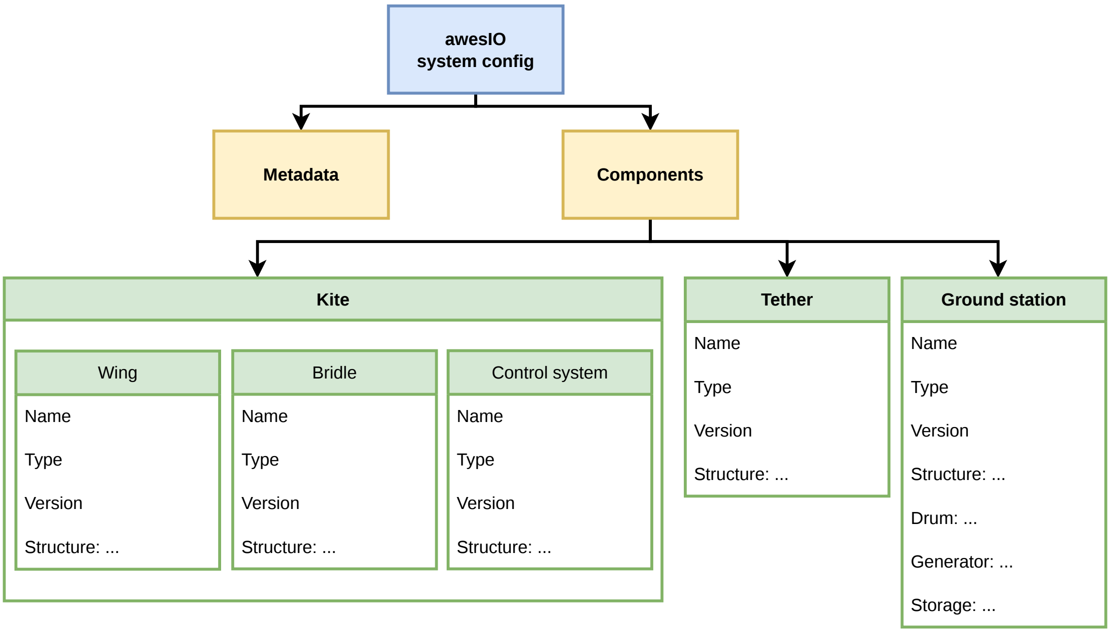

Figure 4.2: Structure of the awesIO system configuration file

metadata are also included at the top level of the file, so that the full system configuration can be identified and versioned consistently. Each component also includes a connector point identifier. These identifiers are used in the connectivity matrix at the top level of the file to define which components are connected to each other. This makes it possible to describe systems with more complex physical layouts, such as systems with multiple tethers, multiple kites, or multiple ground stations.

The other two schemas in awesIO define the power curve file and the wind resource file. The wind resource file contains the site-specific wind conditions, and the power curve file contains the power curves of the AWES. Standardizing these files, in addition to the system configuration file, allows the main inputs and outputs of a performance assessment to be exchanged between tools in a consistent way. The benefits of this are discussed further in Section 4.6.

4.3. AWESPA as an example implementation

AWESPA (Airborne Wind Energy System Performance Assessment) [8] is the Python toolchain developed in this thesis to use awesIO files in a performance assessment process. It currently connects wind resource processing, power estimation, and AEP calculation through a modular structure. AWESPA does not define new physical models. It provides wrappers that allow existing models to be used through a common interface.

AWESPA can be installed as a Python package. The underlying model packages are installed as dependencies, so the user can access the connected tools from one Python environment. This reduces the setup effort needed to run a performance assessment.

4.4. Architecture of AWESPA

The architecture of AWESPA is shown in Figure 4.3. The current implementation is organized around separate modules that each perform one part of the performance assessment.

- **Wind resource assessment** Raw wind data is processed to obtain representative wind profiles and associated probability distributions. The output is a wind resource file in awesIO format.
- **Power estimation** A selected power model computes power curves using the system configuration and wind resource information. The output is a power curve file in awesIO format.

- **Annual energy production estimation** The power curve file is combined with the wind resource file to compute AEP.

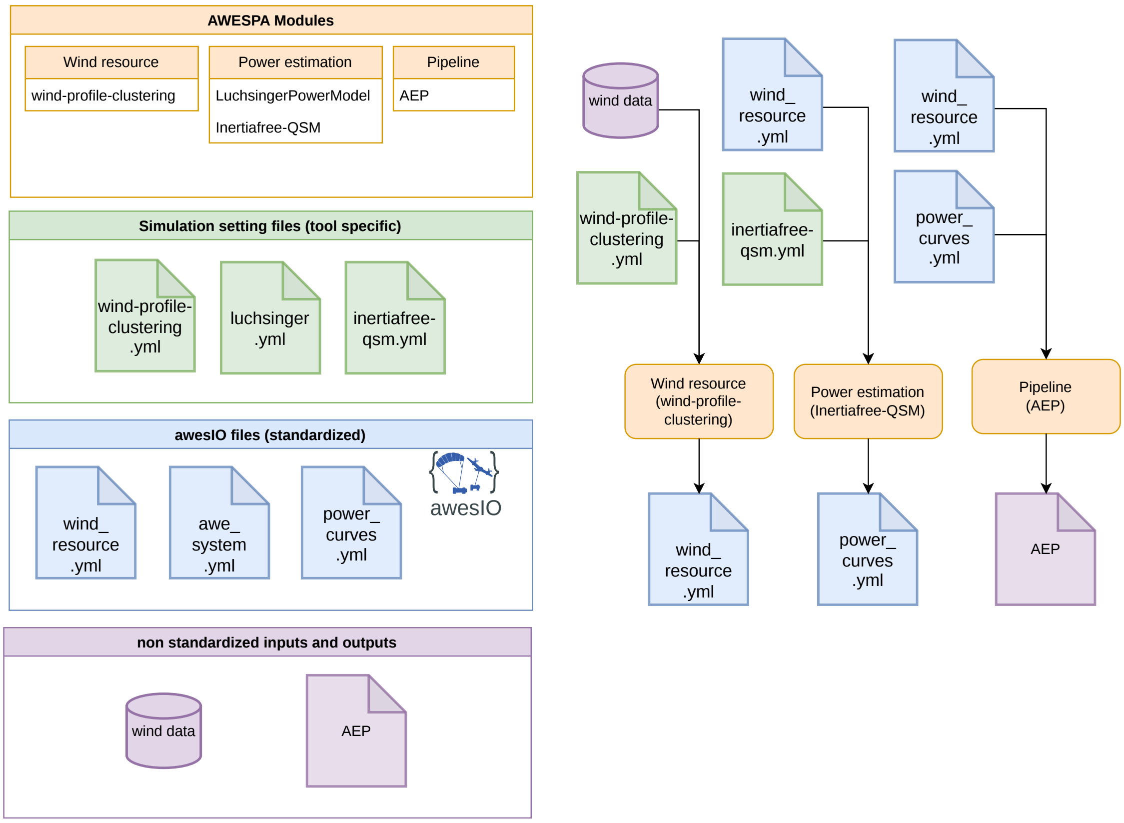

Figure 4.3: Architecture of AWESPA. The orange blocks represent tool-specific wrappers, the green files represent tool-specific settings files, the blue files represent standardized awesIO files, and the purple files represent tool inputs or outputs that are not currently standardized in awesIO. The right side shows example inputs and outputs for each AWESPA module.

AWESPA uses a wrapper-based architecture. Each external tool is connected to AWESPA through a wrapper. The wrapper is responsible for reading the required settings, preparing the inputs, executing the tool, and writing the outputs. The wrapper also converts awesIO inputs to the tool-specific input format and converts the outputs back to awesIO format when the tool does not use awesIO files directly. This means that the user can interact with each module through AWESPA instead of calling each underlying code repository directly.

Simulation settings are provided through YAML configuration files. These settings files are shown in green in Figure 4.3. They remain specific to the tool being used, because different tools can require different numerical settings, solver options, or model parameters. In contrast, the main files exchanged between modules follow the awesIO standard and are shown in blue. This distinction allows tools to keep their own settings while still exchanging common data through standardized files.

The interface of each module type is defined using Python abstract base classes (ABCs). For example, all power estimation models must implement the methods defined by the `PowerEstimationModel` interface. This gives all models within the same module category the same external behavior from the perspective of AWESPA. As a result, a different power model can be selected without changing the rest of the process, provided that the model wrapper follows the same interface.

Some tools can also require inputs or produce outputs that are not currently part of the awesIO standard. These files are shown in purple in Figure 4.3. They are still handled by the relevant wrapper, but they are not used as standardized exchange files between AWESPA modules. This allows AWESPA to work with existing tools while awesIO remains focused on the main files needed for performance assessment.

Reproducibility is supported through Python package management. AWESPA is installed as a Python package, and the external model packages are installed as dependencies. The versions of these dependencies are fixed in the `pyproject.toml` file using commit hashes. This makes it possible to repeat an analysis with the same versions of the underlying tools.

4.5. Extensibility and integration of external tools

AWESPA is designed so that new tools can be added. A new tool is integrated by implementing a wrapper that follows the AWESPA interface conventions.

To be compatible with the performance assessment process, a developer must write a wrapper for their tool that:

- accepts inputs defined through YAML configuration files,
- reads and writes data in awesIO format, or provides conversion functions,
- implements the appropriate abstract base class interface.

This structure also allows tools written in other programming languages to be used in AWESPA. For example, models implemented in Julia or MATLAB can be integrated through a Python wrapper that handles data exchange and execution.

4.6. Benefits of the unified framework

The combined use of awesIO and AWESPA provides several practical benefits:

- **Reproducibility:** the system description, wind resource, power curves, settings files, and tool versions are stored explicitly, which makes analyses easier to repeat.
- **Modularity:** individual modules can be replaced without redefining the data interface around them.
- **Comparability:** different modeling approaches can be evaluated using the same system configuration, wind resource, and AEP calculation.
- **Reduced integration effort:** tool-specific data conversion is handled by wrappers, reducing the amount of manual conversion required by the user.
- **Collaboration:** a shared data format allows institutions to exchange input and output files more consistently.

These benefits are mainly related to how models are connected and compared. Since the main exchange files follow the same structure, a module can be replaced without changing the surrounding code. For example, two power estimation models can use the same system configuration and wind resource file, after which power curve files are generated. These power curve files can then be easily compared directly or via the AEP calculation functionality of AWESPA.

The use of standardized files is also useful for case studies in which only one part of the analysis is changed. For example, the system configuration can be kept fixed while changing the wind resource, or the wind resource can be kept fixed while changing the power estimation model. This supports controlled comparisons between sites, systems, and modeling assumptions.

4.7. Future development

The framework is still under development. Future work can extend AWESPA with additional module types such as farm modeling, economic analysis, and life cycle assessment. These extensions would require either new awesIO schemas or extensions of the existing schemas, depending on the data that must be exchanged between tools.

Further development of awesIO should focus on more detailed system configuration descriptions, so that higher-fidelity tools can use the same configuration files. Other AWE system types, such as fixed-wing aircraft using fly-generation, should also be integrated into the schema by the research groups or companies working on these systems.

The main direction for future work is therefore to keep the framework modular while expanding the number of supported tools and file types. This would allow awesIO and AWESPA to cover more parts of the AWE assessment process while preserving the modular structure described in this chapter.

5 Wind Profile Clustering

The wind profile clustering module processes different types of wind data, including ERA5, DOWA, and Lidar measurements, to identify representative wind profile shapes for use in performance modelling. In this work, ERA5 and Lidar data are used. The workflow consists of height calculation for ERA5 model levels, horizontal and vertical interpolation, preprocessing of wind data, dimensionality reduction, clustering, and export of representative clustered profiles for further analysis. In the second part of this chapter, the ERA5 data are compared with a Lidar dataset provided by Kitepower for the test site in Bangor Erris, Ireland.

5.1. Height calculation for ERA5 model levels

For airborne wind energy system simulations, wind speeds must be expressed at physical altitudes above ground level. ERA5 model levels are defined using a surface-following pressure coordinate system and therefore cannot be interpreted directly as geometric heights. Two approaches are implemented depending on data availability: Method 1 (hydrostatic reconstruction using thermodynamic variables) [37] and Method 2 (standard-atmosphere lookup approximation). Method 1 provides a physically consistent solution and is preferred whenever the required thermodynamic variables are available. Method 2 introduces approximation errors because it assumes a static atmospheric structure.

5.1.1. Method 1: Hydrostatic reconstruction using thermodynamic variables

Geopotential altitudes may be reconstructed from thermodynamic state variables through application of the hydrostatic equation. This method combines temperature and humidity at model levels with surface pressure and surface geopotential.

ERA5 defines pressure at model half-level interfaces through hybrid coefficients A_k and B_k, which relate the pressure at interface k to the surface pressure as shown in Equation 5.1 [37], [38]:

$$p_k = A_k + B_k\, p_{\mathrm{s}} \tag{5.1}$$

where p_{s} denotes surface pressure.

Assuming hydrostatic equilibrium, the geopotential difference across a layer depends on the virtual temperature and the pressure ratio between its bounding interfaces, as expressed in Equation 5.2:

$$\Delta\Phi_k = R_{\mathrm{d}}\, T_{\mathrm{v}}\, \ln\!\left(\frac{p_k}{p_{k-1}}\right) \tag{5.2}$$

The virtual temperature is computed according to Equation 5.3:

$$T_{\mathrm{v}} = T\,(1 + 0.61\, q) \tag{5.3}$$

where T is temperature, q is specific humidity, and R_d is the specific gas constant for dry air.

Starting from the known surface geopotential Φ_s, the geopotential at model level l is obtained by vertically accumulating the layer contributions:

$$\Phi_l = \Phi_\mathrm{s} + \sum_{k=1}^{l} \Delta\Phi_k \tag{5.4}$$

The height above ground level then follows from removing the surface contribution and converting to geometric height, as shown in Equation 5.5:

$$z_l = \frac{\Phi_l - \Phi_\mathrm{s}}{g_\mathrm{ref}} \tag{5.5}$$

This approach produces altitudes that vary with atmospheric stratification and moisture content, thereby yielding a physically consistent representation of the vertical spacing between model levels.

5.1.2. Method 2: Standard-atmosphere lookup approximation

If the required thermodynamic variables are unavailable, model levels are mapped to the predefined geopotential altitudes derived from the ICAO standard atmosphere [38]. These tabulated values assume globally uniform stratification and neglect meteorological variability. Heights obtained in this manner are constant in time and independent of local weather conditions.

5.1.3. Application in this work

Both approaches are implemented to maintain compatibility with varying ERA5 data availability. Hydrostatic reconstruction is preferred whenever the necessary input variables are present because it provides a more realistic representation of the atmospheric vertical structure. The standard-atmosphere approximation is only employed when the required input data are unavailable.

Finally, the wind components are interpolated to the heights used during clustering using linear interpolation, ensuring that wind profiles are expressed at consistent altitudes for subsequent analysis and clustering.

5.2. Horizontal interpolation procedure

ERA5 variables are defined on a regular latitude–longitude grid, whereas the requested location for analysis will, in general, not coincide with a grid point. The atmospheric state at the requested location is therefore obtained through bilinear interpolation using the 4 ERA5 grid points that surround the target position.

Let the requested location be denoted by (ϕ^*, λ^*), where ϕ is latitude and λ is longitude. Suppose this location lies inside the grid cell bounded by the latitudes ϕ_1 and ϕ_2, and the longitudes λ_1 and λ_2, such that the 4 surrounding ERA5 grid points are

$$(\phi_1, \lambda_1), \qquad (\phi_1, \lambda_2), \qquad (\phi_2, \lambda_1), \qquad (\phi_2, \lambda_2).$$

The relative position of the requested location within this grid cell is described by the normalised coordinates

$$\alpha = \frac{\phi^* - \phi_1}{\phi_2 - \phi_1}, \qquad \beta = \frac{\lambda^* - \lambda_1}{\lambda_2 - \lambda_1}, \tag{5.6}$$

with $\alpha, \beta \in [0, 1]$. These parameters define the bilinear weights associated with the 4 surrounding grid points:

$$w_{11} = (1-\alpha)(1-\beta), \qquad w_{12} = (1-\alpha)\beta, \qquad w_{21} = \alpha(1-\beta), \qquad w_{22} = \alpha\beta, \tag{5.7}$$

which satisfy

$$w_{11} + w_{12} + w_{21} + w_{22} = 1. \tag{5.8}$$

A grid point located close to the requested position therefore receives a larger weight, while more distant grid points contribute less. If the requested location coincides exactly with a grid point, the corresponding weight becomes 1 and all other weights vanish.

For a generic variable X defined at the 4 surrounding grid points, the interpolated value at the requested location is given by

$$X^* = w_{11}X_{11} + w_{12}X_{12} + w_{21}X_{21} + w_{22}X_{22}, \tag{5.9}$$

where X_{ij} denotes the value at grid point (ϕ_i, λ_j).

In this work, horizontal interpolation is not applied directly to the raw model-level atmospheric variables prior to the reconstruction of geometric heights. Instead, the vertical wind profile is first reconstructed independently at each of the 4 surrounding ERA5 grid points, after which the resulting profiles are combined using the bilinear weights defined in Equation 5.7.

This procedure is required for consistency when using hydrostatic reconstruction. The conversion from model levels to geometric height depends on temperature, humidity, surface pressure, and surface geopotential. Averaging these variables prior to the height reconstruction would, in general, not yield the same result as reconstructing the vertical structure at each grid point individually. Therefore, the physically consistent approach is to first compute the vertical profiles at each surrounding grid point and subsequently apply horizontal interpolation.

Let $u_{ij}(t, z)$ and $v_{ij}(t, z)$ denote the eastward and northward wind components at time t and altitude z for the grid point (ϕ_i, λ_j). The interpolated wind components at the requested location are then given by

$$u^*(t, z) = w_{11}u_{11}(t, z) + w_{12}u_{12}(t, z) + w_{21}u_{21}(t, z) + w_{22}u_{22}(t, z), \tag{5.10}$$

$$v^*(t, z) = w_{11}v_{11}(t, z) + w_{12}v_{12}(t, z) + w_{21}v_{21}(t, z) + w_{22}v_{22}(t, z). \tag{5.11}$$

The resulting wind profile at the requested location is thus a distance-weighted combination of the 4 surrounding ERA5 profiles, with the largest contribution originating from the grid point closest to the target position.

5.3. Preprocessing of wind data

Each wind profile consists of vertical wind speed measurements at multiple height levels for a given time and location. To isolate the shape of the profile from directional and magnitude effects, the data are processed in 2 stages.

First, the wind velocity components are transformed into components parallel and perpendicular to the wind direction at a reference height, which in this work is $200\,\mathrm{m}$. This removes directional dependence, since the perpendicular component becomes zero at the reference height [19].

Second, the wind speed components are normalised using the 90th percentile of the wind-speed magnitude of each profile. Using a percentile rather than the maximum reduces sensitivity to outliers and yields a more compact representation of the wind resource. The normalised parallel and perpendicular components together define the wind profile shape [19].

Under low-wind conditions, normalisation by the 90th-percentile wind speed can amplify noise, producing profile shapes that are not physically representative. Profiles with a mean wind speed below

$5\,\mathrm{m\,s^{-1}}$ are therefore excluded when training the clustering model. This is justified because such conditions contribute negligibly to annual energy production and their shapes are of limited relevance for performance modelling.

Nevertheless, low-wind profiles are not discarded entirely. After the model is trained, all profiles, including those excluded during training, are preprocessed and passed through the trained model, which assigns each to its nearest cluster centroid. This ensures that cluster frequencies account for the full dataset and that every time step receives a cluster label.

5.4. Dimensionality reduction and cluster number selection

To reduce dimensionality while retaining the dominant variability in profile shapes, principal component analysis (PCA) is applied. The first 5 principal components are retained, as they capture the majority of the variance in the wind-profile dataset while maintaining computational efficiency [19].

Clustering is performed using the k-means algorithm [20], which partitions the dataset into clusters by minimising the within-cluster sum of squared Euclidean distances, also referred to as inertia.

The selection of the number of clusters k involves balancing statistical representativeness with interpretability and computational cost.

In Schelbergen et al. [19], 8 clusters were selected based on several criteria. First, the mean wind-profile reconstruction error was evaluated to determine how accurately the cluster representatives reproduce the original wind profiles. Second, the silhouette score was used to assess the separation and cohesion of the clusters. In addition, representation validity was examined to ensure that the clusters capture the main variability present in the dataset. Finally, the number of clusters was chosen to provide a concise and interpretable set of representative wind-profile types.

A similar trade-off is observed in Sommerfeld et al. [18]. Their analysis shows that the elbow method and silhouette score suggest a larger number of clusters, with $k \approx 20$, for an optimal statistical representation. For visualisation and interpretability purposes, however, they present results using 10 clusters.

Based on these considerations, 8 clusters are adopted in this work. This choice provides sufficient diversity in wind-profile shapes while keeping the dataset manageable for analysis and interpretation.

It should be noted that the optimal number of clusters depends on the intended application. Larger cluster sets improve statistical fidelity, while smaller sets improve clarity and computational efficiency. The clustering module in AWESPA allows the number of clusters to be modified if higher resolution is required.

5.5. Exporting clustered profiles

The clustered wind profiles are exported to an awesIO wind-resource YAML file for AEP calculations. Each cluster centroid represents a characteristic wind-profile shape described by its parallel and perpendicular wind-speed components.

At the reference height, the perpendicular wind component is zero by definition and the parallel component is normalised to 1 to standardise the profile shape. This second normalisation is applied after the earlier percentile-based scaling and therefore does not depend on it; the percentile normalisation does not, in general, guarantee a unit value at the reference height. The enforced unit value produces a compact and consistent representation of profile shape that can be scaled to any reference wind speed in subsequent performance simulations. The exported YAML file additionally records, for each cluster, a 3×3 matrix giving the frequency of occurrence across wind-speed and wind-direction bins.

5.6. Comparison with Lidar data

This section compares wind measurements obtained from a WindCube at the Kitepower test site in Bangor Erris, Ireland, with ERA5 data. This comparison is important because ERA5 data are used to estimate the annual energy production of the case studies presented in Chapter 8. In particular, case study 1 is located at the Kitepower test site itself, making the present comparison especially relevant for assessing the suitability of ERA5 for the subsequent analyses.

5.6.1. WindCube WLS7-130

The measurement device used at the site was a WindCube WLS7-130. A Doppler wind Lidar measures the wind by emitting laser pulses into the atmosphere and detecting the Doppler shift of the backscattered light from aerosols and other small particles suspended in the air. This frequency shift provides the component of the wind velocity along the laser beam, often referred to as the line-of-sight velocity. By combining measurements from multiple beam directions, the instrument reconstructs the horizontal wind speed and wind direction at a set of discrete heights above ground level. [39]

An important distinction between Lidar measurements and ERA5 data is that the Lidar provides a local observation of the wind field at the measurement site, whereas ERA5 represents a spatially averaged and model-smoothed atmospheric state over a finite grid cell. The Lidar data therefore contain site-specific effects more directly, including local turbulence, shear variations, and flow distortions induced by terrain or nearby obstacles. At the Bangor Erris test site, the WindCube was located approximately $30\,\mathrm{m}$ from a hangar. In addition, a two-story building was situated to the north of the WindCube at a distance of approximately $67\,\mathrm{m}$.

Figure 5.1: Satellite image of the Kitepower Bangor Erris test site. The blue arrow indicates the location of the WindCube.

5.6.2. Wind speed data

The Lidar dataset made available for this analysis was provided by Kitepower and spans the period from 20 October 2024 to 11 June 2025. However, a substantial number of days contain either no measurements or only partial-day coverage. This is shown in Figure 5.2.

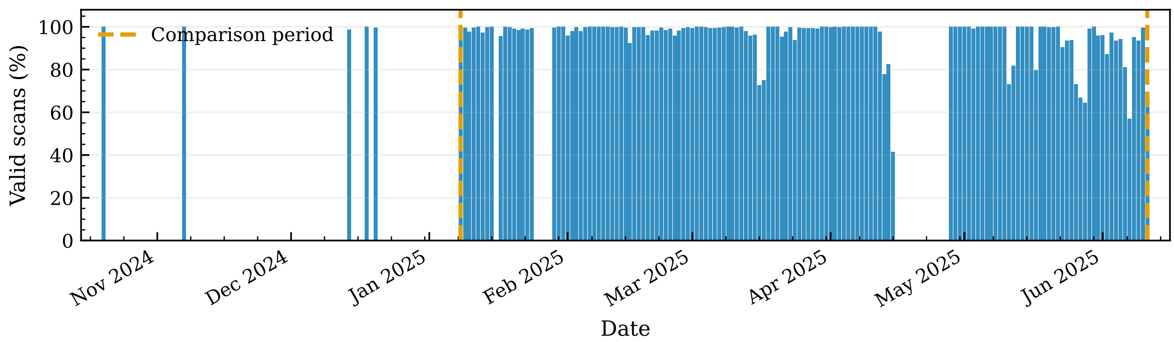

Figure 5.2: Lidar data availability. The period used for the comparison is bounded by the two orange vertical lines.

Because of the limited data availability, only the period from 8 January 2025 to 11 June 2025 was kept for the comparison. The ERA5 dataset was filtered accordingly for all subsequent analyses, such that all data points for which no Lidar data were available were excluded.

The ERA5 data were downloaded from the ECMWF database. The dataset has a horizontal resolution of $0.25°$ and covers the domain $(65°/-20°/30°/20°)$, which includes most of Europe. The data extend to model level 100 ($\sim 4341\,\mathrm{m}$) and include the thermodynamic variables required to apply Method 1 described in Section 5.1. The temporal resolution is 1 hour.

Because ERA5 has an hourly temporal resolution, the Lidar data were averaged to hourly means to ensure temporal consistency between both datasets.

Figure 5.3 shows the 4 surrounding ERA5 grid points and the location of the Kitepower test site.

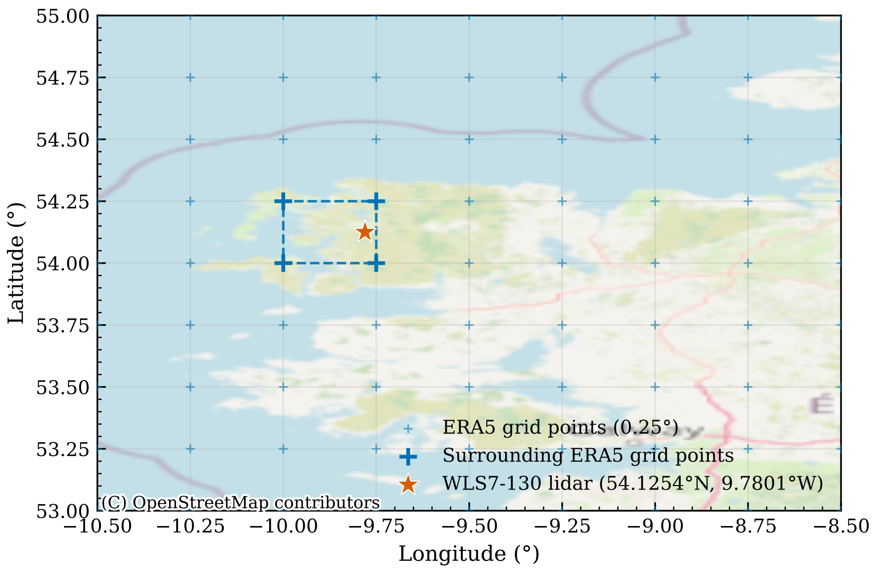

Figure 5.3: ERA5 grid points surrounding the test site in Bangor Erris, Ireland.

5.6.3. Mean wind speed, Weibull distribution and wind rose

Figure 5.4, which shows the mean wind speed and standard deviation at every altitude measured by the WindCube, reveals 2 consistent features. First, the mean wind speed obtained from ERA5 is lower than that measured by the Lidar at nearly all altitudes. The only exceptions occur at $240\,\mathrm{m}$, where both means are approximately equal, and at $250\,\mathrm{m}$, where the ERA5 mean is slightly higher. Overall, this indicates that ERA5 at this site exhibits a small negative wind-speed bias over most of the height range considered.

Second, the standard deviation is larger for the Lidar at every altitude. This indicates that the Lidar records a wider spread of wind speeds than ERA5. Such behaviour is physically plausible, because ERA5 represents a spatially averaged and model-smoothed estimate of the atmospheric state over a grid cell, whereas the Lidar measures the wind locally at a single site. As a result, the Lidar captures site-specific variability more directly, including local turbulence, shear variations, and terrain-induced effects that are not fully resolved by ERA5.

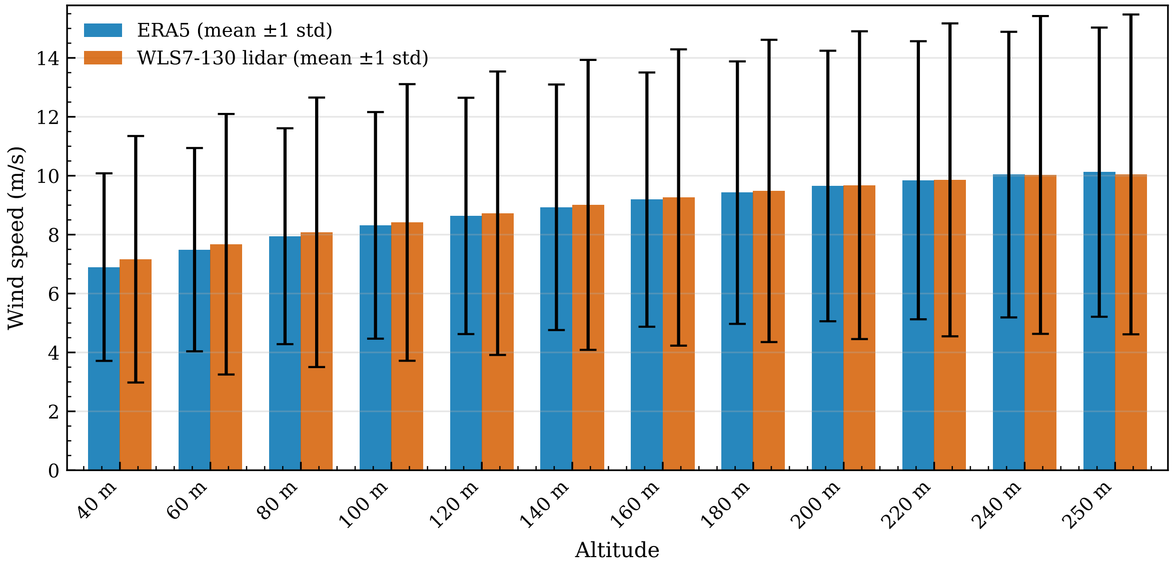

Figure 5.4: Comparison of mean wind speed and standard deviation for Lidar and ERA5 at all altitudes measured by the WindCube.

To determine whether the larger standard deviation of the Lidar data could be explained primarily by outliers, the interquartile range (IQR) was also evaluated. In contrast to the standard deviation, the IQR measures only the spread of the central 50% of the distribution and is therefore much less sensitive to extreme values. If the difference between Lidar and ERA5 were mainly caused by outliers, then the ratio of the standard deviations would be expected to remain noticeably above 1, while the ratio of the IQRs would move much closer to 1. This is not observed. As shown in Figure 5.5, the IQR of the Lidar data is also consistently larger than that of ERA5 at all altitudes. This demonstrates that the broader spread in the Lidar measurements is not merely a tail effect caused by a limited number of extreme events. Instead, the central bulk of the Lidar wind-speed distribution is also wider. The most important conclusion is therefore that ERA5 underrepresents the actual variability of the local wind field at this site, rather than simply missing a small number of rare wind-speed spikes.

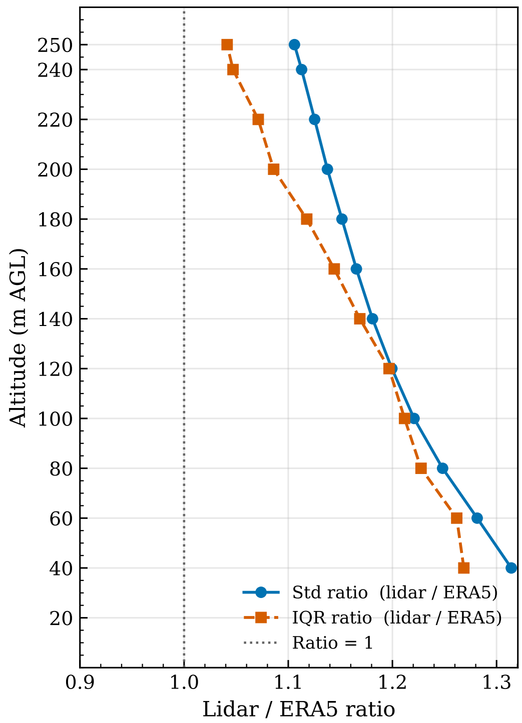

Figure 5.5: Comparison of standard deviation and interquartile range for Lidar and ERA5.

This interpretation is supported by the Weibull fits shown in Figure 5.6. At $100\,\mathrm{m}$, the fitted parameters are $k = 1.83$ and $c = 9.45\,\mathrm{m\,s^{-1}}$ for the Lidar data, and $k = 2.30$ and $c = 9.39\,\mathrm{m\,s^{-1}}$ for ERA5. At $200\,\mathrm{m}$, the fitted parameters are $k = 1.92$ and $c = 10.90\,\mathrm{m\,s^{-1}}$ for the Lidar data, and $k = 2.22$ and $c = 10.89\,\mathrm{m\,s^{-1}}$ for ERA5. The scale parameter c is thus nearly identical for both datasets at both heights, indicating that they represent a similar characteristic wind-speed level. However, ERA5 yields consistently larger values of the shape parameter k, which implies a narrower and more strongly peaked wind-speed distribution. In contrast, the smaller k values obtained from the Lidar data indicate a broader distribution and therefore greater wind-speed variability. ERA5 therefore captures the general magnitude of the wind climate reasonably well, but underrepresents the spread of local wind conditions.

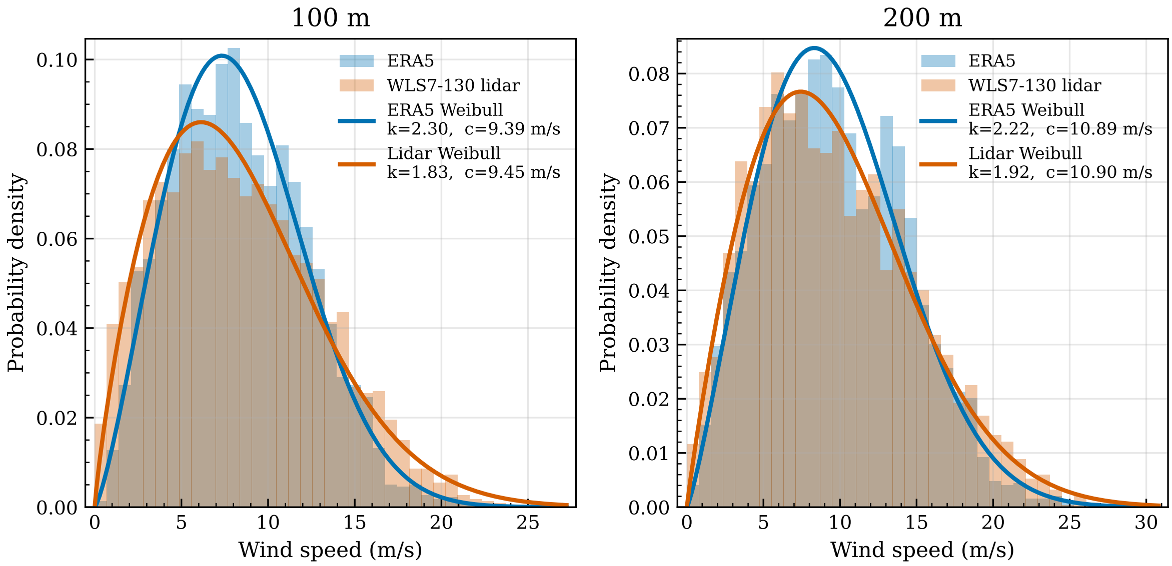

Figure 5.6: Weibull comparison of Lidar and ERA5 data at altitudes $100\,\mathrm{m}$ and $200\,\mathrm{m}$.

The wind roses shown in Figure 5.7 indicate good overall agreement between the Lidar and ERA5 data at both $100\,\mathrm{m}$ and $200\,\mathrm{m}$. Although differences remain in the relative frequency of individual directional sectors, both datasets identify the same dominant wind-direction sector. This suggests that ERA5 reproduces the prevailing flow direction at the site reasonably well.

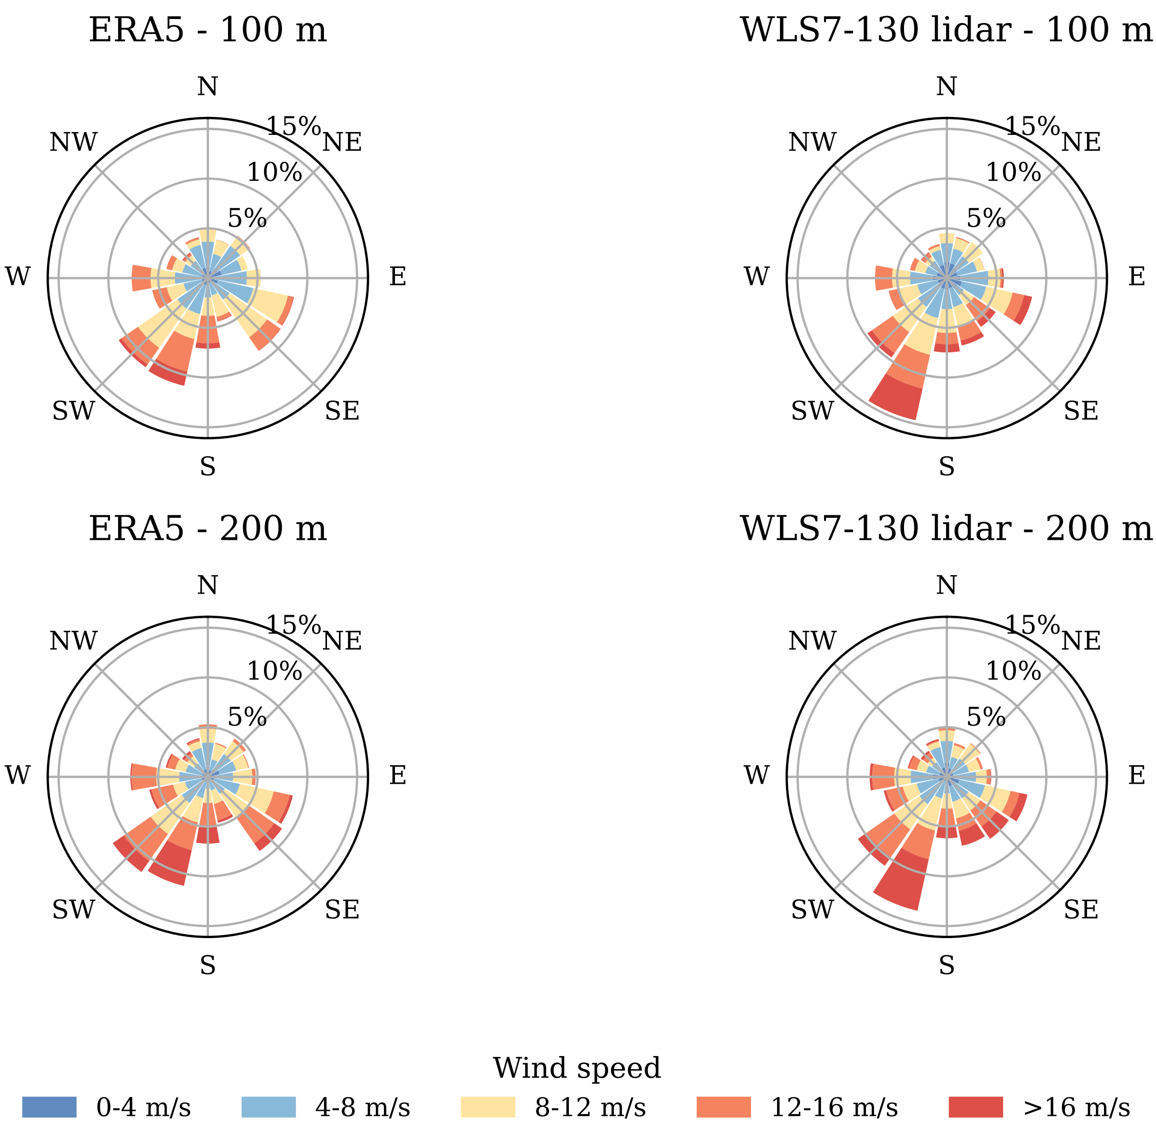

Figure 5.7: Wind rose comparison of Lidar and ERA5 data at altitudes $100\,\mathrm{m}$ and $200\,\mathrm{m}$.

5.6.4. Wind profiles

The analysis now shifts from wind-speed statistics to the vertical structure of the wind profiles. First, the mean wind profile is considered by evaluating the mean wind speed at each altitude measured by the WindCube and subsequently fitting a logarithmic profile to this mean profile.

Figure 5.8 shows the dimensional mean profile together with the logarithmic fit for both the Lidar and ERA5 data. The same general behaviour observed in the previous subsection is visible here as well: ERA5 yields lower wind speeds than the Lidar over most of the height range. However, apart from this offset in magnitude, the overall shape of the mean profile is similar. This suggests that ERA5 reproduces the average vertical structure of the flow reasonably well.

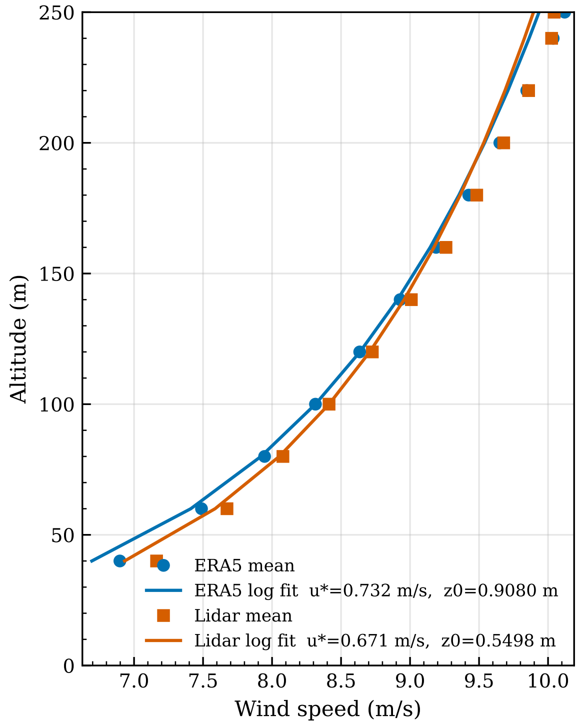

Figure 5.8: Mean wind profile and logarithmic fit of Lidar and ERA5 wind data.

A second way to compare profile shape is through the shear exponent. For each individual profile in both the Lidar and ERA5 datasets, a power-law profile is fitted and the corresponding shear exponent is extracted. The resulting distributions of these exponents are shown in Figure 5.9.

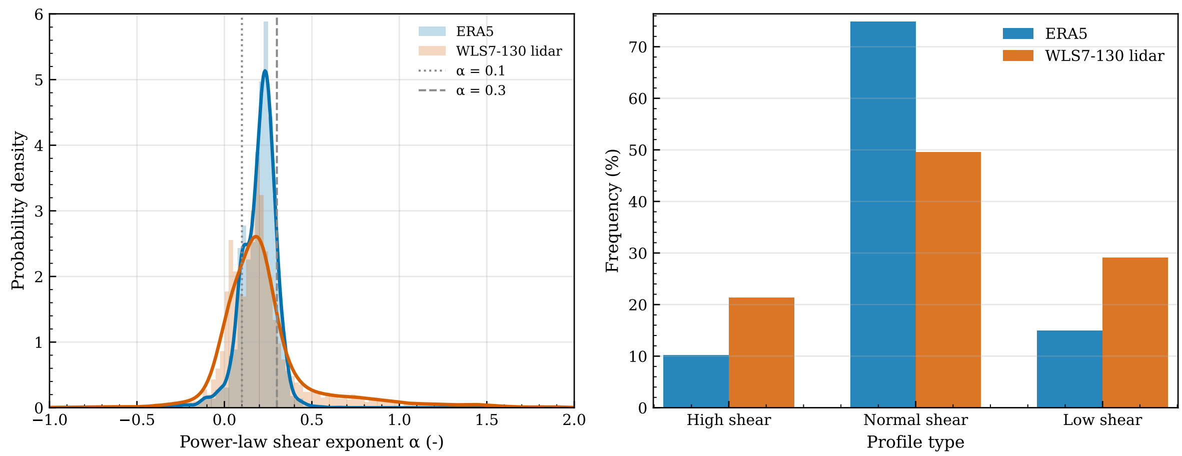

Figure 5.9: Shear exponent distribution for Lidar and ERA5. High shear is defined as $\alpha \geq 0.3$ and low shear as $\alpha \leq 0.1$.

The shear-exponent distribution is considerably wider for the Lidar data than for ERA5. This indicates that the Lidar captures many more cases with both very low shear and very high shear, whereas ERA5 produces a narrower and more concentrated range of profile shapes. This is fully consistent with the earlier finding that ERA5 underrepresents local variability. It suggests that ERA5 not only smooths the wind-speed magnitude but also smooths the vertical structure of the wind field.

It should be noted, however, that a single shear exponent is not always an ideal descriptor of the profile shape. In particular, profiles containing a low-level jet are not well represented by a power-law fit. In such cases, the fitted exponent can be misleading, because the presence of a jet introduces curvature and a local wind-speed maximum that cannot be captured by a single shear value. The shear-exponent analysis should therefore be interpreted as a broad indicator of vertical variability rather than as a complete description of all profile types.

To investigate the profile structure in more detail, the clustering model is used to identify the typical wind profiles encountered at the test site. First, the model is trained on the ERA5 profiles, producing 8 representative clusters, which are shown in Figure 5.10. The Lidar profiles are then assigned to these ERA5-based clusters using the same nearest-centroid procedure as in the clustering algorithm itself, but with the cluster centroids fixed to those obtained from the ERA5 training data. The resulting frequencies are compared in Figure 5.11. The same procedure is then repeated with the model trained on the Lidar data. The resulting Lidar-based clusters are shown in Figure 5.12, and the corresponding frequencies for both datasets are presented in Figure 5.13.

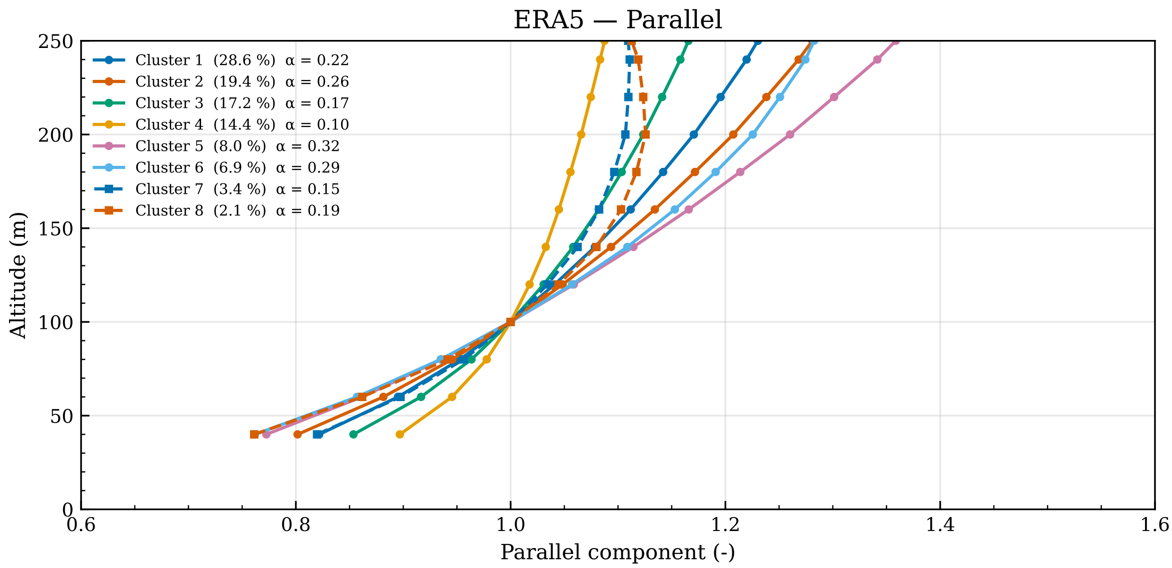

Figure 5.10: Resulting wind profiles from clustering the filtered ERA5 data profiles.

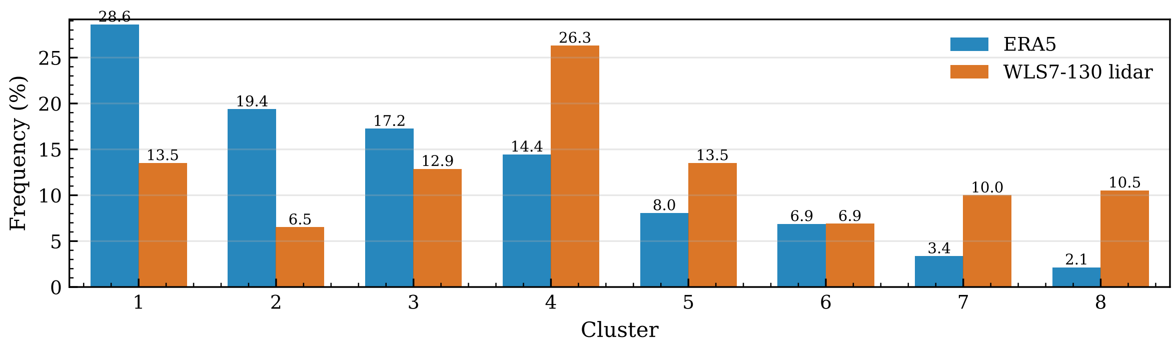

Figure 5.11: Cluster frequencies of the Lidar and ERA5 data for clusters trained on ERA5 data.

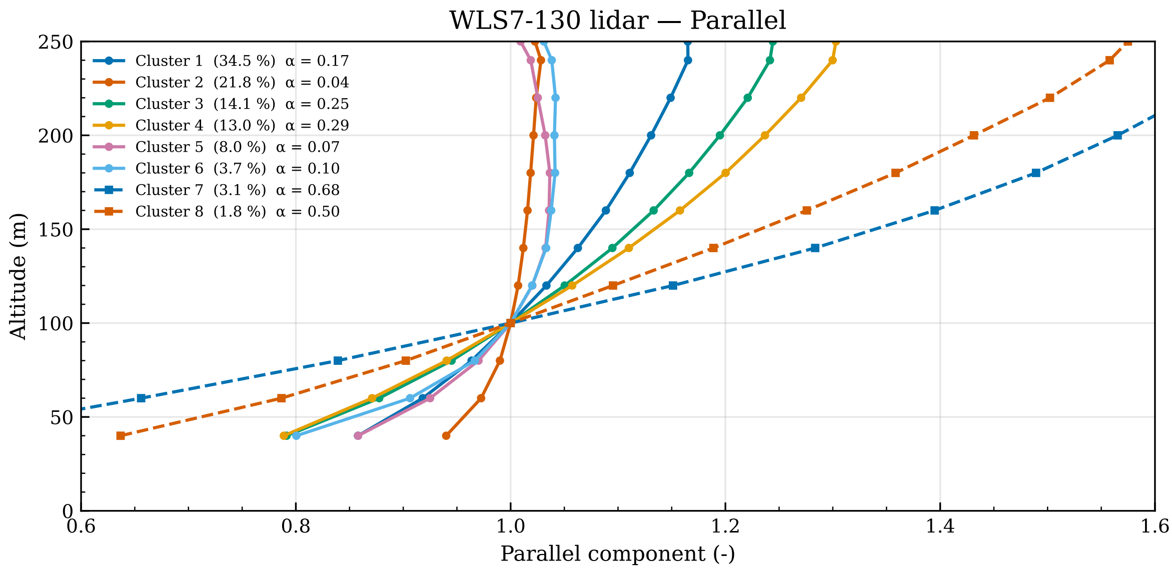

Figure 5.12: Resulting wind profiles from clustering the Lidar data profiles.

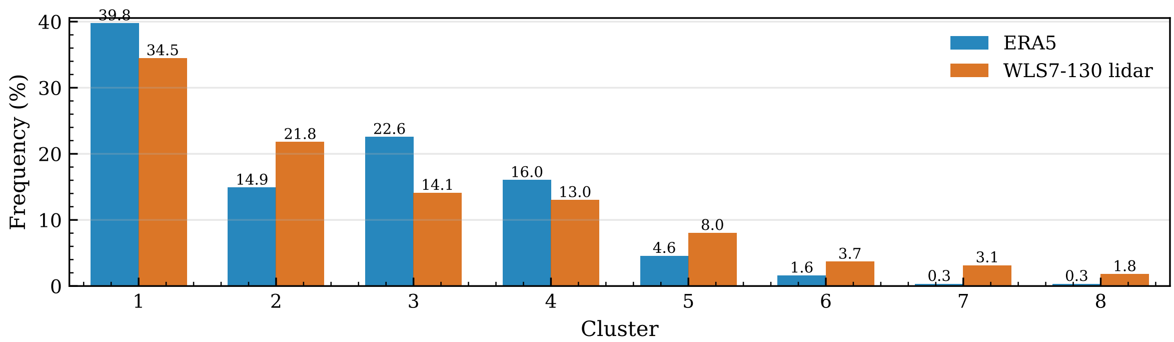

Figure 5.13: Cluster frequencies of the Lidar and ERA5 data for clusters trained on Lidar data.

The ERA5-trained clusters reveal several important differences between the 2 datasets. Profiles 7 and 8, which resemble low-level-jet-type structures, occur much more frequently in the Lidar data than in ERA5. For profile 7, the occurrence increases from 3.4% in ERA5 to 10.0% in the Lidar data, while for profile 8 it increases from 2.1% to 10.5%. A similar behaviour is found for profile 5, which represents a high-shear profile and occurs with a frequency of 8.0% in ERA5 and 13.5% in the Lidar data. A comparable pattern is observed for profile 4, which represents the lowest-shear profile among the 8 ERA5-based clusters. Its frequency increases from 14.4% in ERA5 to 26.3% in the Lidar data. These differences indicate that ERA5 underrepresents both strongly sheared and weakly sheared profile types observed at the site, while favouring more moderate profiles.

The Lidar-trained clusters lead to the same overall conclusion. The more extreme profile types again occur more frequently in the Lidar data than in ERA5. In particular, the very high-shear profiles 7 and 8 are almost absent in ERA5, with frequencies of only 0.3% for both profiles. Similarly, the low-shear profile 2 occurs 6.9% more often in the Lidar data than in ERA5. By contrast, the more moderate and conventional profiles, such as profiles 1, 3, and 4, occur more frequently in ERA5 than in the Lidar data. This again indicates that ERA5 favours smoother and more average profile shapes, while the Lidar measurements capture a broader range of profile variability.

5.7. Conclusion

This chapter presented the wind profile clustering procedure implemented in AWESPA and evaluated the suitability of ERA5 data for representing the wind resource at the Kitepower test site in Bangor Erris. The clustering workflow consists of reconstructing physical heights for ERA5 model levels, interpolating the data to the requested location and analysis heights, preprocessing the wind profiles, reducing their dimensionality using PCA, and grouping them into representative profile classes using k-means clustering. The resulting clustered profiles provide a compact description of the wind resource that can be used in later performance and AEP calculations.

The comparison with Lidar data showed that ERA5 reproduces the first-order characteristics of the wind climate reasonably well. In particular, it captures the overall vertical shape of the mean profile, yields Weibull scale parameters that are close to those obtained from the Lidar data, and reproduces the dominant wind-direction sector at both $100\,\mathrm{m}$ and $200\,\mathrm{m}$. The mean-profile comparison further showed that, although ERA5 yields lower wind speeds than the Lidar over most of the height range, the general shape of the vertical profile remains similar. This indicates that ERA5 provides a reasonable representation of the average vertical structure of the flow, even if it underestimates the site-specific wind-speed level.

However, the comparison also revealed clear limitations. ERA5 underestimates the mean wind speed at most heights and consistently exhibits a smaller standard deviation and interquartile range than the Lidar data, indicating that the reanalysis represents a narrower spread of wind-speed conditions. This interpretation is further supported by the Weibull analysis, in which ERA5 yields consistently larger shape parameters, corresponding to a more peaked and less variable wind-speed distribution. A similar tendency was reported by Gualtieri [17], who found that ERA5 wind-speed predictions generally exhibit less variability than observations. This reduced variability is also visible in the wind-profile analysis: the distribution of shear exponents is narrower in ERA5, and the clustering results show that more extreme profile types, such as high-shear profiles and low-level-jet-like profiles, occur much more frequently in the Lidar data than in ERA5. More broadly, the representative ERA5 profiles are more regular and less diverse than those derived from the Lidar data. Whereas ERA5 mainly produces smoother and more standardised profile shapes, the Lidar data span a wider range of wind-profile behaviour, including both very high-shear cases and more pronounced low-level-jet structures.

This difference is particularly relevant for airborne wind energy applications, because such systems are sensitive not only to the mean wind speed but also to the detailed profile shape over a large altitude range. Underrepresentation of high-shear and low-level-jet cases in ERA5 may therefore lead to an incomplete characterisation of the local wind resource if the reanalysis is used without site-specific validation or correction. Overall, the results show that ERA5 is suitable as a first-order wind-resource description for the present work, but that it smooths both wind-speed variability and profile-shape variability at the Bangor Erris site. Site-specific measurements such as Lidar data therefore remain valuable for validation and for identifying local wind features that are not fully captured by the reanalysis.

6 Power estimation

This chapter describes the power-estimation models used in AWESPA and the system parameters used for the TU Delft V3 kite. First, the choice of power models is motivated. The V3 system configuration is then introduced. Secondly, the Luchsinger model implementation is described and compared with its soft-kite variant. Finally, the improvements and results from the inertia-free QSM are discussed.

6.1. Choice of power model

For the present work, two power models are implemented in AWESPA: the Luchsinger model [4], [10] and the more detailed inertia-free quasi-steady model of Schelbergen [5], [12], which is based on the model of Van der Vlugt et al. [16]. These models are included for different purposes. The Luchsinger model is mainly used as a simple implementation case to test the modular structure of the AWESPA toolchain and the awesIO file-based input and output workflow discussed in Chapter 4. The inertia-free quasi-steady model is used for the power-production estimates in the case-study analysis presented in Chapter 8.

The Luchsinger model is included because it is simple, fast, and requires only a limited set of system and operational parameters. This makes it suitable for testing whether AWESPA can switch between power models while reusing the same awesIO system configuration and wind-resource input structure. Its simplified representation of the pumping cycle, using representative reel-out and reel-in states, makes it useful for checking the modularity of the implementation, but it is not used as the main model for estimating the power production of the case studies.

The inertia-free quasi-steady model is selected as the main power-estimation model because it represents the pumping cycle in more detail while remaining computationally efficient. Compared with the Luchsinger model, it includes a more detailed treatment of the flight path and it accounts for kite mass through gravitational loading, while still neglecting inertial accelerations, making it more suitable for estimating power production. For this reason, the inertia-free quasi-steady model is used for the power-production estimates in the case-study chapter.

6.2. TU Delft V3 kite

The AWE system used to generate the results in this chapter is based on the TU Delft V3 kite. This leading edge inflatable kite has a flat surface area of 25 m^2 and a projected surface area of 19.75 m^2, which is the reference area used in the simulations. The combined mass of the wing and bridle system is 10.6 kg [40]. The kite control unit (KCU), which is suspended below the kite and controls its flight, has a mass of 8.4 kg [41]. This gives a total airborne mass of 19 kg.

Both power models require the lift and drag coefficients in the powered and depowered aerodynamic states. These coefficients are taken from Cayon et al. [31]. In the powered state, the lift and drag coefficients are $C_\mathrm{L} = 0.63$ and $C_\mathrm{D} = 0.14$, respectively. In the depowered state, the corresponding values are $C_\mathrm{L} = 0.40$ and $C_\mathrm{D} = 0.12$.

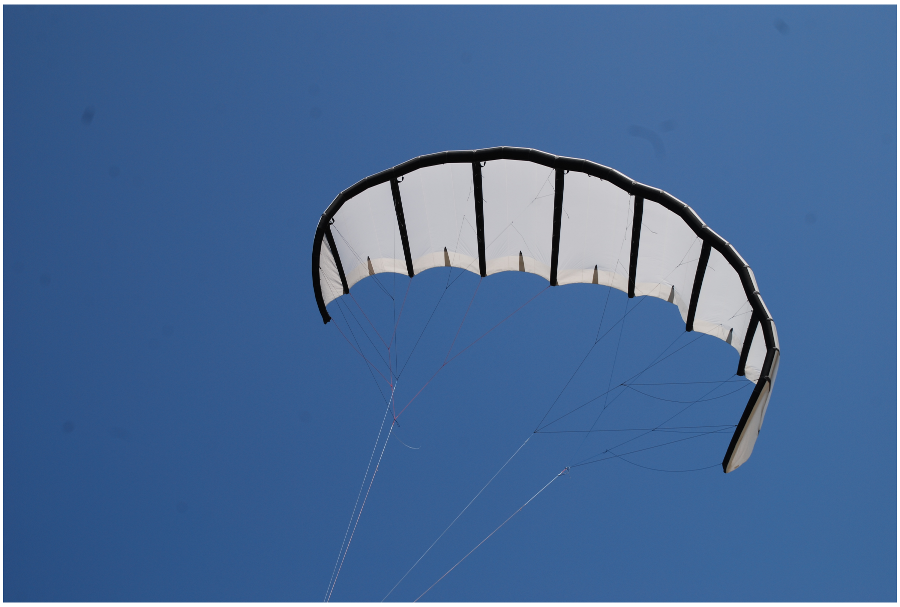

Figure 6.1: V3 kite flying behind the Aerospace Engineering Faculty of TU Delft (2012)

The tether diameter is 6 mm, the tether density is 724 kgm^{-3}, and the tether drag coefficient is 1.1. The tether force is constrained between 750 N and 8.2 kN. In this study, the drum reeling speed limit is set to 10 ms^{-1}, and the generator power limit is set to 40 kW. [5] The maximum tether length is set at 500m.

The awesIO system configuration file corresponding to this is shown in Appendix A and all variables used are summarized in Table 6.1.

Table 6.1: System and operational parameters used for the TU Delft V3 kite simulations.

Parameter	Symbol	Value
Flat kite surface area	A_{flat}	$25\ \mathrm{m}^2$
Projected kite surface area	A	$19.75\ \mathrm{m}^2$
Wing and bridle mass	m_{kite}	$10.6\ \mathrm{kg}$
KCU mass	m_{KCU}	$8.4\ \mathrm{kg}$
Total airborne mass	m_{airborne}	$19\ \mathrm{kg}$
Powered lift coefficient	$C_{\text{L,o}}$	0.63
Powered drag coefficient	$C_{\text{D,o}}$	0.14
Depowered lift coefficient	$C_{\text{L,i}}$	0.40
Depowered drag coefficient	$C_{\text{D,i}}$	0.12
Tether diameter	d_{t}	$6\ \mathrm{mm}$
Tether density	ρ_{t}	$724\ \mathrm{kgm}^{-3}$
Tether drag coefficient	$C_{\text{D,t}}$	1.1
Maximum tether length	l_{max}	$500\ \mathrm{m}$
Minimum tether force	$F_{\text{t,min}}$	$750\ \mathrm{N}$
Maximum tether force	$F_{\text{t,n}}$	$8.2\ \mathrm{kN}$
Maximum drum reeling speed	$v_{\text{r,max}}$	$10\ \mathrm{ms}^{-1}$
Mechanical power limit	P_{n}	$40\ \mathrm{kW}$

6.3. Luchsinger power model

This section presents the Luchsinger pumping-cycle model as implemented in AWESPA. The model is a compact analytical power model in which the pumping cycle is represented by two representative states, one for each of the two phases, reel-out and reel-in. The transition phases between reel-out and reel-in are not resolved.

The implementation contains two variants: the original formulation of Luchsinger [10], which was primarily developed for fixed-wing kites and uses a fixed elevation angle reel-in phase, and the soft-kite formulation of Schmehl et al. [4] which uses a fixed lift over drag ratio and a variable elevation angle during the reel-in phase. Both variants use the same pumping-cycle structure, the same wind-profile input, and the same three-regime strategy. The cycle power is optimized in both models by optimizing the reeling factors according to the equations shown in Section 6.3.2 and Section 6.3.3.

The notation in this section follows Schmehl et al. [4]. The dimensionless reeling factor is denoted by f, with $f_\mathrm{o} > 0$ during reel-out and $f_\mathrm{i} < 0$ during reel-in. The aerodynamic force factors are denoted by γ_o and γ_i, where the subscripts o and i refer to reel-out and reel-in, respectively.

6.3.1. Common formulation

Both variants are quasi-steady. Kite mass, tether elasticity, and inertial effects are neglected. In the present implementation, the original Luchsinger model is modified by adding tether drag through the corrected reel-out drag coefficient in Equation 6.6. The same correction is applied to the soft-kite variant, so the comparison between the variants focuses on the aerodynamic force-factor definitions and reel-in phase modeling rather than on whether tether drag is included. For the original model, the reel-in elevation angle β_i is prescribed as an input. For the soft-kite model, the reel-in angle follows from the constant lift-to-drag assumption and is therefore not prescribed independently.

For one selected vertical wind profile i, with normalized wind speed $u_{\mathrm{norm},i}(z)$ and reference wind speed v_ref, the model first computes an average wind speed over the operating altitude interval,

$$\bar{v}_\mathrm{w} = v_\mathrm{ref} \frac{1}{z_\mathrm{max} - z_\mathrm{min}} \int_{z_\mathrm{min}}^{z_\mathrm{max}} u_{\mathrm{norm},i}(z)\,\mathrm{d}z, \tag{6.1}$$

with

$$z_\mathrm{min} = l_\mathrm{min} \sin\beta_\mathrm{o}, \qquad z_\mathrm{max} = l_\mathrm{max} \sin\beta_\mathrm{o}. \tag{6.2}$$

The same averaged wind speed $\bar{v}_\mathrm{w}$ is then used in both phases of the cycle.

The reel-out and reel-in speeds are written as

$$v_\mathrm{o} = f_\mathrm{o}\bar{v}_\mathrm{w}, \qquad v_\mathrm{i} = f_\mathrm{i}\bar{v}_\mathrm{w}, \tag{6.3}$$

where $f_\mathrm{o} > 0$ and $f_\mathrm{i} < 0$. The reeling length is

$$l_\mathrm{reel} = l_\mathrm{max} - l_\mathrm{min}. \tag{6.4}$$

The reel-out factor is bounded by

$$0 < f_\mathrm{o} \le \min\left(\frac{v_\mathrm{o,max}}{\bar{v}_\mathrm{w}}, 1\right). \tag{6.5}$$

In both model variants, tether drag is included in the reel-out drag coefficient as

$$r_\mathrm{m,o} = \frac{1}{2}\left(l_\mathrm{min} + l_\mathrm{max}\right), \qquad C^*_\mathrm{D,o} = C_\mathrm{D,o} + \frac{1}{4} C_\mathrm{D,t} \frac{d_\mathrm{t} r_\mathrm{m,o}}{A}. \tag{6.6}$$

The reel-out tether force is then

$$F_\mathrm{t,o} = \frac{1}{2}\rho A \bar{v}_\mathrm{w}^2 \gamma_\mathrm{o} \left(\cos\beta_\mathrm{o} - f_\mathrm{o}\right)^2, \tag{6.7}$$

where γ_o is the model-dependent reel-out force factor. The reel-in tether force is model dependent and is given in Section 6.3.2 and Section 6.3.3.

The phase durations are

$$t_\mathrm{o} = \frac{l_\mathrm{reel}}{v_\mathrm{o}}, \qquad t_\mathrm{i} = \frac{l_\mathrm{reel}}{|v_\mathrm{i}|}. \tag{6.8}$$

The net mechanical energy per cycle is

$$E_\mathrm{c} = \left(F_\mathrm{t,o} - F_\mathrm{t,i}\right) l_\mathrm{reel}, \tag{6.9}$$

and the average mechanical cycle power is

$$P_\mathrm{c} = \frac{E_\mathrm{c}}{t_\mathrm{o} + t_\mathrm{i}}. \tag{6.10}$$

Using the wind power density

$$P_\mathrm{w} = \frac{1}{2}\rho \bar{v}_\mathrm{w}^3, \tag{6.11}$$

the average mechanical cycle power can be written as

$$P_\mathrm{c} = p_\mathrm{c}\gamma_\mathrm{o} P_\mathrm{w} A, \tag{6.12}$$

where p_c is the normalized mechanical cycle-power factor. The optimization of the model is therefore performed by maximizing p_c.

For each wind profile, the reference wind speeds are processed in ascending order. The model starts in region 1, switches to region 2 when the reel-out tether force reaches $F_\mathrm{t,n}$, and switches to region 3 when the mechanical reel-out power reaches P_n. In region 2, the force-limit operating point defines

$$\mu_F = \frac{\bar{v}_\mathrm{w}}{\bar{v}_\mathrm{w}^{n,F}}, \qquad f_\mathrm{o} = \frac{(\mu_F - 1)\cos\beta_\mathrm{o} + f_{n,F}}{\mu_F}, \tag{6.13}$$

where $\bar{v}_\mathrm{w}^{n,F}$ and $f_{n,F}$ are the wind speed and reel-out factor at the onset of the tether-force constraint. In region 3, the power-limit operating point defines

$$\mu_P = \frac{\bar{v}_\mathrm{w}}{\bar{v}_\mathrm{w}^{n,P}}, \qquad f_\mathrm{o} = \frac{f_{n,P}}{\mu_P}, \tag{6.14}$$

where $\bar{v}_\mathrm{w}^{n,P}$ and $f_{n,P}$ are the wind speed and reel-out factor at the onset of the mechanical power limit. In this region, the reel-out speed is fixed and the reel-out force is kept at $F_\mathrm{t,n}$, so the mechanical reel-out power is clamped at P_n.

6.3.2. Original Luchsinger model [10]

The original variant uses the reel-out and reel-in force factors

$$\gamma_\mathrm{o} = \frac{C_\mathrm{L,o}^3}{\left(C_\mathrm{D,o}^*\right)^2}, \qquad \gamma_\mathrm{i} = C_\mathrm{D,i}. \tag{6.15}$$

The reel-out force factor uses the corrected drag coefficient introduced in Equation 6.6, while the drag-based reel-in formulation of the original model is retained.

The reel-in elevation angle β_i is a prescribed input, and the reel-in factor is

$$-\frac{v_\mathrm{i,max}}{\bar{v}_\mathrm{w}} \leq f_\mathrm{i} < 0. \tag{6.16}$$

The reel-in tether force becomes

$$F_\mathrm{t,i} = \frac{1}{2}\rho A \bar{v}_\mathrm{w}^2 \gamma_\mathrm{i}\left(1 - 2f_\mathrm{i}\cos\beta_\mathrm{i} + f_\mathrm{i}^2\right). \tag{6.17}$$

In region 1, both reeling factors are optimized:

$$p_\mathrm{c}^{(1)} = \max_{f_\mathrm{o}, f_\mathrm{i}}\left\{\left[(\cos\beta_\mathrm{o} - f_\mathrm{o})^2 - \frac{\gamma_\mathrm{i}}{\gamma_\mathrm{o}}\left(1 - 2f_\mathrm{i}\cos\beta_\mathrm{i} + f_\mathrm{i}^2\right)\right]\left(\frac{f_\mathrm{i} f_\mathrm{o}}{f_\mathrm{i} - f_\mathrm{o}}\right)\right\}. \tag{6.18}$$

This is the unconstrained operating region. Both the reel-out and reel-in factors are free optimization variables.

In region 2, the reel-out factor is prescribed by the tether-force constraint and only the reel-in factor is optimized:

$$p_c^{(2)} = \max_{f_i} \left\{ \left[\frac{1}{\mu_F^2} \left(\cos\beta_o - f_{n,F}\right)^2 - \frac{\gamma_i}{\gamma_o} \left(1 - 2 f_i \cos\beta_i + f_i^2\right) \right] \left(\frac{f_i \left[(\mu_F - 1)\cos\beta_o + f_{n,F}\right]}{\mu_F f_i - \left[(\mu_F - 1)\cos\beta_o + f_{n,F}\right]} \right) \right\}. \tag{6.19}$$

In this region, the reel-out speed is increased with increasing wind speed to maintain the nominal tether force.

In region 3, the reel-out speed is fixed at its nominal power-limit value and only the reel-in factor is optimized:

$$p_c^{(3)} = \max_{f_i} \left\{ \left[\frac{1}{\mu_P^2} \left(\cos\beta_o - f_{n,P}\right)^2 - \frac{\gamma_i}{\gamma_o} \left(1 - 2 f_i \cos\beta_i + f_i^2\right) \right] \left(\frac{f_i f_{n,P}}{\mu_P f_i - f_{n,P}} \right) \right\}. \tag{6.20}$$

In this region, the reel-out force and mechanical reel-out power are both constrained. The reel-out side is therefore no longer optimized directly, while the reel-in factor remains the only free optimization variable.

6.3.3. Soft-kite Luchsinger model [4]

The soft-kite variant keeps the same region-based control structure, but it replaces the drag-only reel-in model by a constant lift-to-drag reel-in model. With $C_{D,o}^*$ from Equation 6.6, the reel-out and reel-in lift-to-drag ratios are

$$E_o = \frac{C_{L,o}}{C_{D,o}^*}, \qquad E_i = \frac{C_{L,i}}{C_{D,i}}. \tag{6.21}$$

The force factors become

$$\gamma_o = C_{L,o}\sqrt{1 + \frac{1}{E_o^2}} \left(1 + E_o^2\right), \qquad \gamma_i = C_{L,i}\sqrt{1 + \frac{1}{E_i^2}}. \tag{6.22}$$

The admissible reel-in factor is

$$-\min\left(\frac{v_{i,\max}}{\bar{v}_w}, \sqrt{1 + \frac{1}{E_i^2}} \right) \leq f_i < 0. \tag{6.23}$$

The reel-in tether force becomes

$$F_{t,i} = \frac{1}{2}\rho A \bar{v}_w^2 \gamma_i \frac{\left(\sqrt{1 + E_i^2\left(1 - f_i^2\right)} - f_i\right)^2}{1 + E_i^2}. \tag{6.24}$$

Unlike in the original Luchsinger model, the reel-in elevation angle is not fixed. It follows from the constant lift-to-drag ratio and the radial reel-in flight condition:

$$\beta_i = \arccos\left(\frac{\sqrt{1 + E_i^2\left(1 - f_i^2\right)} + f_i E_i^2}{1 + E_i^2} \right). \tag{6.25}$$

This relation gives the elevation angle that a kite with lift-to-drag ratio E_i can maintain during steady-state radial flight for a given reeling factor f_i. It can therefore also be used to determine the maximum achievable reel-in elevation angle for the selected depowered aerodynamic state. This formulation is more suitable for soft kites because their depowered state is constrained by the need to maintain stable positive lift.

In region 1, the objective is

$$p_c^{(1)} = \max_{f_o, f_i} \left\{ \left[\left(\cos\beta_o - f_o\right)^2 - \frac{\gamma_i}{\gamma_o} \frac{\left(\sqrt{1 + E_i^2\left(1 - f_i^2\right)} - f_i\right)^2}{1 + E_i^2} \right] \left(\frac{f_i f_o}{f_i - f_o} \right) \right\}. \tag{6.26}$$

In region 2, the objective becomes

$$p_c^{(2)} = \max_{f_i} \left\{ \left[\frac{1}{\mu_F^2} \left(\cos\beta_o - f_{n,F}\right)^2 - \frac{\gamma_i}{\gamma_o} \frac{\left(\sqrt{1+E_i^2\left(1-f_i^2\right)} - f_i\right)^2}{1+E_i^2} \right] \left(\frac{f_i\left[\left(\mu_F - 1\right)\cos\beta_o + f_{n,F}\right]}{\mu_F f_i - \left[\left(\mu_F - 1\right)\cos\beta_o + f_{n,F}\right]} \right) \right\}. \tag{6.27}$$

In region 3, the objective is

$$p_c^{(3)} = \max_{f_i} \left\{ \left[\left(\cos\beta_o - \frac{f_{n,P}}{\mu_P}\right)^2 - \frac{\gamma_i}{\gamma_o} \frac{\left(\sqrt{1+E_i^2\left(1-f_i^2\right)} - f_i\right)^2}{1+E_i^2} \right] \left(\frac{f_i f_{n,P}}{\mu_P f_i - f_{n,P}} \right) \right\}. \tag{6.28}$$

It should be noted that Equation 6.28 differs from the expression given by Schmehl et al. [4]. The equation used here is the corrected form. Using the expression as published by Schmehl et al. leads to unrealistically low cycle powers in region 3, which was identified as an error in the original formulation.

6.3.4. Comparison of the two variants

The two variants mainly differ in their aerodynamic force-factor definitions and in their treatment of the reel-in phase. In the original Luchsinger model, the reel-in phase is described by a drag-based force factor and a prescribed reel-in elevation angle, representing a fixed-wing kite that dives toward the ground station during retraction. In the soft-kite model, the reel-in phase is instead described by a prescribed depowered lift-to-drag ratio, from which the reel-in elevation angle follows. This distinction is important for soft kites, because their depowered state is constrained by the need to maintain positive lift.

Because the same tether-drag correction is applied to both variants in this implementation, the comparison below isolates the remaining model differences. On the reel-out side, the original variant uses the classical Luchsinger force-factor structure, whereas the soft-kite variant expresses the force factor through the reel-out lift-to-drag ratio. On the reel-in side, the original variant prescribes the elevation angle directly, whereas the soft-kite variant derives it from the depowered aerodynamic state.

Both model variants are compared using the same logarithmic wind profile, with a friction velocity of $0.4\ \mathrm{ms^{-1}}$ and a surface roughness of $0.03\ \mathrm{m}$. The wind speed shown on the horizontal axis of the following figures is the reference wind speed at $200\ \mathrm{m}$. In the original Luchsinger model, the reel-in elevation angle must be prescribed. For this comparison, it is set to $80°$, which is representative of a high-elevation reel-in phase for a soft kite and for both models the reel-out elevation angle is set to $30°$.

Figure 6.2 shows that the soft-kite formulation gives a slightly higher reel-out power and reel-out tether force over most of the operating range until the reel-out power reaches its limit. This is caused by the different definition of the reel-out force factor. The largest difference between the two variants occurs during reel-in. Up to a wind speed of approximately $10\ \mathrm{ms^{-1}}$, where the reel-in speed reaches its maximum, the soft-kite formulation gives much lower reel-in power and tether force than the original formulation. This is linked to the high reel-in elevation angle predicted by the soft-kite model, as shown in Figure 6.4. As a result, the two variants also produce different cycle powers. After the reel-in speed reaches its maximum value, the reel-out speeds and reeling factors also show a change in slope. This occurs because the reel-in factor can no longer reach its unconstrained optimum. Since the optimum cycle power depends on the combination of reel-in and reel-out factors, the reel-out factor also adjusts in response.

The very low reel-in tether forces predicted by the soft-kite formulation at low wind speeds should also be interpreted in light of the model constraints. In this Luchsinger implementation, no minimum tether-force constraint is imposed during reel-in. In the inertia-free quasi-steady model, such a constraint is included to ensure that the kite remains sufficiently tensioned, which is especially important for soft kites because they require a minimum load to maintain their shape and stable flight. If a similar minimum tether-force constraint were included in the Luchsinger model, the low-wind reel-in tether forces would likely become more realistic for soft-kite operation. This limitation is less relevant for fixed-wing kites, which do not rely on tether tension in the same way to maintain their structural shape.

The optimized reeling speeds and reeling factors are shown in Figure 6.3. The reel-out quantities are similar in trend, but the soft-kite model reaches the tether-force and power limits earlier because of its higher reel-out force and power. The reel-in quantities differ most strongly at low wind speeds, where the soft-kite model predicts a different reel-in state. Above approximately $10\ \mathrm{ms}^{-1}$, the reel-in factors and reel-in speeds become more comparable. The cycle times in Figure 6.4 are also relatively similar, showing that the main difference in cycle power is caused more by the force and power levels than by large differences in cycle duration.

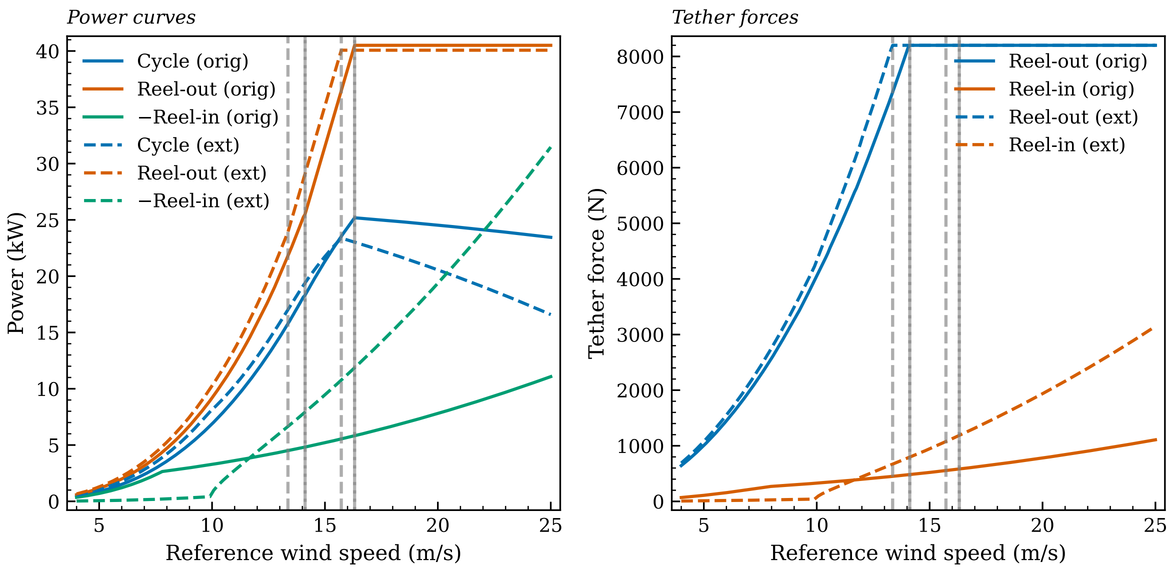

Figure 6.2: Comparison of mechanical phase power, cycle power, and tether force for the original and soft-kite Luchsinger variants. The vertical lines indicate the operating-region transitions for each model variant.

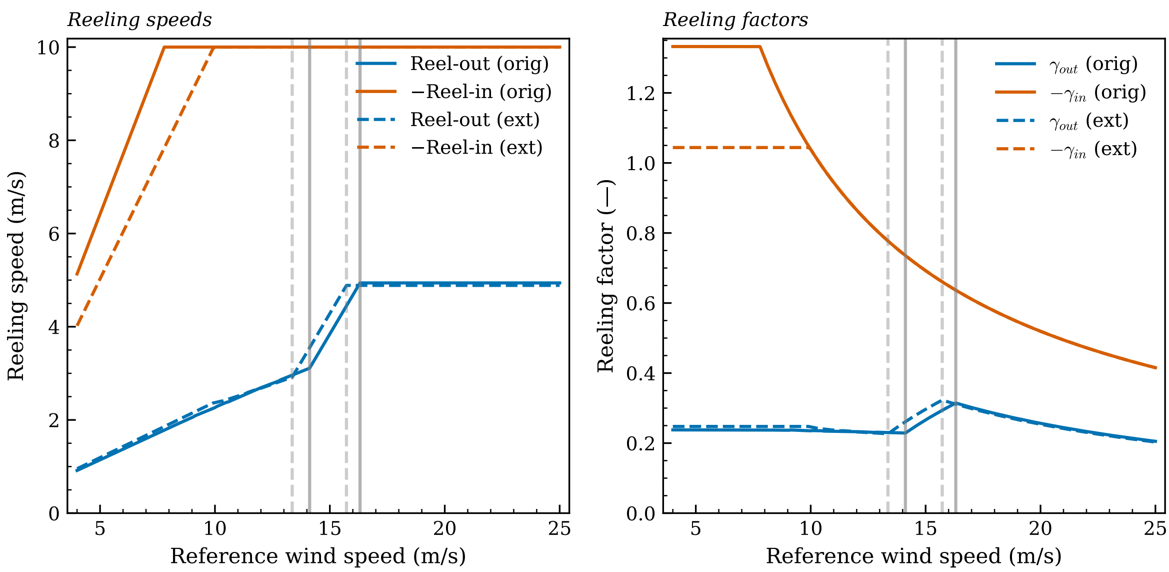

Figure 6.3: Comparison of optimized reeling speeds and dimensionless reeling factors for the original and soft-kite Luchsinger variants. The vertical lines indicate the operating-region transitions for each model variant.

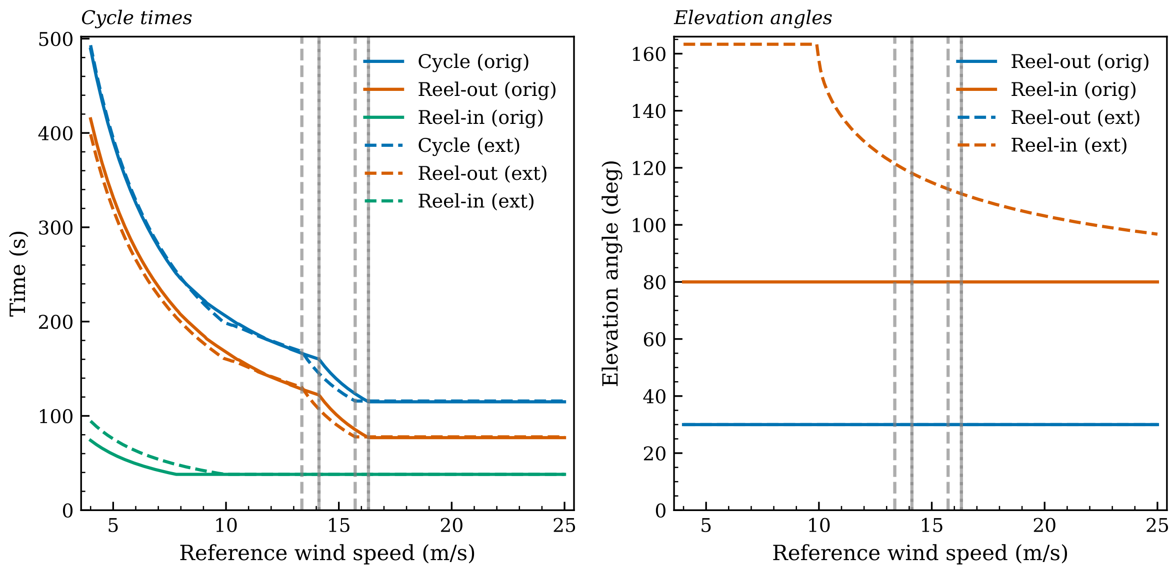

Figure 6.4: Comparison of cycle duration and reel-out and reel-in elevation angles for the original and soft-kite Luchsinger variants. The vertical lines indicate the operating-region transitions for each model variant.

6.4. Inertia-free quasi-steady model

The power estimation model used for the case studies is an inertia-free quasi-steady model (QSM) based on the pumping-kite model of Van der Vlugt et al. [16] and on the code implementation developed by Schelbergen [5], [12]. In this thesis, the existing code is refactored for use in AWESPA, the model input and output structure is adapted to awesIO, and changes are made to the cycle representation, the regime-3 handling, and the optimization procedure. The main additions are the implementation of an explicit reel-out to reel-in transition, the option to use a variable elevation angle during traction, the relocation of regime-3 handling from the optimizer to the core QSM code, and a constrained optimization procedure for power-curve generation.

6.4.1. Model formulation

Mass and tether representation

The term inertia-free means that translational acceleration states are not included in the equations of motion. This does not mean that mass is fully neglected. The kite mass and tether mass contribute through gravitational loading in the static force balance.

The kite is represented as a point mass attached to a straight tether. Tether sag, tether elasticity and, structural dynamics are not resolved. The kite position is described using spherical coordinates: tether length, azimuth angle, and elevation angle. The flight direction is described by the course angle. This gives a compact description of the cycle geometry, but it does not reproduce the full three-dimensional motion of the real kite.

Aerodynamic states and tether drag

The aerodynamic model uses two kite states. The powered coefficients are used during traction and both transition phases, while the depowered coefficients are used during retraction. This means that the model does not simulate a continuous change in aerodynamic coefficients during depowering or powering. Instead, each phase is assigned one of the two aerodynamic states.

Tether drag is included through an added drag contribution that depends on tether diameter and tether length. This gives an approximate representation of tether drag without using a distributed tether model.

Wind profile and air density

The environment is represented by a one-dimensional vertical wind profile and a height-dependent air density. The wind profile is supplied as a normalized lookup table with a reference height. The local wind speed is then obtained by interpolation at the instantaneous kite altitude. Air density decreases exponentially with altitude.

Horizontal wind-direction changes, turbulence, and gust dynamics are not modeled. The model output should therefore be interpreted as deterministic performance under prescribed mean wind conditions.

Steady-state solver and cycle power

The core of the model is the steady-state solver. For a given kite position, wind speed, aerodynamic state, and control setting, the solver determines a quasi-steady flight state from the force equilibrium shown in Figure 6.5. The aerodynamic force must be consistent with the prescribed lift-to-drag ratio G. Since gravitational loading causes the kinematic ratio κ to differ from G, the solver iterates on κ until the lift-to-drag ratio obtained from the force balance matches the target value. [16]

$$\kappa = \frac{v_{\mathrm{a},\tau}}{v_{\mathrm{a},r}}. \tag{6.29}$$

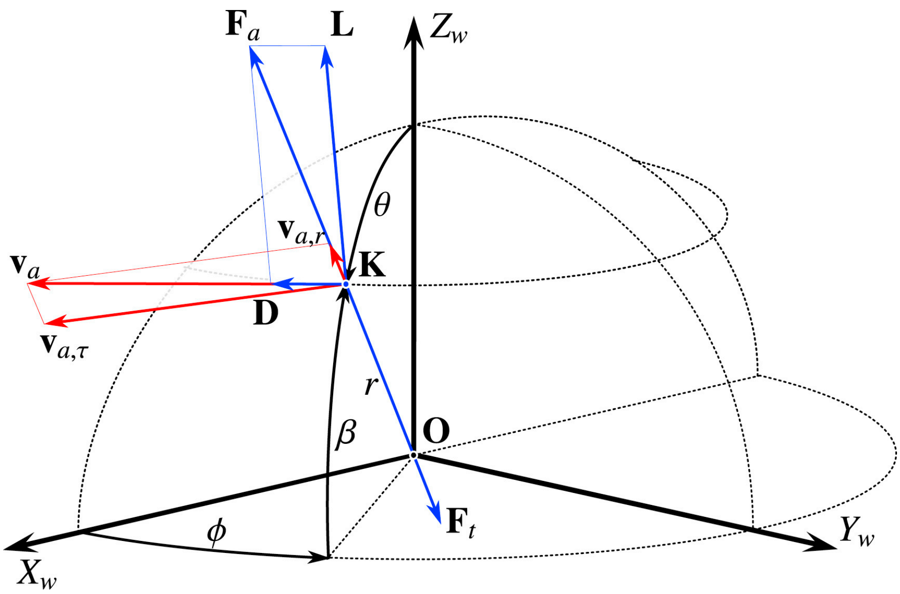

Figure 6.5: Steady force equilibrium including gravitational loading [16].

For quasi-steady motion, the tangential aerodynamic force component balances the tangential component of the lumped gravitational force:

$$F_{\mathrm{a},\theta} = -F^*_{\mathrm{g},\theta} = -\left(\frac{1}{2}m_\mathrm{t} + m\right) g \sin\theta. \tag{6.30}$$

The radial aerodynamic force component is then obtained from the aerodynamic force magnitude:

$$F_{\mathrm{a},r} = \sqrt{F_\mathrm{a}^2 - F_{\mathrm{a},\theta}^2}. \tag{6.31}$$

For a given iteration i, the apparent wind components and aerodynamic force components are first computed using the current value of κ_i. The corresponding lift-to-drag ratio is then evaluated as

$$G_i = \sqrt{\left(\frac{F_\mathrm{a} v_\mathrm{a}}{\mathbf{F}_\mathrm{a} \cdot \mathbf{v}_\mathrm{a}}\right)^2 - 1}, \tag{6.32}$$

where F_a and v_a are the magnitudes of the aerodynamic force and apparent wind velocity. Starting from $\kappa_1 = G$, the kinematic ratio is updated according to

$$\kappa_{i+1} = \kappa_i \sqrt{\frac{G}{G_i}}. \tag{6.33}$$

The iteration is repeated until G_i is sufficiently close to G. The resulting steady state provides the reeling speed, tether forces, apparent wind speed, mechanical power, and angular rates used to advance the kite position.

6.4.2. Traction phase

During traction, the kite is powered and the tether length increases until the end-of-traction tether length is reached. In real operation, this phase consists of crosswind figure-of-eight motion. The QSM does not resolve the individual turns of this trajectory. Instead, the reel-out phase is represented by an average crosswind state that approximates the mean force and velocity behavior of the figure-of-eight path.

This representative reel-out state is defined by fixed parameters. The values used in this thesis are taken from the validation study of Schelbergen and Schmehl [12]. During reel-out, the azimuth angle is set to $11.5°$, and the course angle is set to $93°$. These values are used both in this chapter and in the case studies.

The traction phase follows the same three-regime interpretation as the Luchsinger model. In regime 1, the system operates below its force and power limits, so the prescribed reel-out speed is maintained. In regime 2, the maximum tether-force limit is reached. The control setting is then changed to maximum-force control, and the steady-state solver finds an operating state with a higher reel-out speed so that power can continue to increase while the tether force remains at its limit. In regime 3, the mechanical power limit is reached in addition to the tether-force limit, so the reel-out power is constrained.

In the previous code structure, regimes 1 and 2 were handled in the QSM through the phase control settings. Regime 3, however, was handled in the optimizer. During optimization, the optimizer checked whether the reel-out power exceeded the maximum admissible value. If the limit was exceeded, the optimizer increased the traction elevation angle until the power limit was no longer violated.

In the present implementation, this regime-3 handling is moved from the optimizer to the core QSM code. The QSM now checks the mechanical power limit directly during the traction calculation. If the power limit is exceeded, the reel-out speed is first reduced to

$$v_\mathrm{r} = \frac{P_\mathrm{max}}{F_\mathrm{max}}, \tag{6.34}$$

subject to the maximum reeling-speed limit. If this is not sufficient to satisfy the power or tether-force constraint, the model increases the elevation angle until a feasible operating point is found. Moving this step into the core QSM means that power-limited operation is treated as part of the model rather than as a separate optimizer correction. In real systems, staying below the power limit can also be done via depowering the kite or flying to the side of the wind window.

6.4.3. Retraction phase

During retraction, the kite is depowered and the tether length decreases until the end-of-retraction tether length is reached. The retraction phase is represented by a radial return state rather than by a crosswind figure-of-eight motion. The azimuth angle is set to $0°$, and the course angle is set to $180°$, following the same validation setup [12]. This represents flying toward zenith.

Normally, retraction control is speed based and uses a prescribed negative reel-in speed. If this setting violates the tether force limits, the system will be controlled by setting a tether force as control setting.

6.4.4. Transition phases

Two transition phases connect traction and retraction. The reel-out to reel-in transition is added in this thesis. It starts at the final traction elevation and raises the kite to the elevation at which retraction begins. With this transition included, the model no longer switches directly from powered traction to depowered retraction.

The reel-in to reel-out transition lowers the kite back toward the traction operating condition.

Because the simulation starts with the reel-out to reel-in transition, the model first has to determine the elevation angle at which the preceding traction phase ended. This value can differ from the prescribed traction elevation angle, especially in regime 3, where the elevation angle may be increased to satisfy the power limit. The same consistency requirement applies to the reel-in to reel-out transition: the code has to determine the elevation angle at the start of the following traction phase so that the transition connects to the actual traction state rather than to the nominal setpoint.

The transition phases are speed controlled with a nominal zero reeling-speed setpoint. If the resulting tether force is outside the admissible range, the tether force is clamped to the corresponding limit.

6.4.5. Optimization framework

The pumping cycle depends on several operating settings, so an optimizer is used to find a feasible cycle with high average power for each wind speed. The optimizer is based on Sequential Least Squares Programming (SLSQP). The objective is to maximize the average cycle power. The optimized variables are the reel-out and reel-in speeds, the tether-length fractions, the traction elevation angle, and the end elevation of the reel-out to reel-in transition. The variables and their constraints used in this study are listed in Table 6.2.

A new feature in this implementation is that the traction elevation angle can vary during the reel-out phase. Instead of prescribing one constant elevation angle, the user can provide an array of elevation-angle setpoints. For example, if 4 setpoints are used, the traction phase is divided into 3 equally spaced tether-length intervals. The setpoints define the elevation angle at the interval boundaries, so each interval has a start and end elevation angle. These setpoints are included in the optimization. This allows the optimizer to trade off lower-elevation operation, which reduces cosine losses, against higher-elevation operation, where the wind speed is usually higher. In this work, 4 setpoints are used to limit the computational cost while still allowing variation in the traction elevation profile.

The optimization is subject to constraints that ensure a feasible pumping cycle. These constraints enforce a minimum reel-in length, a minimum altitude, and a maximum tether length. They also ensure that the reel-out to reel-in transition connects consistently to the traction phase. In particular, the transition end elevation angle must not be below the traction elevation angle. When several traction elevation angles are optimized, the difference between consecutive setpoints is also bounded to avoid unrealistically sharp changes in the elevation profile. The coupled constraints are summarized in Table 6.3. During power-curve generation, the optimizer also accounts for the active operating regime. If the traction phase is fully force limited, the reel-out speed is removed from the decision vector because the phase is then controlled by tether force. If the traction phase is fully power limited, the traction elevation angle is removed because the required elevation angle is determined directly by the core QSM.

Table 6.2: Decision variables and constraints used in the QSM optimization.

Parameter	Symbol	Constraint
Traction reeling speed	v_{tr}	$0.01 \leq v_{\mathrm{tr}} \leq 15.0\,\mathrm{m\,s^{-1}}$
Retraction reeling speed	v_{ret}	$-15.0 \leq v_{\mathrm{ret}} \leq -0.01\,\mathrm{m\,s^{-1}}$
Tether-length fraction at end of retraction	$f_{\mathrm{ret,end}}$	$0.20 \leq f_{\mathrm{ret,end}} \leq 0.80$
Tether-length fraction at end of traction	$f_{\mathrm{tr,end}}$	$0.50 \leq f_{\mathrm{tr,end}} \leq 0.95$
Traction elevation angle	$\beta_{\mathrm{tr},i}$	$20^\circ \leq \beta_{\mathrm{tr},i} \leq 80^\circ$
Reel-out to reel-in transition end elevation angle	$\beta_{\mathrm{RORI,end}}$	$30^\circ \leq \beta_{\mathrm{RORI,end}} \leq 80^\circ$

Table 6.3: Constraints used in the QSM optimization.

Constraint	Mathematical form
Minimum reel-in amount	$f_{\mathrm{tr,end}} - f_{\mathrm{ret,end}} \geq 0.10$
Minimum altitude	$h_{\min} \geq 100\ \mathrm{m}$
Maximum tether length	$l_{\max} \leq l_{\mathrm{tether,max}}$
Transition end elevation angle must not be below the traction elevation angle	$\beta_{\mathrm{RORI,end}} \geq \beta_{\mathrm{tr}}$
Maximum change between consecutive traction elevation angles	$\lvert\beta_{\mathrm{tr},i+1} - \beta_{\mathrm{tr},i}\rvert \leq 10^\circ$

6.4.6. Optimized operation for representative wind profiles

The optimizer is tested with a wind resource file containing 8 representative wind profiles. These profiles are the same profiles used in case study 1 and are shown in Figure 8.1. The simulation settings use a minimum height of $100\ \mathrm{m}$, the decision-variable bounds from Table 6.2, and the coupled constraints from Table 6.3. It should be noted that this minimum-height constraint is applied to the representative QSM trajectory. Since the figure-of-eight motion is not resolved, local altitude variations within the actual crosswind path are not captured, and parts of a real figure-of-eight trajectory could fall below this representative height. The power curves are simulated from $4\ \mathrm{m\,s^{-1}}$ to $25\ \mathrm{m\,s^{-1}}$ with a step size of $1\ \mathrm{m\,s^{-1}}$. For all profiles, the cut-in wind speed was found to be $5\ \mathrm{m\,s^{-1}}$.

The wind-profile shape affects both the rated wind speed and the maximum cycle power. Figure 6.6 shows the power curves for the 8 profiles. The maximum rated cycle power is found for profile 4, with $14.687\ \mathrm{kW}$ at $16\ \mathrm{m\,s^{-1}}$. The lowest rated cycle power is found for profile 3, with $13.407\ \mathrm{kW}$ at $18\ \mathrm{m\,s^{-1}}$. The absolute difference between these two profiles is $1.280\ \mathrm{kW}$, corresponding to a $9.4\ \%$ difference.

Profile 3 gives a lower rated cycle power than profile 4 because of the difference in wind shear. Profile 3 has relatively high shear, while profile 4 has the lowest shear. Since the reference wind speed is defined at $200\ \mathrm{m}$, and the kite mostly operates below this height during traction, profile 3 encounters a lower wind speed during the traction phase than profile 4 for the same reference wind speed. During retraction, the kite can operate above $200\ \mathrm{m}$, where profile 3 gives a higher wind speed. This combination reduces the generated traction power and increases the retraction penalty, which explains the lower rated cycle power of profile 3.

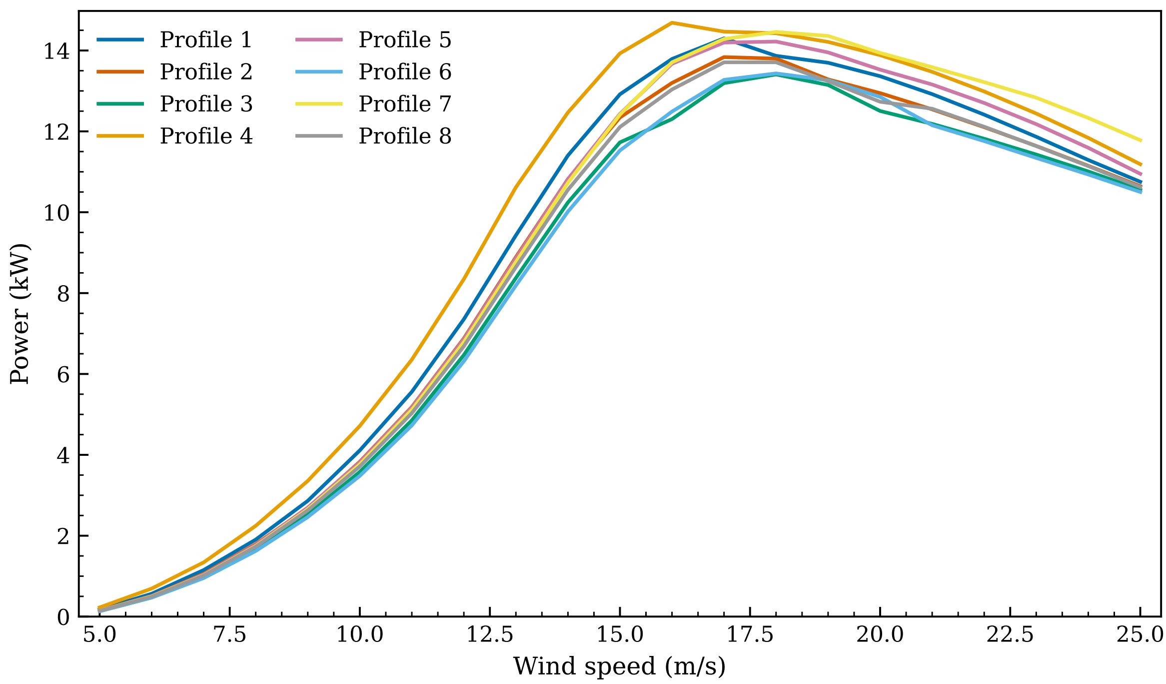

Figure 6.6: Cycle power curves for the 8 representative wind profiles using the TU Delft V3 kite. The wind speed is the reference wind speed at $200\ \mathrm{m}$.

The optimized trajectories at $12\ \mathrm{m\,s^{-1}}$ show how the optimizer responds to differences in wind shear. Figure 6.7 shows that profile 4 uses the lowest starting tether length and keeps a lower elevation angle than the other profiles. For this low-shear profile, flying higher gives a smaller wind-speed benefit, so lower-elevation operation is favored because it reduces cosine losses. Profile 6 shows the highest starting tether length, which is consistent with higher-shear operation where the wind-speed increase with altitude is larger.

The trajectories also show a common pattern across the profiles. Most profiles operate near the minimum allowed altitude during the early part of traction. After a certain elevation angle is reached, the elevation angle is no longer reduced further. At that point, the wind-speed gain from flying higher balances the reduction in cosine losses from flying lower.

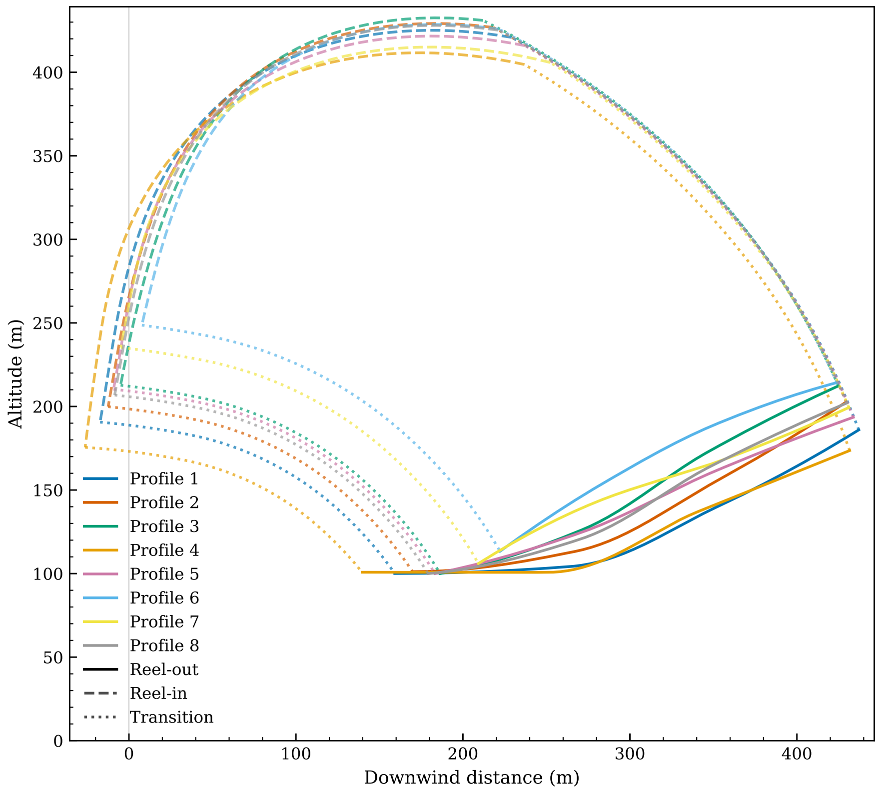

Figure 6.7: Optimized kite trajectories at $12\ \mathrm{m\,s^{-1}}$ for the 8 representative wind profiles. The minimum allowed height is $100\ \mathrm{m}$.

The optimized trajectories for profile 1 show how the operating strategy changes with wind speed. Figure 6.8 shows the trajectories, while Figure 6.9 shows the corresponding optimized cycle parameters. At the cut-in wind speed of $5\ \mathrm{m\,s^{-1}}$, the optimized starting tether length is relatively large. This allows the kite to operate at a higher elevation angle while remaining above the minimum height constraint. At this low wind speed, the kite also starts retraction shortly after the reel-out to reel-in transition.

The intermediate wind-speed cases show a gradual change in the end of the traction phase. At $7\ \mathrm{m\,s^{-1}}$, $9\ \mathrm{m\,s^{-1}}$, and $11\ \mathrm{m\,s^{-1}}$, the elevation angle increases near the end of traction. This can reduce the required elevation change during the transition and retraction phases. It also avoids ending the traction phase at a low elevation angle where the instantaneous power would be higher.

The high wind-speed cases are controlled by the power limit. At $18\ \mathrm{m\,s^{-1}}$ and $22\ \mathrm{m\,s^{-1}}$, the elevation angle is set by regime-3 operation, in which the QSM increases the elevation angle until the power limit is satisfied. At $22\ \mathrm{m\,s^{-1}}$, the optimized maximum tether length is reduced. At this high elevation angle, a longer tether would increase the altitude and therefore the wind speed. This gives no extra benefit during traction because the system is already power limited, while it can increase the required retraction power.

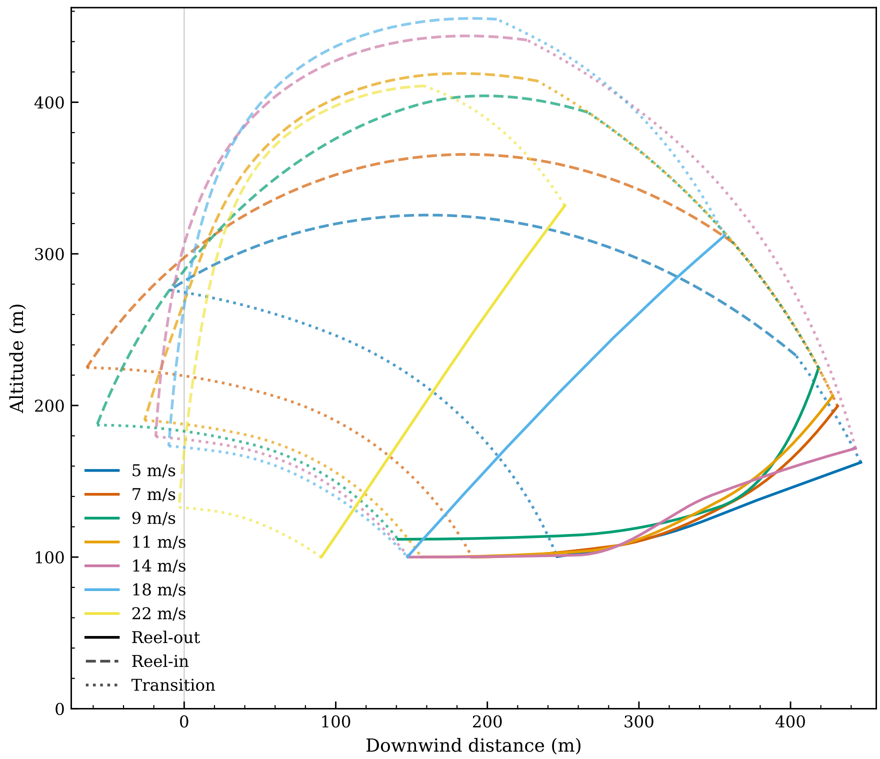

Figure 6.8: Optimized kite trajectories for profile 1 at different reference wind speeds.

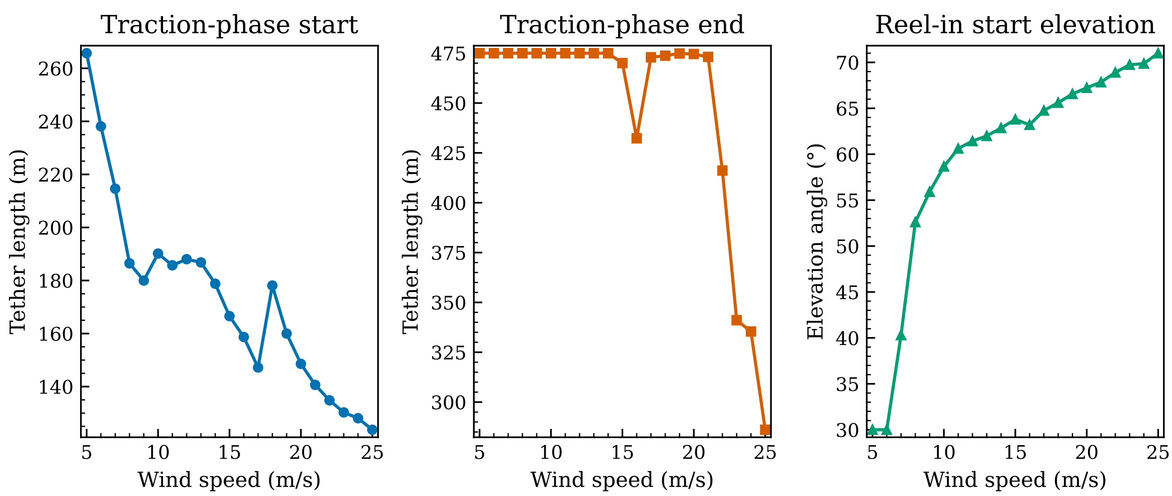

Figure 6.9: Optimized cycle parameters for profile 1 as a function of reference wind speed.

6.5. Comparison between Luchsinger and the inertia-free QSM

To compare the soft-kite Luchsinger model with the inertia-free QSM, the QSM settings are chosen to match the Luchsinger setup as closely as possible. The traction elevation angle is kept constant at $30°$, and the minimum tether length is set to $120\ \mathrm{m}$ in both models. All other system parameters are taken from Table 6.1. In the inertia-free QSM, only the reel-out speed, reel-in speed, and retraction-start elevation angle are optimized. The other optimizer variables are fixed to match the soft-kite Luchsinger setup.

Figure 6.10 compares the power curves and cycle energies predicted by both models. The Luchsinger model reaches rated power at $15.72\ \mathrm{m\,s^{-1}}$, where its cycle power is $64\ \%$ higher than the QSM cycle power at the same wind speed. The rated wind speed of the QSM is above $17\ \mathrm{m\,s^{-1}}$. Since the QSM includes transition phases, their influence is shown by comparing the solid and dashed blue lines in Figure 6.10. The difference between these two lines is small, so the transition phases are not the main cause of the difference in cycle power between the models.

The average reel-out power of the QSM is lower than that of the Luchsinger model until the power limit is reached. The QSM also reaches this limit at a higher wind speed. This difference can partly be explained by the inclusion of airborne and tether mass in the QSM. In the QSM, part of the aerodynamic force is used to balance gravitational loading, which affects the quasi-steady force balance and reduces the reel-out force and power. The cut-in behavior also differs between the models. In this comparison, the QSM starts producing power at approximately $5.5\ \mathrm{m\,s^{-1}}$, whereas the Luchsinger model starts at the prescribed cut-in wind speed of $4\ \mathrm{m\,s^{-1}}$. This lower cut-in wind speed is possible in the Luchsinger model because the model does not include mass or a corresponding gravitational feasibility constraint.

The largest difference between the models occurs during retraction, as shown in Figure 6.11 and Figure 6.12. The QSM predicts a much lower average reel-in speed, which increases the reel-in time. It also predicts a higher average reel-in power than the Luchsinger model. Together, the longer duration and higher power consumption lead to a much larger energy consumption during retraction. Although gravity can reduce the required reel-in force, this effect does not offset the difference caused by the idealized Luchsinger reel-in formulation. In the Luchsinger model, the reel-in phase reaches very high elevation angles with low tether forces and low power consumption, which leads to a lower predicted retraction energy.

Overall, the comparison shows that the Luchsinger model gives a higher and more optimistic power estimate for this setup, mainly because of its retraction phase and the absence of gravitational loading.

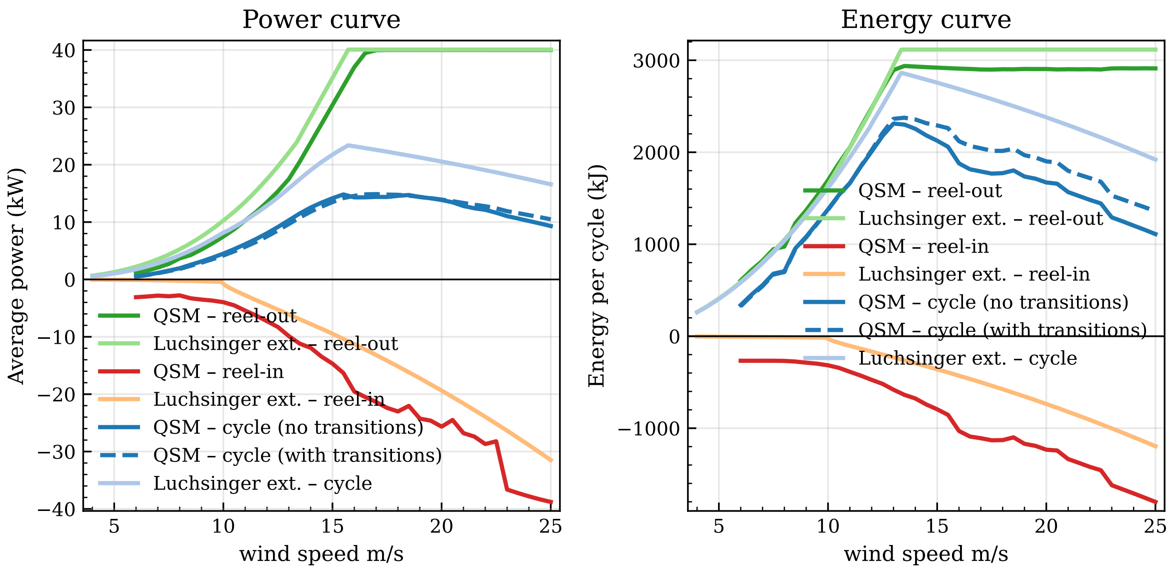

Figure 6.10: Comparison of the predicted power curves and cycle energies for the soft-kite Luchsinger model and the inertia-free QSM. The reel-out and reel-in powers are phase-averaged values.

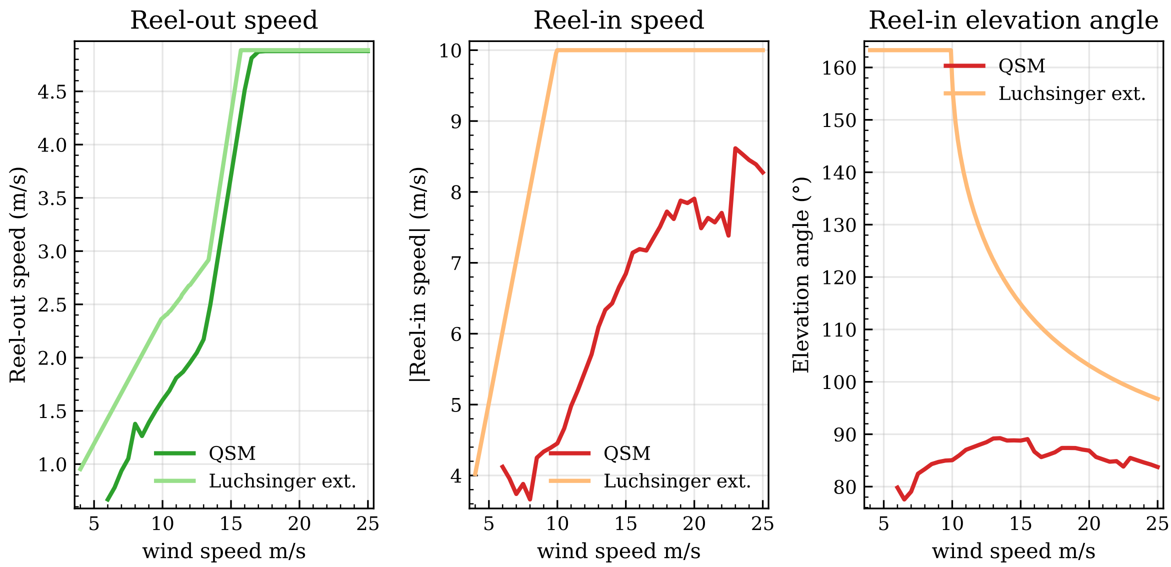

Figure 6.11: Comparison of optimized reeling speeds and average reel-in elevation angle for the soft-kite Luchsinger model and the inertia-free QSM.

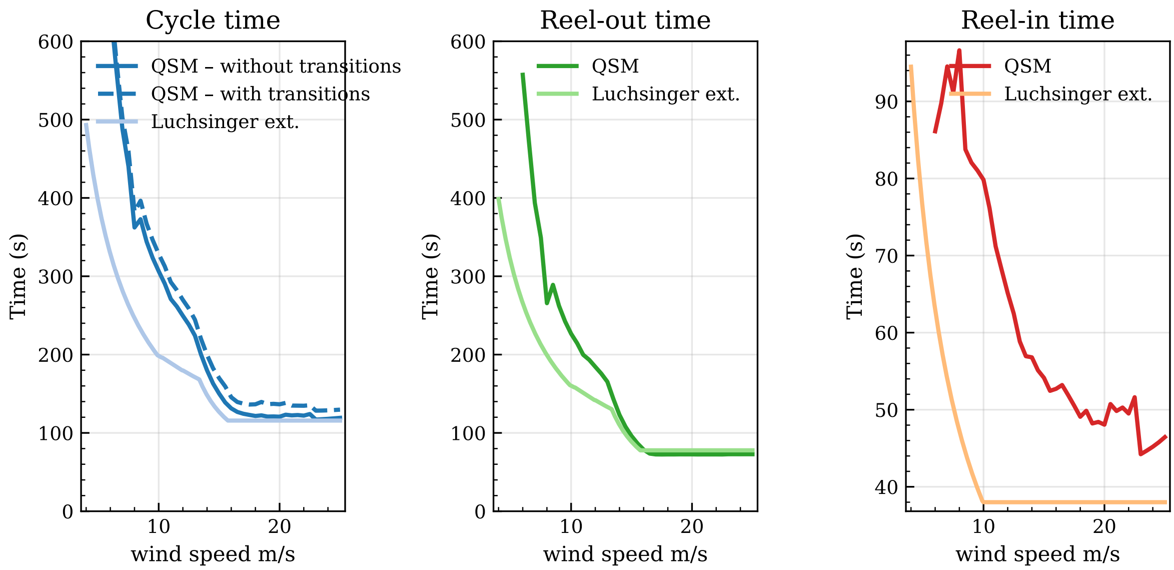

Figure 6.12: Comparison of cycle duration and phase durations for the soft-kite Luchsinger model and the inertia-free QSM.

6.6. Conclusion

This chapter described the power-estimation models implemented in AWESPA and the TU Delft V3 kite configuration used for the power-curve calculations. Two models were considered: the Luchsinger model and the inertia-free quasi-steady model. The Luchsinger model was mainly used as a simple implementation case to test the awesIO-based workflow and the modular structure of AWESPA. The inertia-free QSM was selected as the main model for the power-production estimates in the case studies.

The Luchsinger implementation included both the original fixed-wing formulation and the soft-kite formulation. Tether drag was added to the original formulation to make the comparison between the two variants more consistent. The comparison showed that the soft-kite formulation mainly changes the reel-in phase, because the reel-in elevation angle follows from the depowered lift-to-drag ratio instead of being prescribed directly.

The inertia-free QSM was refactored for use in AWESPA and extended in several ways. The implemented extensions include a reel-out to reel-in transition, variable traction elevation angles, regime-3 handling inside the core QSM, and a constrained optimizer for power-curve generation.

The optimized QSM results showed that wind-profile shape affects both the power curve and the optimized trajectory. For the 8 representative profiles, profile 4 gave the highest rated cycle power, $14.687\ \mathrm{kW}$, while profile 3 gave the lowest rated cycle power, $13.407\ \mathrm{kW}$. This difference was linked to the vertical wind shear. Since the reference wind speed is defined at $200\ \mathrm{m}$, high-shear profiles can give lower wind speeds during traction below this height and higher wind speeds during retraction above this height. This shows that the vertical wind profile should be included in AWE performance assessment instead of using only a single reference wind speed.

The comparison between the soft-kite Luchsinger model and the inertia-free QSM showed that the Luchsinger model gives a higher power estimate for the same system setup. At the Luchsinger rated wind speed of $15.72\ \mathrm{m\,s^{-1}}$, its cycle power was $64\ \%$ higher than the QSM cycle power. The transition phases in the QSM had only a small effect on the cycle power, so they were not the main cause of this difference. The difference mainly came from the treatment of mass and the reel-in phase. The QSM includes gravitational loading and predicts lower reel-in speeds, longer reel-in times, and higher retraction energy consumption than the idealized Luchsinger formulation.

7 Asynchronous wind-aligned farm operation

Previous studies on AWE farm layouts generally require closely spaced wind-aligned systems to operate in sync when their reel-in and reel-out regions can interfere [27], [32], [33], [34]. This synchronization reduces collision risk, but it also makes the power cycles of neighboring systems occur at the same time. As a result, the farm power output can have larger differences between the maximum and minimum power during a cycle.

This chapter investigates whether this synchronization requirement can be relaxed. A geometric maximum phase offset is first derived for two wind-aligned systems with adjacent reel-in and reel-out zones. This offset gives the largest allowed timing difference before the reel-out envelope of the upstream system reaches the reel-in envelope of the downstream system. The derived limit is then used as a constraint in a farm phase optimization. The optimization shifts the power cycles of the systems to reduce the peak-to-peak farm power variation while keeping the phase offsets within the geometric limit.

The method is applied to two cases. The first case uses a two-system farm to show the effect of the phase constraint for a single wind-aligned pair. The second case uses a 10-system farm to compare wind-aligned and wind-perpendicular layouts. These cases are used to assess whether constrained asynchronous operation can reduce farm power variation without changing the farm spacing.

7.1. Maximum phase offset derivation

The aim of this section is to derive a maximum phase offset for a two-system wind-aligned farm layout. The phase offset is chosen such that the systems can operate out of sync while avoiding collisions between kites and tethers. The downstream system, system 2, starts its reel-in phase at $t = 0$. The upstream system, system 1, starts its reel-in phase earlier, at $t = -t_s$. The time t_s is the time required for system 1 to reel out from $R_{\min}$ to $R_{\min} + s$, where s is the reel-out distance before the reel-out envelope of system 1 starts to overlap with the reel-in envelope of system 2. The geometry used for this derivation is shown in Figure 7.1.

Using the sine rule for the blue triangle in Figure 7.1 gives

$$\frac{R_{\min} + s}{\sin\gamma} = \frac{d}{\sin\theta} \tag{7.1}$$

Rearranging Equation 7.1 gives the overlap distance

$$s = \frac{d \cdot \sin\gamma}{\sin\theta} - R_{\min} \tag{7.2}$$

The angle γ is defined as

$$\gamma = 180^\circ - (\beta_i + \Delta\beta_i) \tag{7.3}$$

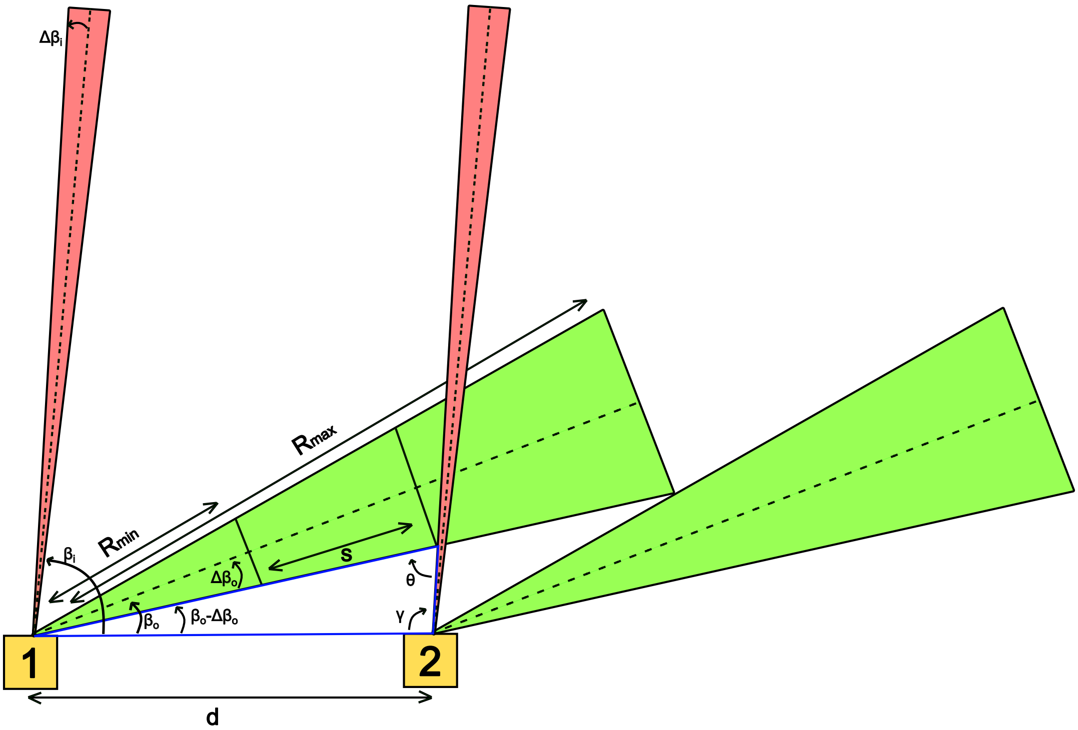

Figure 7.1: Geometry used to derive the maximum phase offset for two wind-aligned systems with overlapping reel-in and reel-out envelopes.

where β_i is the center of the reel-in cone, taken as the maximum elevation angle during reel-in, and $\Delta\beta_i$ is the reel-in half-cone width. The remaining angle θ is defined as

$$\theta = 180^\circ - (\beta_0 - \Delta\beta_0) - \gamma \tag{7.4}$$

where β_0 is the center of the reel-out cone, taken as the minimum elevation angle during reel-out, and $\Delta\beta_0$ is the reel-out half-cone width. Substituting Equation 7.3 into Equation 7.4 gives

$$\theta = \beta_i + \Delta\beta_i - \beta_0 + \Delta\beta_0 \tag{7.5}$$

The time required for system 1 to cover the distance s during reel-out is

$$t_s = \frac{s}{v_{\text{out}}} \tag{7.6}$$

where v_{out} is the reel-out speed. Substituting Equation 7.2, Equation 7.3, and Equation 7.5 into Equation 7.6 gives

$$t_s = \frac{1}{v_{\text{out}}} \left(\frac{d \cdot \sin(180^\circ - \beta_i - \Delta\beta_i)}{\sin(\beta_i + \Delta\beta_i - \beta_0 + \Delta\beta_0)} - R_{\text{min}} \right) \tag{7.7}$$

This value gives the geometric time offset between the start of reel-in for the upstream and downstream systems before the reel-out envelope of the upstream system reaches the reel-in envelope of the downstream system. It is therefore used as the upper limit on asynchronous operation in the wind-aligned case.

Dividing this geometric time offset by the total cycle time gives the geometric phase difference,

$$\Delta\psi_{\text{geom}} = \frac{t_s}{t_{\text{o}} + t_{\text{i}}}, \tag{7.8}$$

where t_{o} is the reel-out time and t_{i} is the reel-in time. Substituting Equation 7.7 into Equation 7.8 gives

$$\Delta\psi_{\text{geom}} = \frac{1}{v_{\text{out}} \cdot (t_{\text{o}} + t_{\text{i}})} \left(\frac{d \cdot \sin(180^\circ - \beta_i - \Delta\beta_i)}{\sin(\beta_i + \Delta\beta_i - \beta_0 + \Delta\beta_0)} - R_{\text{min}} \right). \tag{7.9}$$

7.2. Assumptions and setup

Several assumptions are used in the phase optimization. First, the widths of the reel-in and reel-out cones are kept constant for all considered wind speeds. The cone locations are taken from the cycle time history of each wind speed. The geometric phase constraint is therefore recomputed for each wind speed from the corresponding time history, while the cone widths remain fixed.

Second, the reel-out and reel-in times are taken from the power curve data. The reel-in time includes the reel-in phase and both transition phases. This adds a conservative margin because the transition from reel-in to reel-out, during which the kite moves toward a lower elevation angle, is treated as if it is operating at that high reel-in elevation angle while the kite is actually already at a lower elevation angle. The total cycle time is therefore written as the sum of the reel-out time and the grouped reel-in time.

Third, all systems are assumed to have the exact same single-system power cycle for a given wind speed. Different systems are represented by applying phase shifts to this periodic power signal. The power signal is interpolated on a uniform time grid and treated as periodic. The farm power signal is then obtained by summing the shifted power signals of all systems. The power signal is purely mechanical power, so battery losses and generator efficiencies are not included.

The time histories are produced with the inertia-free QSM described in Chapter 6. The simulations use the Kitepower V3 kite described in Section 6.2. A logarithmic wind profile is used with friction velocity $u^* = 0.4\ \mathrm{m\,s^{-1}}$, roughness length $z_0 = 0.03\ \mathrm{m}$, and reference height $200\ \mathrm{m}$.

The systems are placed at a distance d from each other using the longitudinal spacing relation in Equation 2.2. The spacing is generated by checking the considered wind speeds and using the most restrictive reel-out geometry. This ensures that the reel-out cones do not overlap, as illustrated in Figure 2.2.

The half-cone widths are set to $10°$ during reel-out and $2.5°$ during reel-in. Bertozzi et al. [30] studied power estimation for the TU Delft V3 system and found a maximum elevation-angle deviation of $4.3°$ from the mean elevation angle of $35°$. The reel-out half-cone width is therefore set larger than this value to include a margin for collision avoidance. The reel-in half-cone width of $2.5°$ is used as an additional buffer around the reel-in trajectory.

An additional safety margin of $10\ \mathrm{s}$ is included in the phase constraint to reduce the risk of collision. This margin is subtracted from the available time offset before converting the offset to a phase difference. The maximum allowed phase difference used in the optimization is therefore

$$\Delta\psi_{\max} = \frac{t_{\mathrm{s}} - t_{\mathrm{safe}}}{t_{\mathrm{o}} + t_{\mathrm{i}}}, \tag{7.10}$$

where t_{s} is the geometric time offset from Equation 7.7, $t_{\mathrm{safe}} = 10\ \mathrm{s}$ is the safety margin, t_{o} is the reel-out time, and t_{i} is the grouped reel-in time. If $t_{\mathrm{s}} - t_{\mathrm{safe}}$ is negative, $\Delta\psi_{\max}$ is set to zero.

7.3. Phase optimization

The phase optimization assigns a phase offset to each system in the farm. The phase offset is written as a fraction of the pumping cycle, so that a value of 0 corresponds to the reference timing and a value of 1 corresponds to one full cycle. The first system in the first wind-aligned row is fixed at phase 0.

For wind-aligned rows, the allowed phase shift between neighboring systems is limited by the maximum phase difference $\Delta\psi_{\max}$. For each adjacent pair in a wind-aligned row, the directed phase advance from the upwind system i to the downstream system j must satisfy

$$(\psi_j - \psi_i) \bmod 1 \leq \Delta\psi_{\max}. \tag{7.11}$$

Here, ψ_i and ψ_j are the phase offsets of the upwind and downstream systems, respectively. The modulo operation accounts for the periodic nature of the pumping cycle. Candidate phase combinations that violate this constraint are penalized in the optimization.

The objective is to reduce the peak-to-peak variation of the total farm power signal,

$$J = \max\left(P_{\mathrm{farm}}(t)\right) - \min\left(P_{\mathrm{farm}}(t)\right). \tag{7.12}$$

The peak-to-peak metric is used, this directly captures the amplitude of the peaks and troughs that determine the required power balancing capacity at the farm level.

The synchronized case, where all phase offsets are zero, is used as the reference case. The optimized asynchronous case is compared with this reference case to evaluate the reduction in farm power variation. The optimization is performed globally for all system phase offsets at the same time. For constrained wind-aligned layouts, the phase offsets are optimized using differential evolution. For unconstrained layouts, such as wind-perpendicular layouts without wind-aligned row constraints, the systems are initialized with phase offsets that are approximately evenly distributed over the cycle and then refined with a bounded local optimization.

The optimization is performed for $7\,\mathrm{m\,s^{-1}}$, $15\,\mathrm{m\,s^{-1}}$, and $25\,\mathrm{m\,s^{-1}}$ to represent low, medium, and high wind-speed conditions. These cases are selected because the reel-in and reel-out elevation angles and phase durations vary strongly with wind speed.

The following sections apply this method to two farm cases. First, a two-system farm is analyzed to show how the geometric phase constraint affects a single wind-aligned pair. Second, a 10-system farm is analyzed to compare wind-aligned and wind-perpendicular layouts. These cases are used to evaluate how much the farm output can be smoothed when the phase offsets are restricted by wind-aligned spacing constraints.

7.4. Two-system wind-aligned and wind-perpendicular layouts

Figure 7.2 compares the power output of a two-system farm for wind-perpendicular and wind-aligned layouts. For wind speeds of $7\,\mathrm{m\,s^{-1}}$ and $15\,\mathrm{m\,s^{-1}}$, the wind-perpendicular layout has a lower peak-to-peak power variation than the wind-aligned layout. At $25\,\mathrm{m\,s^{-1}}$, both layouts give the same result.

This difference can be explained by the operational cones shown in Figure 7.3. The reel-out elevation angle, reel-in elevation angle, and tether length range are optimized by the power estimation model for each wind speed. The resulting reel-in and reel-out cones therefore differ between wind speeds. The reel-out and reel-in times also change with wind speed. Together, these effects determine the maximum allowed phase offset for the wind-aligned layout. At some wind speeds, this constraint prevents the wind-aligned systems from reaching the phase offset that would minimize the peak-to-peak power variation. In the wind-perpendicular layout, this constraint does not apply in the same way, so the systems can be shifted closer to the unconstrained optimum.

The remaining peak-to-peak variation should therefore be interpreted as the result of the phase optimization only. The individual system cycles are not re-optimized for farm-level smoothing. Further reductions could be possible if the systems were allowed to operate with different tether lengths, elevation angles, or reeling speeds, so that their cycle times and phase constraints better match the desired farm-level timing. In this thesis, the power model and the farm optimization are used sequentially, and no coupled optimization between the two models is performed.

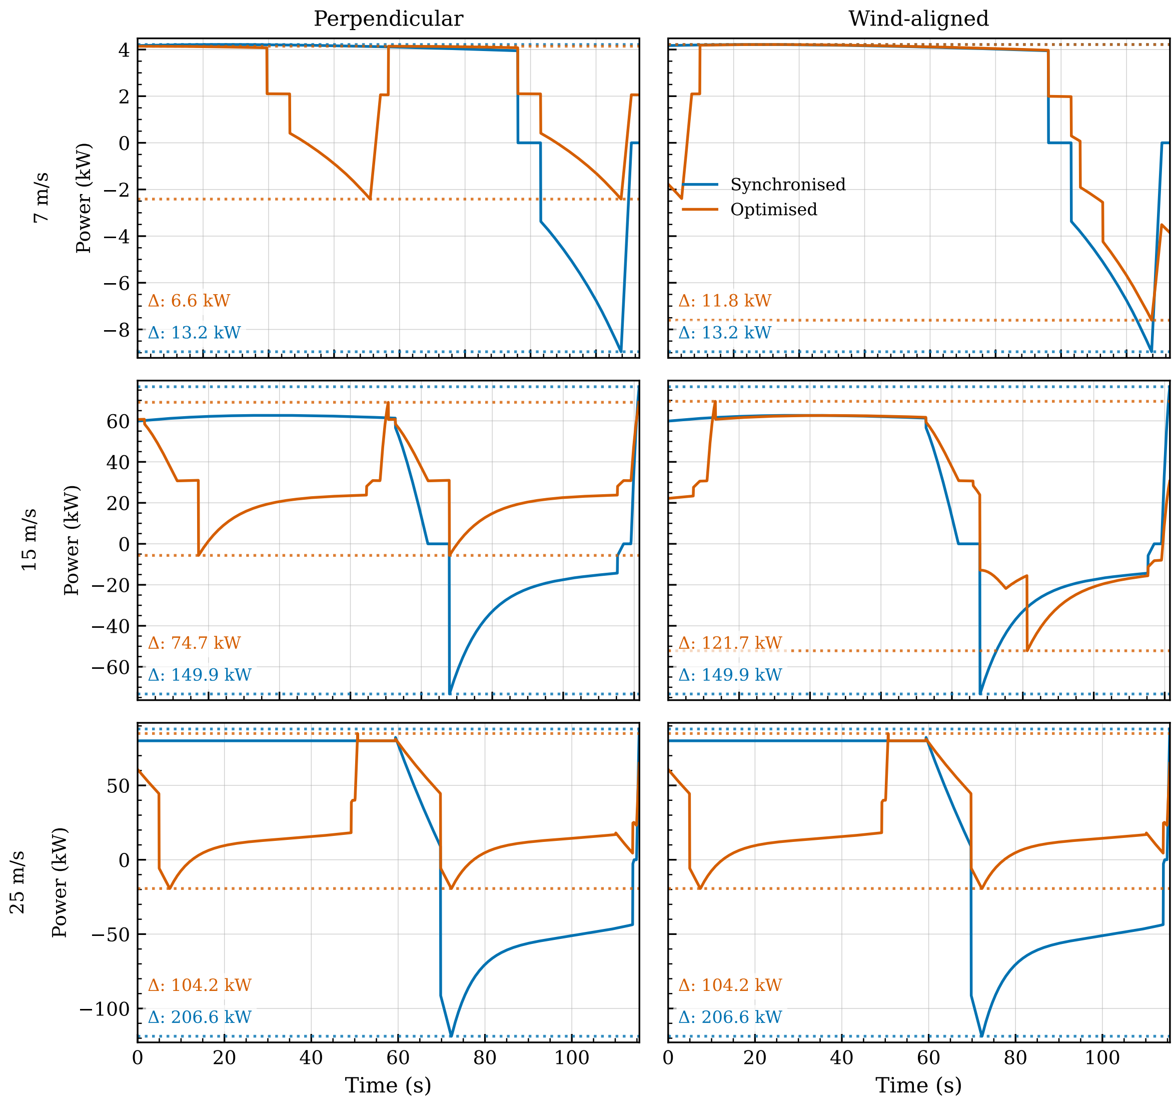

Figure 7.2: Farm power output of a two-system farm at wind speeds of $7\ \mathrm{m\,s^{-1}}$, $15\ \mathrm{m\,s^{-1}}$, and $25\ \mathrm{m\,s^{-1}}$, comparing the peak-to-peak power variation of the wind-perpendicular and wind-aligned layouts. The dashed lines show the maximum and minimum farm power output encountered during a cycle.

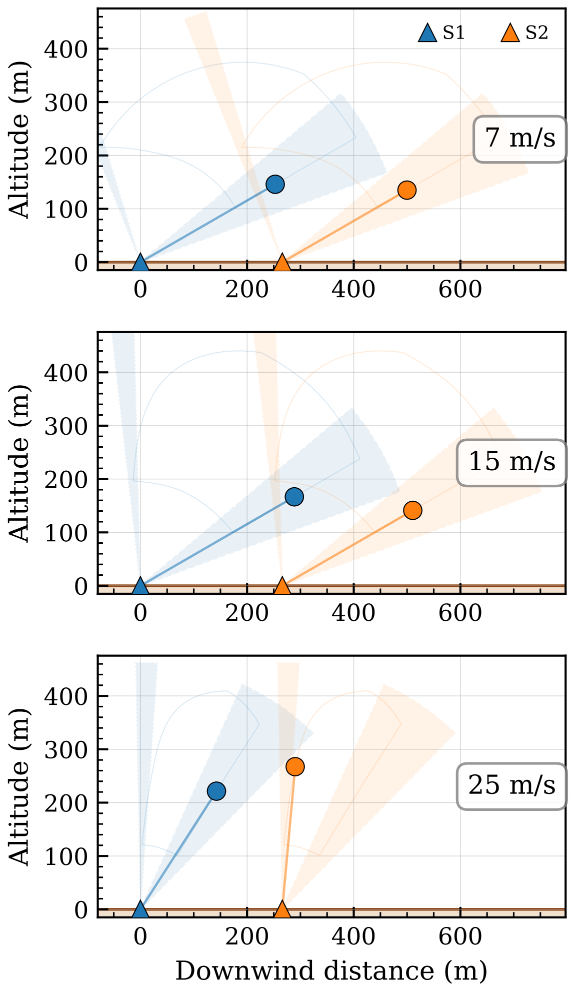

Figure 7.3: Operational reel-in and reel-out cones for the considered wind speeds.

7.5. Ten-system wind-aligned and wind-perpendicular layouts

Figure 7.4 compares the power output of a 10-system farm for wind-perpendicular and wind-aligned layouts. For $15\,\mathrm{m\,s^{-1}}$ and $25\,\mathrm{m\,s^{-1}}$, both layouts give the same peak-to-peak power variation. At $7\,\mathrm{m\,s^{-1}}$, the wind-aligned layout gives a larger peak-to-peak variation than the wind-perpendicular layout. This shows that the wind-aligned phase constraint can still limit the achievable smoothing, even for a larger farm. However, compared with the two-system case, the effect is smaller because the 10 systems require smaller phase differences between neighboring systems to distribute the power cycles over one period.

Figure 7.5 shows the optimized asynchronous operation of the 10-system wind-aligned farm at a wind speed of $15\,\mathrm{m\,s^{-1}}$. The colored tails show the kite positions during the previous $10\,\mathrm{s}$, corresponding to the safety margin used in the phase constraint. The snapshot illustrates that neighboring systems can operate out of phase while maintaining separation between the reel-out and reel-in regions. For example, system 6 has started to reel out and moves toward the reel-in cone of system 7, while system 7 has already left this region $10\,\mathrm{s}$ earlier.

For systems 1 to 6, all systems are in the reel-out phase. Because of the phase offsets, wakes from upstream systems can be convected toward downstream systems, as indicated by the wake boxes behind the kite positions in Figure 7.5. This may reduce the power production of downstream systems. LES simulations of two ground-generation AWE systems found that a downstream kite forced to fly through the wake of an upstream kite had a power loss of about $6\,\%$ [42]. The same study notes that wake avoidance strategies and control methods could reduce such losses in multi-kite operation.

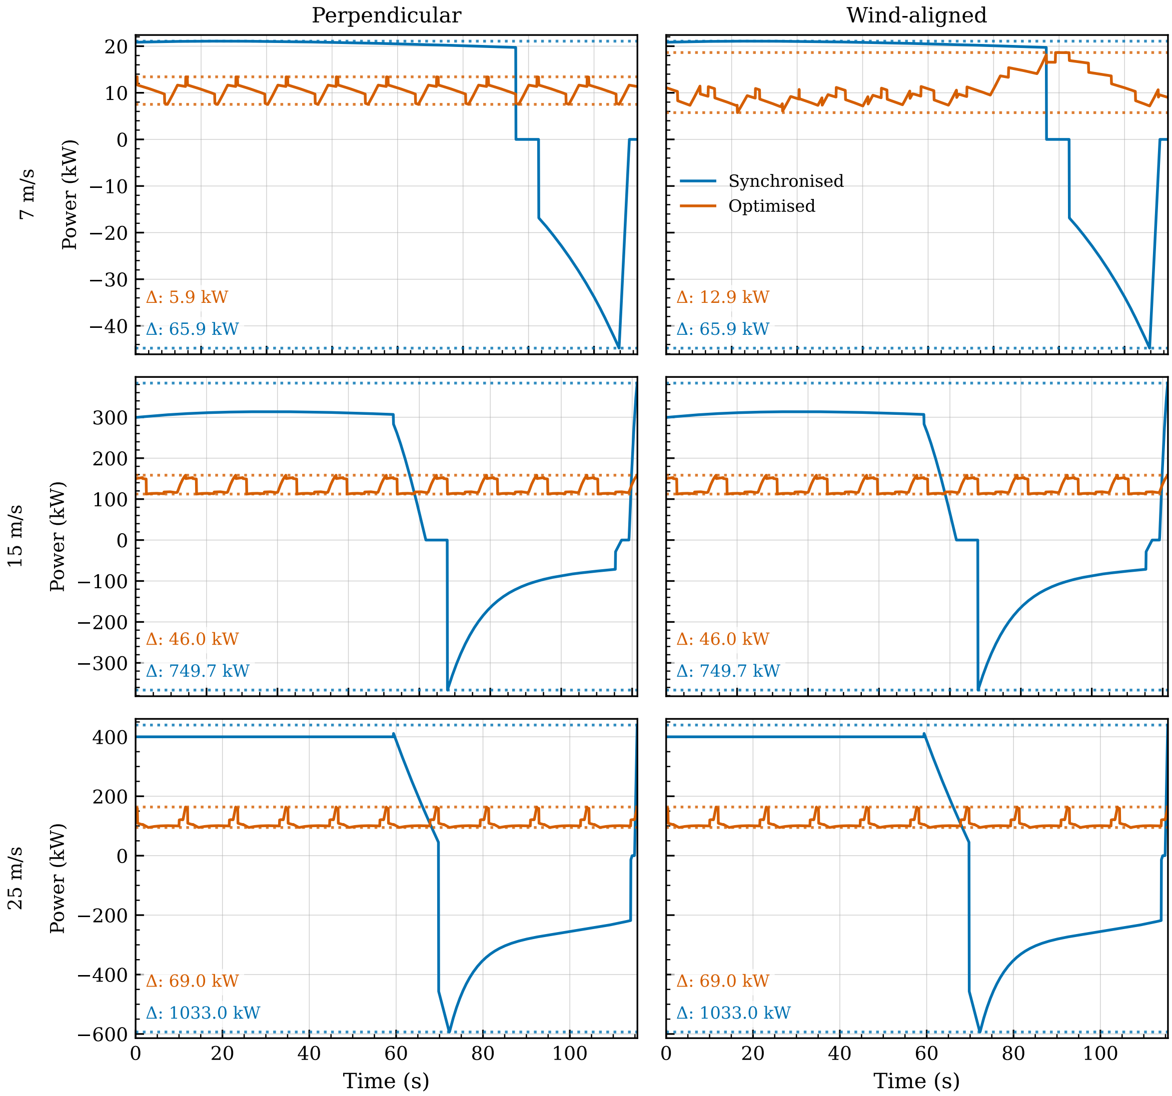

Figure 7.4: Farm power output of a 10-system farm at wind speeds of $7\ \mathrm{m\,s^{-1}}$, $15\ \mathrm{m\,s^{-1}}$, and $25\ \mathrm{m\,s^{-1}}$, comparing the peak-to-peak power variation of the wind-perpendicular and wind-aligned layouts.

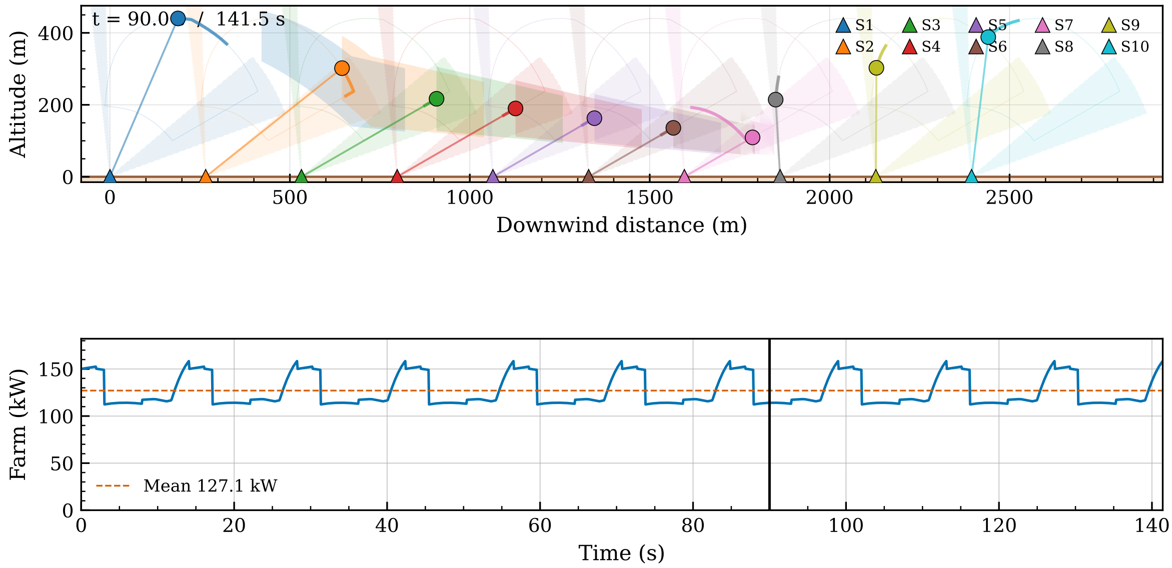

Figure 7.5: Optimized asynchronous operation of the 10-system wind-aligned farm at a wind speed of $15\ \mathrm{m\,s^{-1}}$. The tails in the upper panel show the kite positions during the previous $10\ \mathrm{s}$, which is used as the safety margin. The boxes behind the kite positions indicate the wake regions.

7.6. Conclusion

This chapter investigated whether wind-aligned systems with adjacent reel-in and reel-out envelopes have to be operated in sync. A geometric maximum phase offset was derived and used as a constraint in the farm phase optimization. Within the geometric envelope model used here, the results show that wind-aligned systems can operate asynchronously if the phase offsets remain within this limit.

For the two-system farm, the wind-aligned layout gives a larger peak-to-peak power variation than the wind-perpendicular layout at $7\ \mathrm{m\,s^{-1}}$ and $15\ \mathrm{m\,s^{-1}}$. This occurs because the maximum phase offset limits how far the two systems can be shifted apart. At $25\ \mathrm{m\,s^{-1}}$, both layouts give the same peak-to-peak variation. For the 10-system farm, the wind-aligned and wind-perpendicular layouts give the same peak-to-peak variation at $15\ \mathrm{m\,s^{-1}}$ and $25\ \mathrm{m\,s^{-1}}$, while the wind-aligned layout still gives a larger variation at $7\ \mathrm{m\,s^{-1}}$. The wind-aligned constraint is therefore less restrictive for the larger farm, but it is not removed completely.

The results show that synchronized operation is not always required for wind-aligned systems. Allowing a limited phase offset can reduce the difference between the maximum and minimum farm power output while maintaining the geometric separation imposed by the reel-in and reel-out envelopes. However, asynchronous wind-aligned operation can place downstream systems in the wakes of upstream systems. This effect is not included in the phase optimization and may reduce the power output of downstream systems. Wake effects and wake avoidance should therefore be included when studying wind-aligned AWE farm operation.

8 Case studies

This chapter applies AWESPA to two case studies from the MERIDIONAL project. The aim is to estimate the annual energy production and related performance parameters for an isolated $100\,\mathrm{kW}$ airborne wind energy system and for a $1\,\mathrm{MW}$ airborne wind energy farm. The first case study considers a single inland system. The second case study uses the same system in a 10-system farm layout.

The wind resource is treated in the same way for both case studies. For each site, the wind resource is described using the wind-profile clustering method introduced in Chapter 5. The clustering is based on a 10-year ERA5 dataset from 2010 through 2019. A reference height of $200\ \mathrm{m}$ is used throughout the case studies for the wind-profile statistics, the Weibull distribution, and the power curves.

The clustering procedure first filters out low-wind cases before deriving the representative wind profiles. This filtering is used because low-wind cases contain less information on the profile shape. After the representative profiles have been obtained, all time steps in the full ERA5 dataset are assigned to the closest representative profile. This gives two frequency distributions: one for the filtered dataset used during clustering and one for the full dataset used for the performance assessment.

In the first case study, several combinations of kite size, tether force, and ground-station generator power are compared to define a representative $100\ \mathrm{kW}$ system for the site. The resulting performance is then discussed using the power curve, annual energy production, and capacity factor. In the second case study, the same system is used in a 10-system farm. The farm case study considers the annual energy production, capacity factor, and variation of the aggregated electrical power output.

8.1. Case study 1: $100\ \mathrm{kW}$ isolated inland AWES

The first case study represents a small-scale isolated inland AWES at the RWE Airborne Wind Energy test centre in Bangor Erris, Ireland. Kitepower has tested its system at this site since October 2023. The site is used as a case study for a future $100\ \mathrm{kW}$ system.

8.1.1. Wind resource

Figure 8.1 shows the 8 representative wind profiles obtained from the ERA5 dataset for Bangor Erris. The fitted power-law profile is included for comparison and its shear exponent is 0.189. The representative profiles cover a range of profile shapes, including high-shear profiles, low-shear profiles, and profiles with a local wind-speed maximum. Profiles 5 and 7 show a low-level jet shape.

The corresponding cluster frequencies are shown in Figure 8.2. The difference between the filtered and full dataset frequencies follows from the clustering procedure. The filtered dataset is used to derive the representative profiles, while the full dataset is used to assign occurrence frequencies for the performance assessment. For example, profile 4 shows a clear difference between both frequency distributions. Its frequency increases from $13.8\ \%$ in the filtered dataset to $16.9\ \%$ in the full dataset. This indicates that profile 4 occurs relatively often during low-wind conditions that were excluded before clustering.

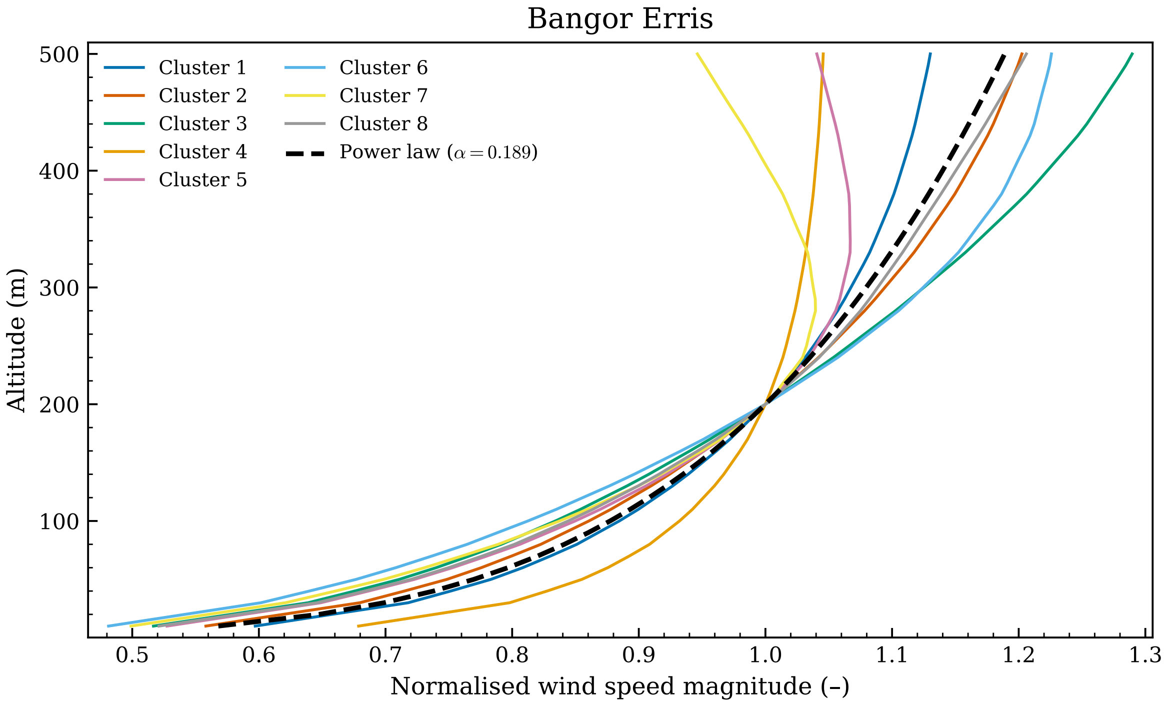

Figure 8.1: Representative wind profiles obtained by clustering the 2010–2019 ERA5 dataset for Bangor Erris, Ireland. The fitted power-law profile is shown for comparison. Wind speeds are referenced to $200\ \mathrm{m}$.

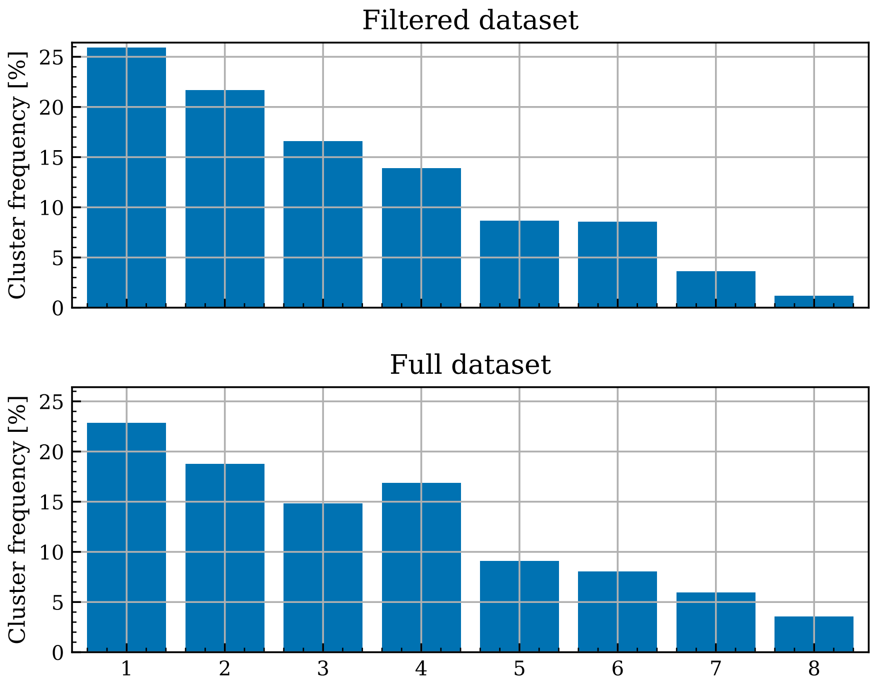

Figure 8.2: Cluster frequencies for the filtered clustering dataset and the full 2010–2019 ERA5 dataset at Bangor Erris, Ireland. The full dataset frequencies are obtained after assigning all time steps to the representative profiles.

The occurrence patterns of the representative profiles are shown in Figure 8.3. The yearly distribution is relatively stable, which indicates that no single year strongly changes the profile statistics. Clearer differences are found in the monthly and hourly distributions. Profile 4 occurs mainly at low wind speeds, during the warmer months, and during daytime. In contrast, profiles 1, 2, 3, and 6 occur more often during winter than during summer. These patterns are consistent with the increase in the full-dataset frequency of profile 4, since this profile is linked to low-wind conditions.

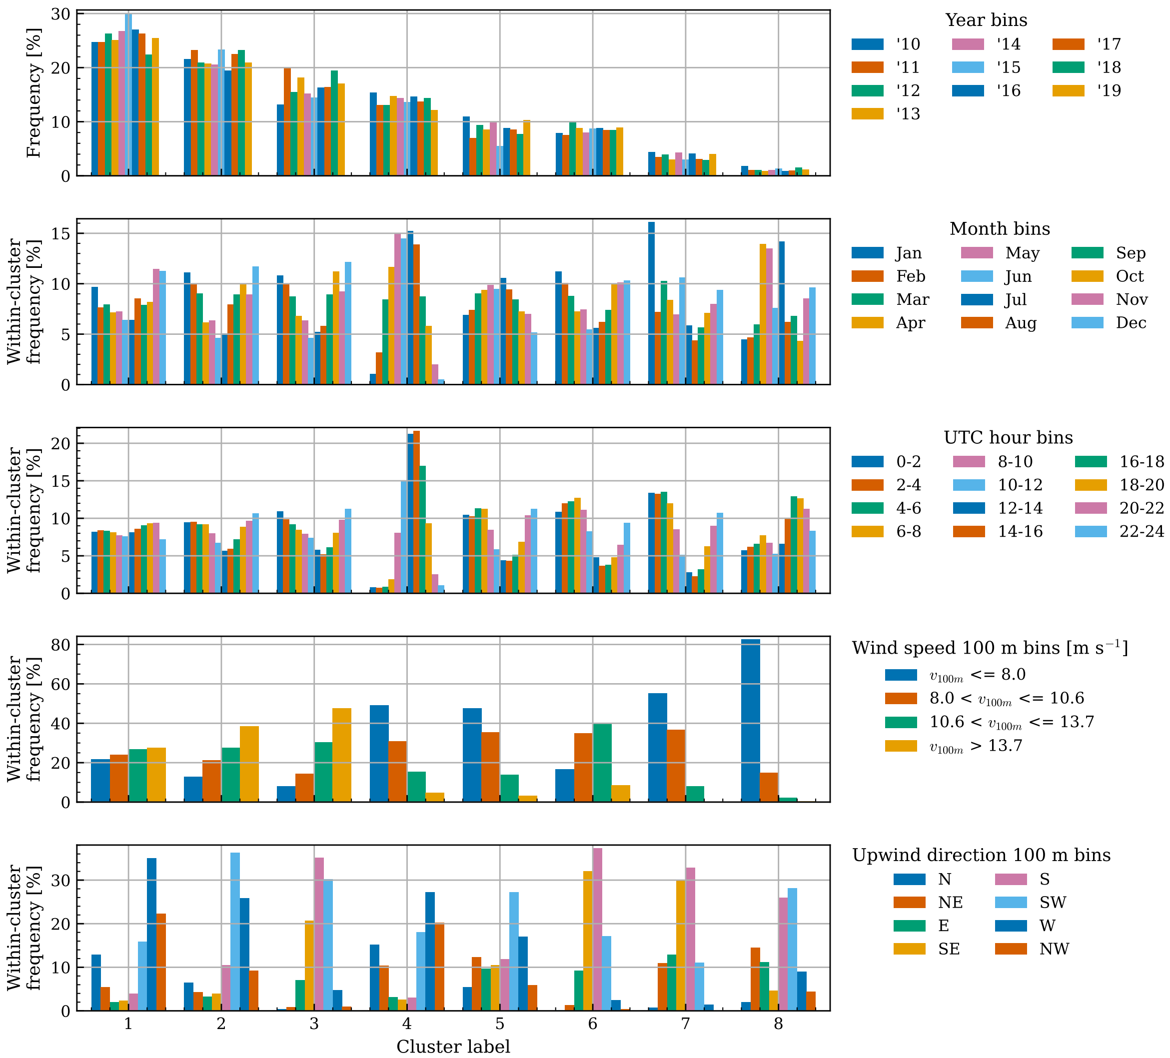

Figure 8.3: Occurrence patterns of the representative wind profiles for Bangor Erris, Ireland, based on the 2010–2019 ERA5 dataset. The profile frequencies are grouped by year, month, hour of day, wind speed at $200\ \mathrm{m}$, and wind direction.

Figure 8.4 shows the wind-speed distribution at the $200\ \mathrm{m}$ reference height. The fitted Weibull distribution has a shape factor of $k = 2.18$ and a scale factor of $c = 11.25\ \mathrm{m\,s^{-1}}$. The mean wind speed at the site is $9.98\ \mathrm{m\,s^{-1}}$ at $200\ \mathrm{m}$. This wind-speed distribution is used later to weight the power curve when calculating the annual energy production.

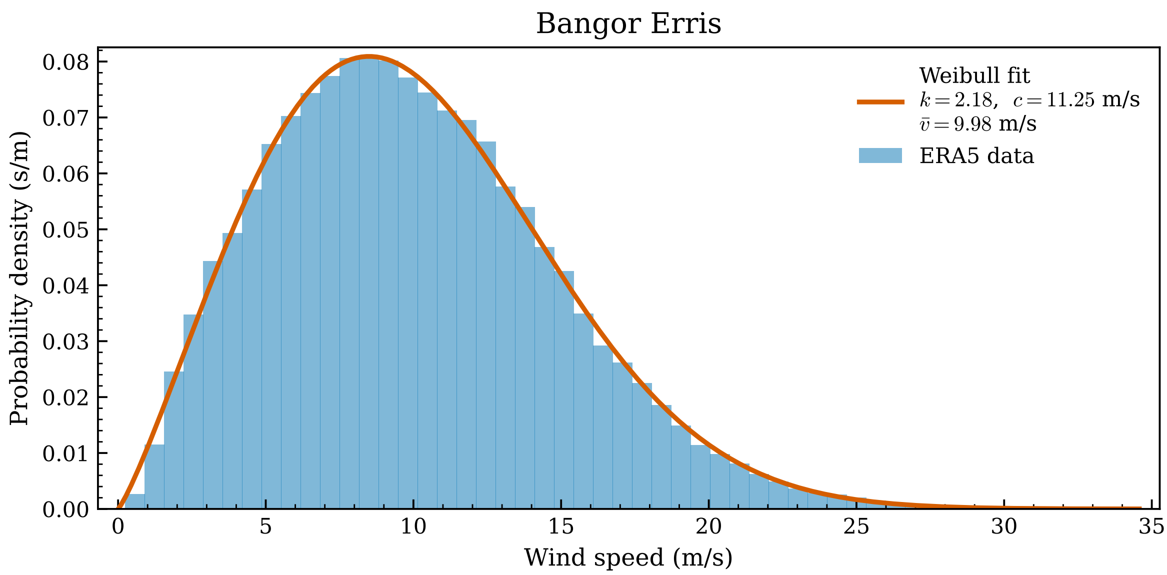

Figure 8.4: Wind-speed distribution at $200\ \mathrm{m}$ for Bangor Erris, Ireland, based on the 2010–2019 ERA5 dataset. The fitted Weibull distribution has a shape factor of $k = 2.18$ and a scale factor of $c = 11.25\ \mathrm{m\,s^{-1}}$.

The ERA5-based wind resource should be interpreted together with the comparison against lidar data in Section 5.6. That comparison showed that the average wind profile and the Weibull distribution from ERA5 were close to the lidar-based results. However, the profile occurrence frequencies did not match as well. ERA5 assigned more cases to standard profile shapes, while the lidar data showed a wider range of profiles, including more frequent high-shear and low-shear profiles and more pronounced low-level jets. Since no long-term lidar dataset is available for this site, the ERA5 dataset is used for the performance assessment in this case study.

8.1.2. $100\ \mathrm{kW}$ system definition

The $100\ \mathrm{kW}$ system is defined by comparing several combinations of kite size, maximum tether force, and maximum generator power. For a pumping ground-generation AWES, the rated power is not equal to the maximum generator power. The average cycle power depends on both the reel-out phase, where electrical power is produced, and the reel-in phase, where electrical power is consumed. The rated power therefore results from the full pumping cycle and depends on the aerodynamic properties, kite size, tether force limits, generator power, and cycle settings. For this reason, several design combinations are compared.

The comparison uses the fitted power-law wind profile from case study 1, shown in Figure 8.1. This keeps the calculations simple while showing how the main design parameters affect the power curve, capacity factor, and annual energy production. After one system has been selected, it is evaluated using the 8 representative wind profiles from the clustering analysis.

Model assumptions

Several model parameters are kept constant during the system comparison. These parameters are listed in Table 8.1. The available tether length is set to $600\ \mathrm{m}$. For each power curve, the minimum and maximum tether lengths of the pumping cycle are optimized within this available tether length. The minimum height is set to $100\ \mathrm{m}$. The elevation angle is kept constant at $30°$ and is not optimized to reduce the computational cost. The tether density is set to $724\ \mathrm{kg\,m^{-3}}$ and the kite control unit mass is kept constant at $20\ \mathrm{kg}$.

Table 8.1: Constant model parameters used for the $100\ \mathrm{kW}$ system comparison.

Parameter	Value
Available tether length	$600\ \mathrm{m}$
Minimum height	$100\ \mathrm{m}$
Elevation angle	$30°$
Generator efficiency	$\eta_{\mathrm{gen}} = 0.95$
Storage efficiency	$\eta_{\mathrm{sto}} = 0.90$
Reel-out lift coefficient	$C_{\mathrm{L,out}} = 0.8$
Reel-out drag coefficient	$C_{\mathrm{D,out}} = 0.16$
Reel-in lift coefficient	$C_{\mathrm{L,in}} = 0.3$
Reel-in drag coefficient	$C_{\mathrm{D,in}} = 0.12$
Tether density	$724\ \mathrm{kg\,m^{-3}}$
Kite control unit mass	$20\ \mathrm{kg}$

The aerodynamic coefficients are based on the TU Delft V3 kite model [31]. During reel-out, the angle of attack is set to $12°$, giving $C_{\mathrm{L,out}} = 0.8$ and $C_{\mathrm{D,out}} = 0.16$. During reel-in, the angle of attack is set to $3°$, giving $C_{\mathrm{L,in}} = 0.3$ and $C_{\mathrm{D,in}} = 0.12$.

Design comparison

The comparison considers three design parameters: kite size, maximum tether force, and generator power. The kite size changes the aerodynamic force, the kite mass, the projected area, and the minimum tether force. The maximum tether force changes the allowable reel-out force and is combined with a larger tether diameter. The generator power changes the maximum reel-out power. The investigated values are listed in Table 8.2, Table 8.3, and Table 8.4.

Table 8.2: Kite parameters considered for the $100\ \mathrm{kW}$ system. The projected kite area is calculated from the flat area using a factor of 0.79, based on the ratio between projected and flat area of the TU Delft V3 kite.

Flat kite area $[\mathrm{m}^2]$	Projected kite area $[\mathrm{m}^2]$	Kite mass $[\mathrm{kg}]$	Minimum tether force $[\mathrm{kN}]$
100	79.0	33.3	3.75
120	94.8	39.0	4.50
140	110.6	47.0	5.25
160	126.4	53.0	6.00

Table 8.3: Maximum tether force values considered for the $100\ \mathrm{kW}$ system. Increasing the maximum tether force is combined with a larger tether diameter.

Maximum tether force $[\mathrm{kN}]$	Tether diameter $[\mathrm{mm}]$
60	12
70	13
80	14

Table 8.4: Maximum generator powers considered for the $100\ \mathrm{kW}$ system.

Maximum generator power $[\mathrm{kW}]$
160
170
180

For each combination of kite size, maximum tether force, and generator power, a power curve is calculated. The power curve is combined with the wind-speed distribution to calculate the annual energy production and capacity factor. For a pumping AWES, the cycle power can decrease above the rated wind speed. This occurs because the reel-out power is limited by the generator power, while the reel-in power can still increase at higher wind speeds. A design with a low rated wind speed can therefore give a lower annual energy production if a large part of the wind-speed distribution lies above the rated wind speed and the power drop-off is significant.

Figure 8.5 shows the power curves for the different designs. The figure is used to compare how kite size, maximum tether force, and generator power affect the rated wind speed and the rated power.

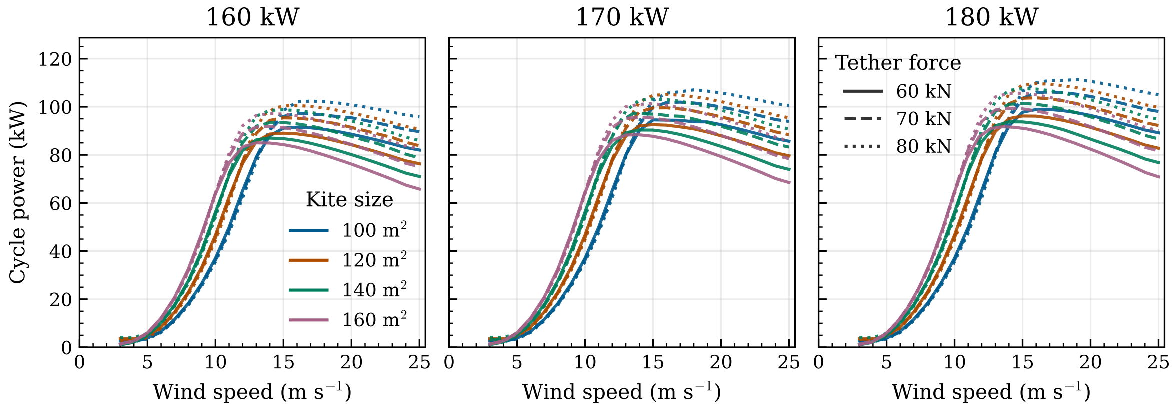

Figure 8.5: Power curves for the different system designs. The curves show the effect of varying kite size, maximum tether force, and generator power. The power curves are calculated using the fitted power-law wind profile for case study 1. The reference wind speed is taken at $200\ \mathrm{m}$.

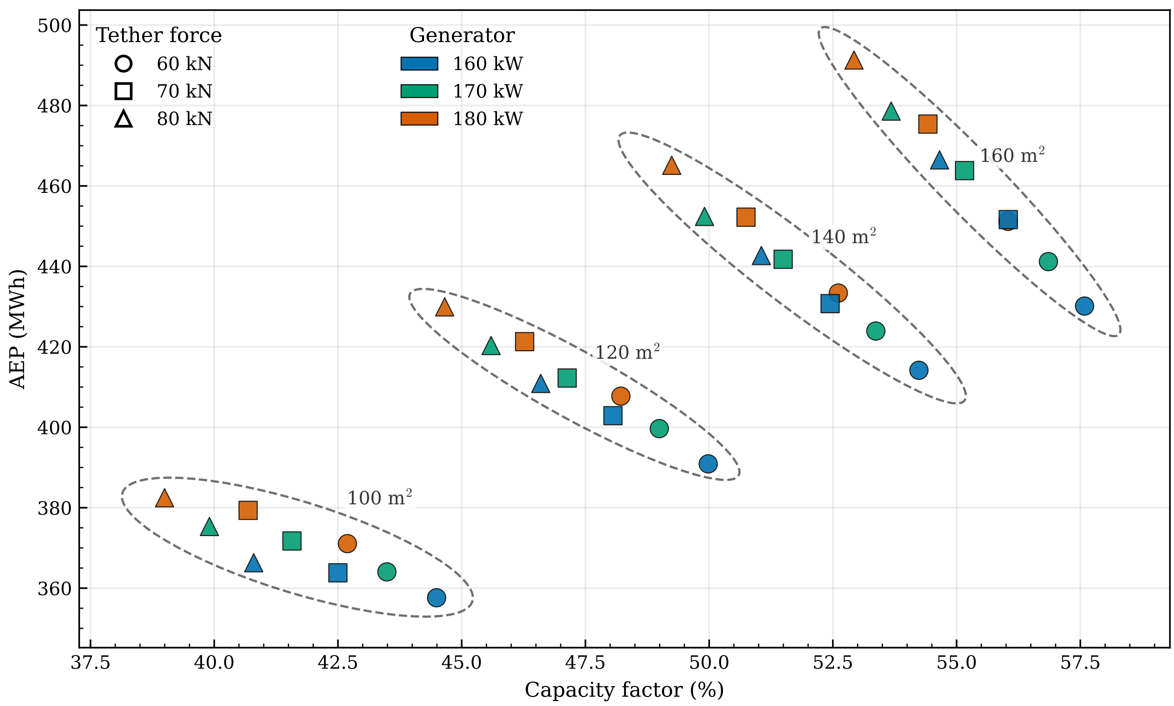

Figure 8.6: Annual energy production and capacity factor for the different system designs. The values are calculated from the case study 1 wind-speed distribution. Labels next to the points show the rated power in kW.

The results in Figure 8.5 and Figure 8.6 show the effect of the three design parameters. For a fixed maximum tether force and generator power, a larger kite gives a lower rated wind speed and higher power below rated wind speed. This is caused by the larger aerodynamic force. As a result, larger kites generally give a higher annual energy production and capacity factor. Smaller kites can reach a higher maximum cycle power, but this maximum occurs at higher wind speeds. This gives a lower annual energy production for most combinations. The higher maximum power of the smaller kites is mainly caused by their lower reel-in energy demand. When the reel-out phase is limited by the generator, the reel-out power is similar, while the smaller kite requires less power during reel-in.

For a fixed kite size and generator power, increasing the maximum tether force has little effect at low wind speeds. In this part of the power curve, regime 1, the tether force limit is not active. Designs with a higher maximum tether force can even have a slightly lower cycle power in regime 1, because the larger tether diameter increases the tether drag and mass. At higher wind speeds, a larger maximum tether force allows unconstrained operation over a wider wind-speed range. The transition to force-constrained operation is delayed, so the cycle power remains higher before rated power is reached. A higher maximum tether force also increases the rated power, because the generator limit can be reached at a lower reel-out speed. This increases the duration of the reel-out phase and increases the reel-out energy per cycle. As a result, increasing the maximum tether force increases the annual energy production, while the capacity factor decreases because the rated power also increases. The increase in annual energy production is larger for the larger kites than for the smaller kites, as shown in Figure 8.6.

For a fixed kite size and maximum tether force, the effect of the generator power is mainly visible above the wind speeds where the generator limit becomes active. At low wind speeds, the power curves are nearly the same because the generator power does not limit the operation. A larger generator power delays the start of the fully generator-limited part of the curve and shifts the rated wind speed to higher wind speeds. It therefore increases the annual energy production, but also increases the rated power used in the capacity factor calculation. For this reason, the capacity factor decreases when the generator power is increased, even though the annual energy production increases.

Not all parameter combinations give a rated power of exactly $100\ \mathrm{kW}$. Since this case study considers a $100\ \mathrm{kW}$ system, only designs with a rated power close to $100\ \mathrm{kW}$ are considered for the final selection. No cost model is included in this analysis thus the selection of the system parameters is based only on the AEP and the capacity factor. The chosen system has a kite size of $160\ \mathrm{m}^2$, a tether with a maximum tether force of $80\ \mathrm{kN}$, and a maximum generator power of $170\ \mathrm{kW}$. The awesIO system configuration file is shown in Section A.2.

8.1.3. System performance

The selected system is evaluated with the 8 representative wind profiles from the clustering analysis, with the elevation angle optimized for each profile instead of being kept fixed. This gives a more detailed estimate of the system performance under the different wind-profile shapes that occur at the site. Figure 8.7 shows the resulting power curves. The differences between the curves show how the vertical wind-profile shape affects the cycle power.

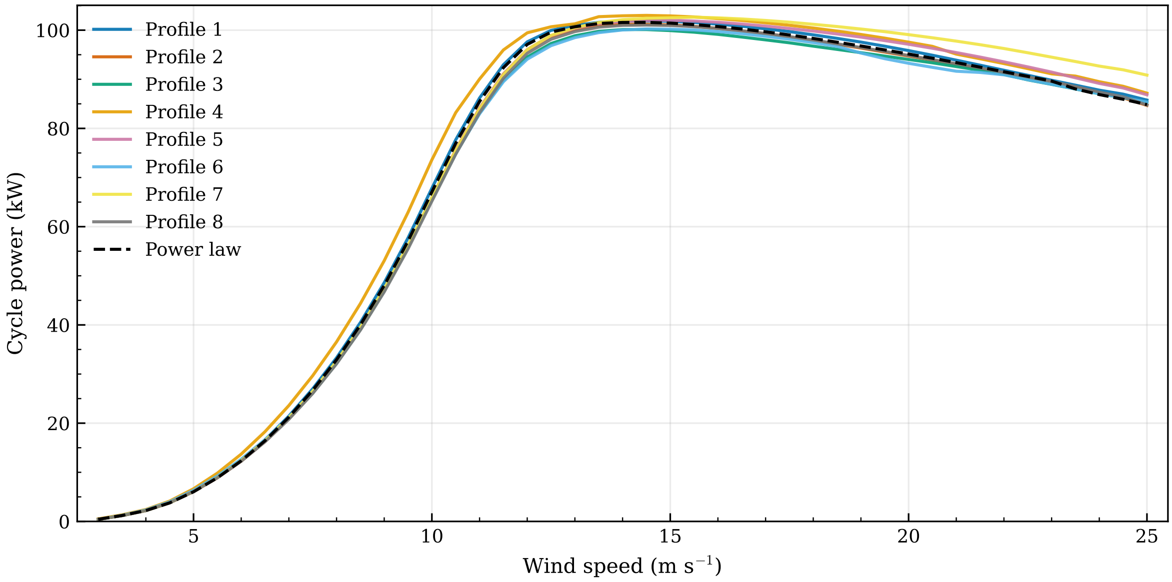

Figure 8.7: Power curves of the selected $100\ \mathrm{kW}$ system for the 8 representative wind profiles of case study 1 and the fitted power-law profile. The reference wind speed is taken at $200\ \mathrm{m}$. The elevation angle is optimized separately for each profile.

The wind-profile shape affects the power curve, although the differences are limited for most profiles. Profile 4 gives a higher cycle power below the rated wind speed than the other representative profiles. This is because the reference wind speed is defined at $200\ \mathrm{m}$, while the traction phase mainly occurs below this height. Since profile 4 has low shear, the wind speed encountered by the kite below $200\ \mathrm{m}$ is higher than for profiles with stronger shear at the same reference wind speed.

Profile 7 gives the highest cycle power above the rated wind speed. This profile corresponds to a low-level jet. Above the rated wind speed, the reel-out power is limited, so differences in cycle power are mainly caused by the reel-in phase. For profile 7, the wind speed at the higher altitudes reached during reel-in is lower than for most other profiles. The reel-in phase therefore consumes less energy, which increases the net cycle power.

The fitted power-law profile is also included in Figure 8.7. The difference between calculating annual energy production with the power-law profile and summing the contributions from the clustered profiles is small. The annual energy production is $490.39\ \mathrm{MWh}$ with the power-law profile and $488.62\ \mathrm{MWh}$ with the clustered profiles. The corresponding capacity factor for the power-law profile is $55.0\ \%$.

Table 8.5: Selected $100\ \mathrm{kW}$ system used for the isolated-system performance assessment and the farm case study.

Parameter	Value
Flat kite area	$160\ \mathrm{m}^2$
Projected kite area	$126.4\ \mathrm{m}^2$
Kite mass	$53.0\ \mathrm{kg}$
KCU mass	$20\ \mathrm{kg}$
Minimum tether force	$6\ \mathrm{kN}$
Maximum tether force	$80\ \mathrm{kN}$
Tether diameter	$14\ \mathrm{mm}$
Maximum generator power	$170\ \mathrm{kW}$
Annual energy production, power-law profile	$490.39\ \mathrm{MWh}$
Annual energy production, clustered profiles	$488.62\ \mathrm{MWh}$
Capacity factor, power-law profile	$55.0\ \%$

8.2. Case study 2: $1\ \mathrm{MW}$ inland AWES farm

The second case study represents a medium-scale inland AWES farm with a total rated power of $1\ \mathrm{MW}$. The farm is located at Maasvlakte 2 in the Port of Rotterdam. Although the site is connected to inland infrastructure and logistics, its exposed coastal location gives it a wind climate that is closer to an offshore site than to a sheltered inland site. The farm is modeled as a 10-system layout using the same $100\ \mathrm{kW}$ system considered in Section 8.1.

8.2.1. Wind resource

Figure 8.8 shows the 8 representative wind profiles obtained from the ERA5 dataset for Maasvlakte 2. The fitted power-law profile is included for comparison. The fitted shear exponent is 0.127, which is lower than the value of 0.189 found for case study 1. This indicates that the average wind profile at Maasvlakte 2 has lower shear than the average wind profile at Bangor Erris.

The representative profiles show a higher occurrence of low-level jet shapes than in case study 1. This is also visible in the cluster frequencies shown in Figure 8.9.

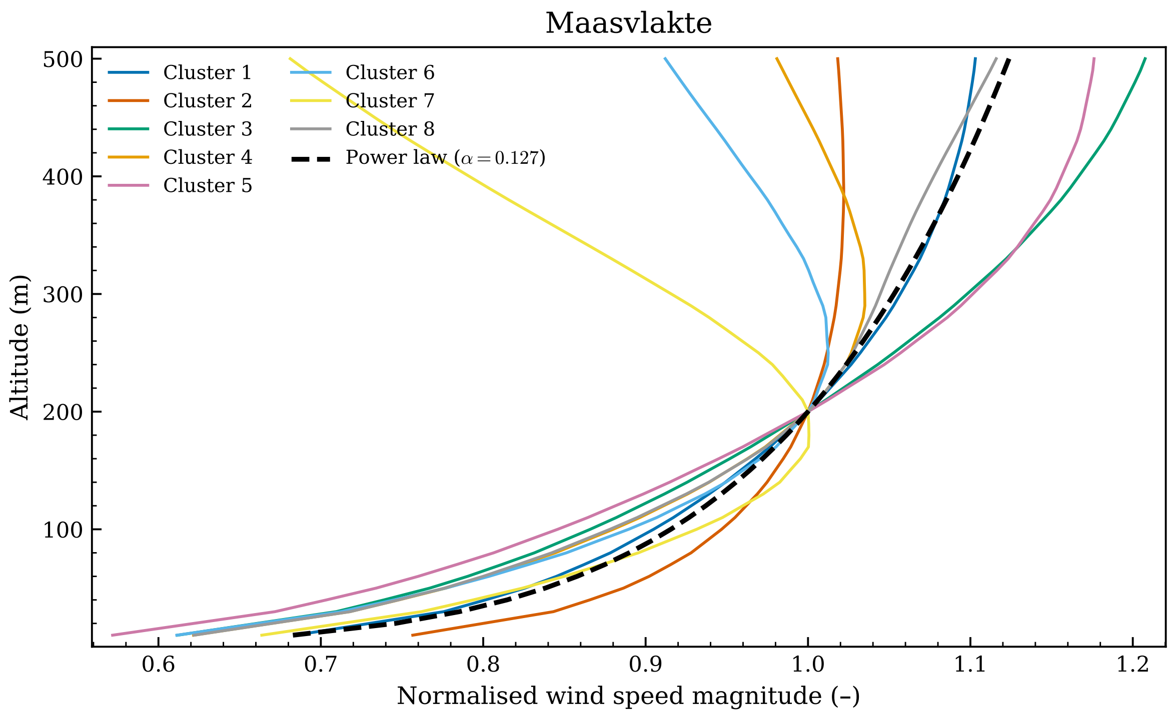

Figure 8.8: Representative wind profiles obtained by clustering the 2010–2019 ERA5 dataset for Maasvlakte 2, Port of Rotterdam. The fitted power-law profile is shown for comparison. Wind speeds are referenced to $200\ \mathrm{m}$.

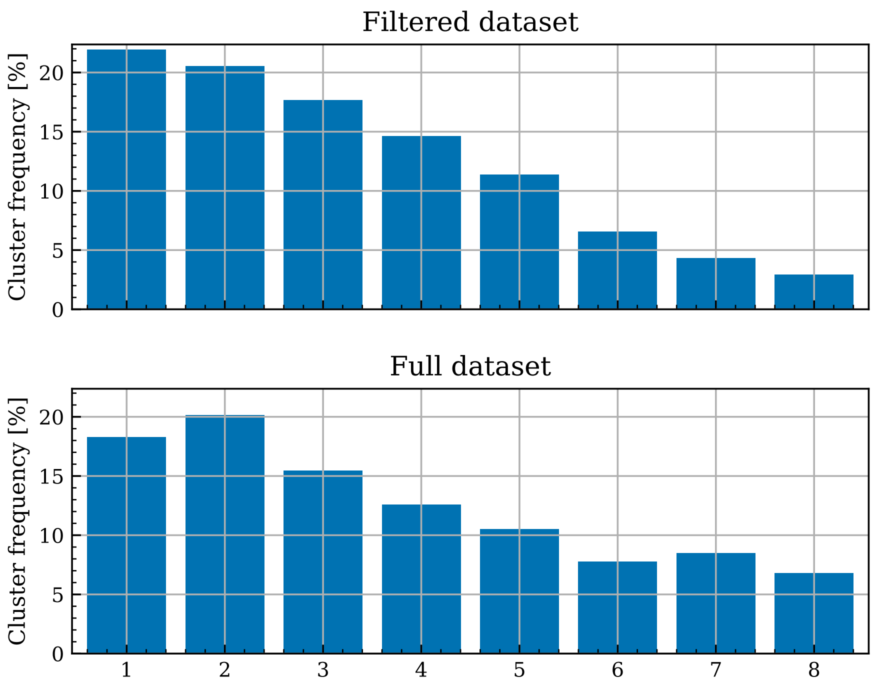

Figure 8.9: Cluster frequencies for the filtered clustering dataset and the full 2010–2019 ERA5 dataset at Maasvlakte 2, Port of Rotterdam. The full dataset frequencies are obtained after assigning all time steps to the representative profiles.

The occurrence patterns of the representative profiles are shown in Figure 8.10. The yearly variation is limited, which indicates that the profile distribution is relatively stable over the 10-year ERA5 period. Larger differences are found across month, hour of day, wind speed, and wind direction. Profile 3 occurs mainly at higher wind speeds, while profiles 7 and 8 occur mainly at lower wind speeds. Several profile shapes also occur more often in specific months and at specific hours of the day. The wind-direction distribution shows that some profiles are linked more strongly to specific inflow sectors. These patterns show that the profile occurrence depends on both the wind-speed regime and the time and direction of occurrence.

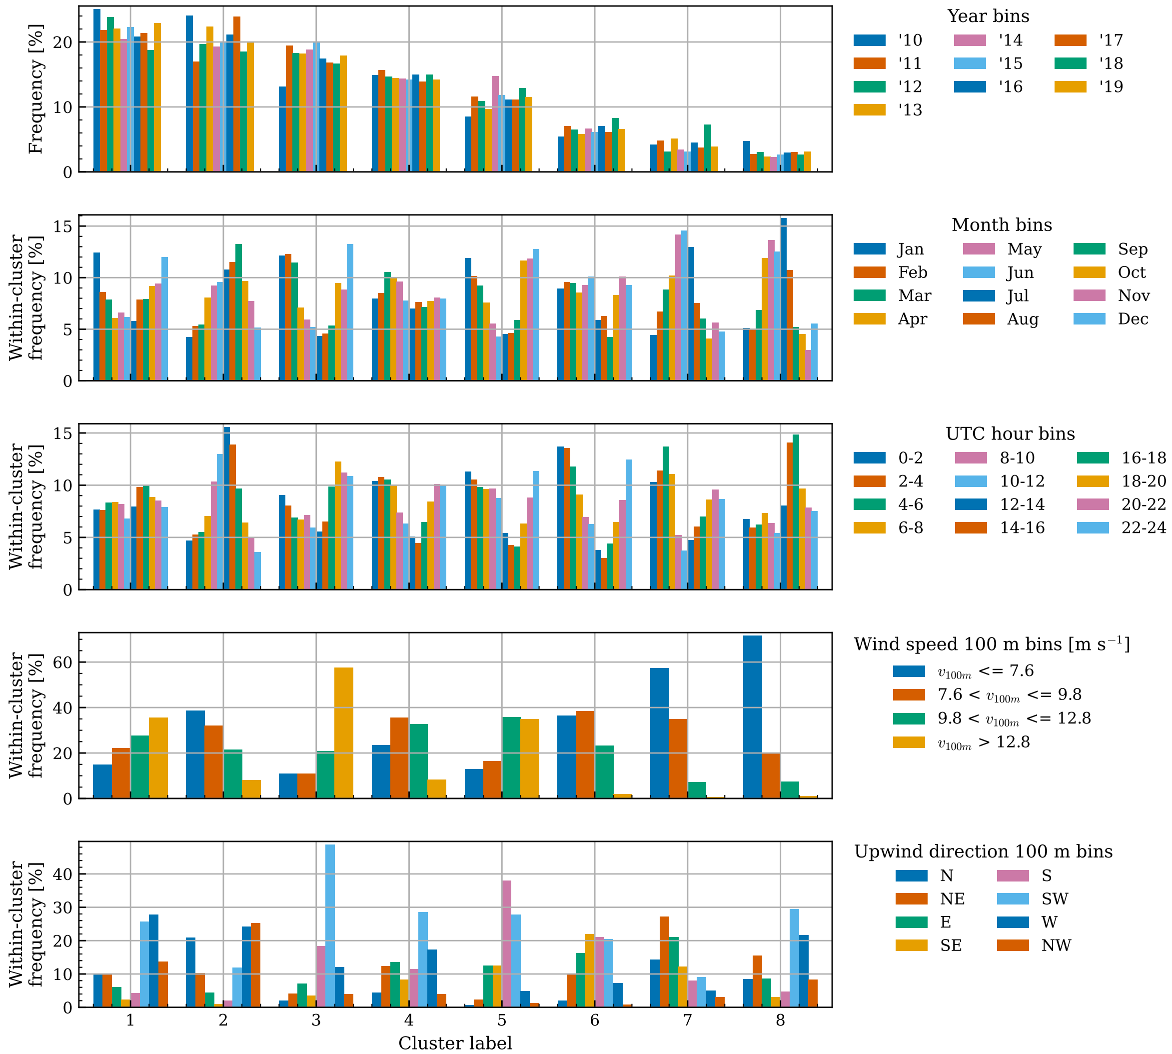

Figure 8.10: Occurrence patterns of the representative wind profiles for Maasvlakte 2, Port of Rotterdam, based on the 2010–2019 ERA5 dataset. The profile frequencies are grouped by year, month, hour of day, wind speed at $200\ \mathrm{m}$, and wind direction.

Figure 8.11 shows the wind-speed distribution at the $200\ \mathrm{m}$ reference height. The fitted Weibull distribution has a shape factor of $k = 2.18$ and a scale factor of $c = 10.40\ \mathrm{m\,s^{-1}}$. The mean wind speed at the site is $9.12\ \mathrm{m\,s^{-1}}$ at $200\ \mathrm{m}$. This mean wind speed is lower than the value found for case study 1, while the lower fitted shear exponent indicates a flatter average vertical wind profile.

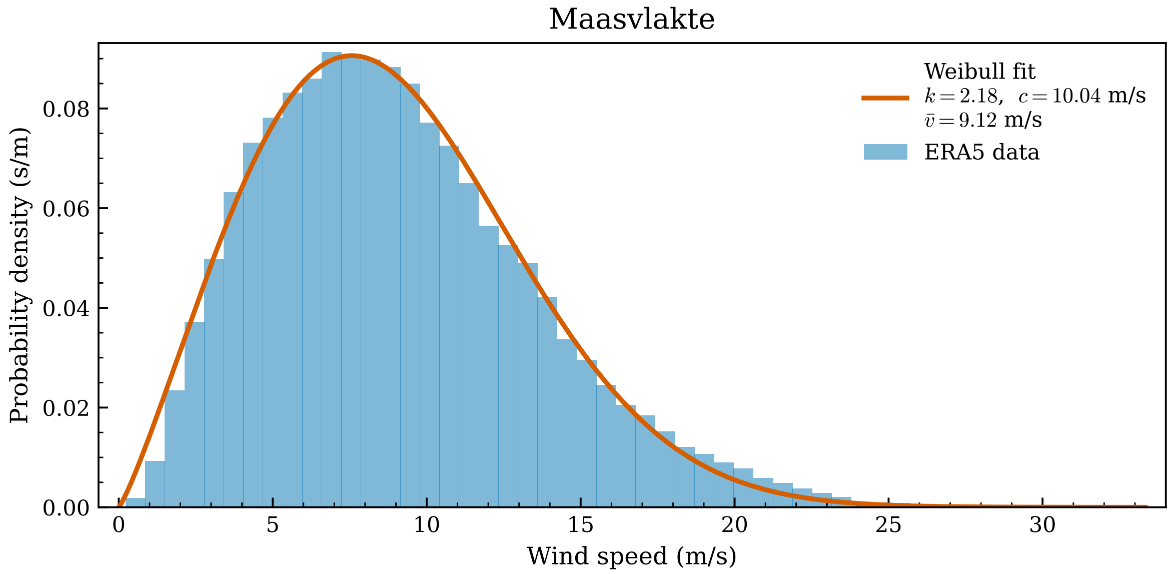

Figure 8.11: Wind-speed distribution at $200\ \mathrm{m}$ for Maasvlakte 2, Port of Rotterdam, based on the 2010–2019 ERA5 dataset. The fitted Weibull distribution has a shape factor of $k = 2.18$ and a scale factor of $c = 10.40\ \mathrm{m\,s^{-1}}$.

8.2.2. System performance

The selected $100\ \mathrm{kW}$ system from Section 8.1.3 is first evaluated for the wind resource of case study 2. Figure 8.12 shows the power curves for the 8 representative wind profiles and for the fitted power-law profile.

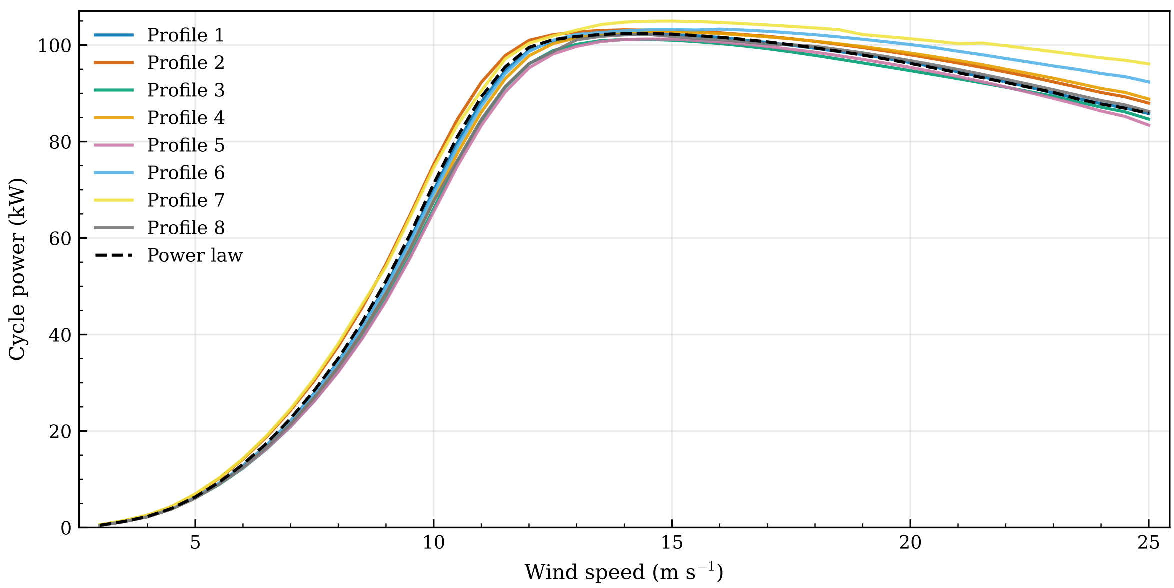

Figure 8.12: Power curves of the selected $100\ \mathrm{kW}$ system for the 8 representative wind profiles of case study 2 and the fitted power-law profile. The reference wind speed is taken at $200\ \mathrm{m}$. The elevation angle is optimized separately for each profile.

The profile shape affects the power curves in a similar way as in case study 1. Profile 2, which has the lowest shear, gives higher cycle power below rated wind speed. This is because the reference wind speed is defined at $200\ \mathrm{m}$, while the traction phase mainly occurs below this height. Profile 7, which has a low-level-jet shape, gives higher cycle power above rated wind speed. In this region, the reel-out power is limited, so differences in cycle power are mainly caused by the reel-in phase. The lower wind speed at the higher altitudes reached during reel-in reduces the reel-in energy demand.

The annual energy production is $441.12\ \mathrm{MWh}$ when the fitted power-law profile is used, with a capacity factor of $49.1\ \%$. When the 8 representative wind profiles are used, the annual energy production is $436.08\ \mathrm{MWh}$. The clustered profiles therefore give a lower annual energy production than the fitted power-law profile for this site.

8.2.3. Farm operation

The $1\ \mathrm{MW}$ farm is modeled as a 10-system farm using the selected $100\ \mathrm{kW}$ system. The analysis focuses on the cyclic variation in aggregated electrical power output. The farm model from Chapter 7 is used to compare synchronized operation with phase-shifted operation, both with and without the wind-aligned phase constraint.

Two operating cases are considered for the phase-shifted operation. In the first case, the systems are spaced far enough apart that their flight envelopes cannot overlap. The phase offsets between systems can then be chosen freely to minimize the peak-to-peak variation in farm power. Since the maximum tether length is $600\ \mathrm{m}$, this case requires a spacing larger than $600\ \mathrm{m}$ between neighboring systems.

In the second case, the 10 systems are placed in a wind-aligned line with the spacing defined by Equation 2.2. This gives a spacing of $416.29\ \mathrm{m}$ between neighboring systems. For this closer wind-aligned layout, the allowable phase offsets are constrained by the geometric separation condition derived in Chapter 7, with a safety margin of $10\ \mathrm{s}$. This case is included to investigate whether the systems can be placed closer together to increase farm power density while still limiting the variation in aggregated farm power.

The farm power variation is evaluated at reference wind speeds of $7\ \mathrm{m\,s^{-1}}$, $15\ \mathrm{m\,s^{-1}}$, and $25\ \mathrm{m\,s^{-1}}$. The power cycle of the selected $100\ \mathrm{kW}$ system is calculated using the fitted power-law profile for case study 2. The farm power is then obtained by summing the time-shifted power cycles of the 10 systems. Figure 8.13 compares the resulting farm power for synchronized operation, phase-shifted operation without wind-aligned constraints, and phase-shifted operation with the wind-aligned phase constraint.

The comparison shows how farm-level phase shifting can reduce the variation in aggregated power output. This reduction would reduce the storage capacity required for a given level of farm-power smoothing, although storage sizing is not modeled explicitly here. The unconstrained and wind-aligned cases give the same farm power output at $15\ \mathrm{m\,s^{-1}}$ and $25\ \mathrm{m\,s^{-1}}$. At $7\ \mathrm{m\,s^{-1}}$, the wind-aligned phase constraint prevents the systems from using the phase offsets that give the lowest peak-to-peak variation. This shows that closer wind-aligned spacing can increase farm power density without increasing farm power variation for some operating conditions, but that the constraint can still be restrictive at some wind speeds. This limitation could be reduced by optimizing the pumping cycle together with the farm phase offsets, so that the maximum allowable phase offset increases.

In this case study, all systems are placed on a single wind-aligned line. The effect of placing systems next to each other in the wind-perpendicular direction with overlapping operational envelopes is not considered.

8.2.4. Farm performance

Wake losses are not included in this case study. The farm's annual energy production is therefore calculated by multiplying the single-system annual energy production by the number of systems. The phase shifting used in the farm operation analysis changes the timing of the power output, but it does not change the total energy production in this model. The resulting farm annual energy production is $4411\ \mathrm{MWh}$ when the fitted power-law profile is used and $4361\ \mathrm{MWh}$ when the clustered wind profiles are used. The difference between these values comes only from the wind-profile representation. Since wake losses are neglected, the farm performance results should be interpreted as an idealized estimate of the energy production of the 10-system layout.

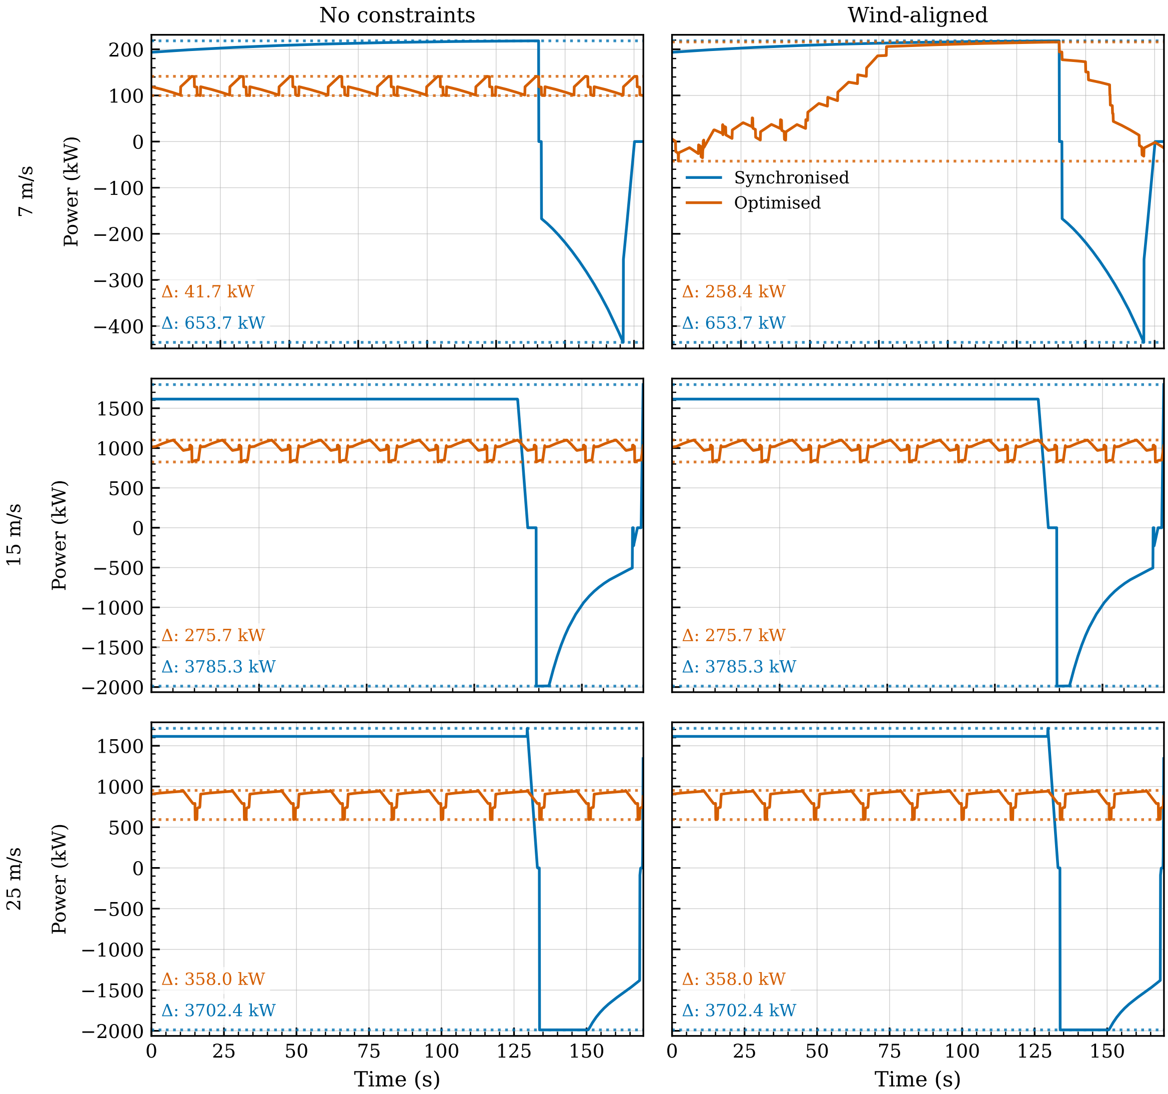

Figure 8.13: Aggregated electrical power output of the 10-system farm at reference wind speeds of $7\ \mathrm{m\,s^{-1}}$, $15\ \mathrm{m\,s^{-1}}$, and $25\ \mathrm{m\,s^{-1}}$. Comparing synchronized operation, phase-shifted operation without wind-aligned constraints, and phase-shifted operation for a wind-aligned layout with phase offsets constrained by the geometric separation condition.

9 Conclusion

This thesis developed AWESPA, a modular toolchain for the performance assessment of airborne wind energy systems and farms. The toolchain connects wind resource analysis, power estimation, farm modeling, and annual energy production estimation through a common workflow. The work also developed awesIO as a standardized input and output framework for airborne wind energy tools. This toolchain was then used for performance assessment of two case studies.

This chapter summarizes the main results of the thesis according to the research questions from Chapter 3 and gives recommendations for future work in Section 9.6.

9.1. AWESPA and awesIO

AWESPA separates each step of the performance assessment procedures into modules. These modules work together since they follow the same interface. This makes it possible to swap and compare different models for a certain module as well.

Input and output data were standardized by defining common awesIO file formats for the main data exchanged between modules. The system configuration file defines the physical properties of the AWE system, the wind resource file defines the representative wind profiles and their occurrence frequencies, and the power curve file defines the power output used in later performance calculations. These files use a YAML-based structure, so each module can read and write the same type of data without relying on tool-specific naming conventions and structure.

Validation was included through predefined awesIO schemas. These schemas define the required fields, data structure, and naming conventions for each file type. Before a file is used in AWESPA or after one is created, it can be checked against the corresponding schema. This makes sure the files adhere to the standards.

9.2. Wind resource representation

The wind resource module converts atmospheric data into representative wind profile shapes using k-means clustering. The resulting representative profiles and their occurrence frequencies provide the wind resource input for power and annual energy production calculations.

The comparison with Lidar data at Bangor Erris showed that ERA5 represents the first-order wind climate reasonably well. ERA5 captured the general mean wind profile shape, gave similar Weibull scale parameters, and reproduced the dominant wind-direction sector. This makes ERA5 suitable for first-order performance assessment.

The comparison also showed that ERA5 smooths the local wind resource. ERA5 gave lower wind-speed variability and a narrower range of profile shapes than the Lidar data. High-shear, low-shear, and low-level-jet-like profiles occurred more often in the Lidar data than in ERA5. Site measurements therefore remain useful for validating ERA5-based airborne wind energy assessments.

9.3. Power model comparison

The power model comparison showed that the predicted performance depends strongly on the model assumptions. The original Luchsinger model and the soft-kite variant use the same simplified cycle structure, but differ in their reel-in phase. In the original model, the reel-in elevation angle is prescribed, while in the soft-kite variant it follows from the depowered lift-to-drag ratio.

The soft-kite Luchsinger model gave a more optimistic power estimate than the inertia-free quasi-steady model. At a reference wind speed of $15.72, \mathrm{m}, \mathrm{s}^{-1}$, its cycle power was $64, \%$ higher. This difference was mainly caused by the inclusion of gravitational loading and the more detailed reel-in phase in the quasi-steady model.

The inertia-free quasi-steady model was extended by replacing the single reel-out elevation angle with multiple elevation-angle setpoints along the traction phase. This allows the optimizer to vary the reel-out trajectory and account for the trade-off between lower-elevation operation and higher wind speeds at higher altitude.

9.4. Farm-level operation

The farm model investigated whether closely spaced wind-aligned systems must operate in sync when their reel-in and reel-out regions are adjacent. A geometric maximum phase offset was derived and used as a constraint in the phase optimization. This allowed limited asynchronous operation while maintaining separation between neighboring systems.

The results showed that synchronized operation is not always required for wind-aligned systems. However, the power smoothing benefit depends on the operating condition because the wind speed changes the reel-in and reel-out envelopes, the cycle timing, and therefore the allowed phase offset. Compared with the wind-perpendicular layout, the wind-aligned layout can give a larger peak-to-peak farm power variation when this phase constraint is restrictive.

For larger farms, the difference between the wind-aligned and wind-perpendicular layouts becomes smaller. This is because smaller phase differences between neighboring systems are sufficient to distribute the power cycles over one period. The wind-aligned constraint therefore remains relevant, but it is less restrictive when more systems are included.

9.5. Case studies

The case studies showed how system sizing, wind resource representation, and farm operation affect the predicted performance of the selected airborne wind energy systems. For the $100\ \mathrm{kW}$ isolated system at Bangor Erris, kite size, maximum tether force, and maximum generator power all affected the power curve and annual energy production. Larger kites gave higher power below rated wind speed, while higher tether-force limits and larger generator powers mainly affected the higher wind-speed range. Since no cost model was included, the selected system was chosen as a representative $100\ \mathrm{kW}$ design rather than as a cost-optimal design.

The selected system used a flat kite area of $160\ \mathrm{m}^2$, a maximum tether force of $80\ \mathrm{kN}$, and a maximum generator power of $170\ \mathrm{kW}$. For Bangor Erris, this system gave an annual energy production of $490.4\ \mathrm{MWh}$ with the fitted power-law profile and $488.6\ \mathrm{MWh}$ with the 8 representative wind profiles. The difference between the two wind resource representations was small for this site and system. The corresponding capacity factor for the power-law profile was $55.0\ \%$.

The same $100\ \mathrm{kW}$ system gave a lower annual energy production at Maasvlakte 2 than at Bangor Erris. For Maasvlakte 2, the annual energy production was $441.1\ \mathrm{MWh}$ with the fitted power-law profile and $436.1\ \mathrm{MWh}$ with the 8 representative wind profiles. The clustered profiles therefore gave a lower annual energy production than the fitted power-law profile for this site.

The $1\ \mathrm{MW}$ farm was modeled as a 10-system farm using the selected $100\ \mathrm{kW}$ system. Since wake losses were not included, the annual energy production was calculated by multiplying the single-system result by the number of systems. This gave $4411\ \mathrm{MWh}$ with the fitted power-law profile and $4361\ \mathrm{MWh}$ with the clustered wind profiles. These values are therefore idealized estimates for the 10-system layout.

The farm operation analysis showed that phase shifting reduced the variation in aggregated farm power output. For most of the considered wind speeds, the unconstrained and wind-aligned constrained cases gave the same peak-to-peak variation. Only when the maximum phase offset became restrictive did the wind-aligned layout give a larger peak-to-peak variation than the unconstrained case. This shows that systems can be placed closer together in the wind-aligned direction, increasing farm power density, while still maintaining a smooth farm power output for most operating conditions.

9.6. Recommendations

AWESPA and awesIO should be extended to cover a larger part of the performance assessment process. This includes adding an economic module so that system designs can be compared in terms of levelized cost of energy, as well as annual energy production. It also includes further development of the awesIO system configuration file and the addition of more power models with different levels of fidelity. This would make it possible to compare low-fidelity and higher-fidelity predictions in the same framework and to select models based on the required accuracy and computational cost.

The interaction between wakes in AWE farms should be studied in more detail. In the current farm model, downstream systems operate in wakes from upstream systems, but the effect of these wakes on power production is not included. Future work should therefore include wake interaction between systems and investigate wake avoidance strategies, such as adapted phase timing or trajectory changes.

Coupled optimization between the power model and the farm model should also be considered. In this thesis, the single-system cycle is optimized first, after which the farm phase offsets are optimized. A coupled approach could optimize the cycle timing, tether-length range, elevation angles, and phase offsets together. This may reduce farm power variation even further while maintaining the required separation between systems.

References

[1] K. Calvin et al., “Climate Change 2023: Synthesis Report,” Intergovernmental Panel on Climate Change (IPCC), Tech. Rep., Jul. 2023. DOI: 10.59327/IPCC/AR6-9789291691647.

[2] M. Diehl, “Airborne Wind Energy: Basic Concepts and Physical Foundations,” in *Airborne Wind Energy*, U. Ahrens, M. Diehl, and R. Schmehl, Eds., Berlin, Heidelberg: Springer, 2013, pp. 3–22. DOI: 10.1007/978-3-642-39965-7_1.

[3] R. Schmehl, M. Noom, and R. van der Vlugt, “Traction Power Generation with Tethered Wings,” in *Airborne Wind Energy*, U. Ahrens, M. Diehl, and R. Schmehl, Eds., Berlin, Heidelberg: Springer, 2013, pp. 23–45. DOI: 10.1007/978-3-642-39965-7_2.

[4] R. Schmehl, M. Rodriguez, L. Ouroumova, and M. Gaunaa, “Airborne Wind Energy for Martian Habitats,” in *Adaptive On- and Off-Earth Environments*, A. Cervone, H. Bier, and A. Makaya, Eds., Cham: Springer International Publishing, 2024, pp. 145–197. DOI: 10.1007/978-3-031-50081-7_7.

[5] M. Schelbergen, “Power to the airborne wind energy performance model,” Ph.D. dissertation, TU Delft, 2024. DOI: 10.4233/uuid:353d390a-9b79-44f1-9847-136a6b880e12.

[6] L. Thimm, “Wind Resource Parametrisation and Power Harvesting Estimation using AWERA - the Airborne Wind Energy Resource Analysis tool,” MSc Thesis, University of Bonn, Apr. 2023. DOI: 10.5281/zenodo.7848071.

[7] *Awegroup/AWERA*, Airborne Wind Energy Research Group, Oct. 2024.

[8] *Awegroup/AWESPA*, Airborne Wind Energy Research Group, May 2026.

[9] *Awegroup/awesIO*, Airborne Wind Energy Research Group, May 2026.

[10] R. H. Luchsinger, “Pumping Cycle Kite Power,” in *Airborne Wind Energy*, U. Ahrens, M. Diehl, and R. Schmehl, Eds., Berlin, Heidelberg: Springer, 2013, pp. 47–64. DOI: 10.1007/978-3-642-39965-7_3.

[11] R. van der Vlugt, J. Peschel, and R. Schmehl, “Design and Experimental Characterization of a Pumping Kite Power System,” in *Airborne Wind Energy*, Sep. 2013, pp. 403–425. DOI: 10.1007/978-3-642-39965-7_23.

[12] M. Schelbergen and R. Schmehl, “Validation of the quasi-steady performance model for pumping airborne wind energy systems,” *Journal of Physics: Conference Series*, vol. 1618, no. 3, p. 032 003, Sep. 2020. DOI: 10.1088/1742-6596/1618/3/032003.

[13] H. Hersbach et al., “The ERA5 global reanalysis,” *Quarterly Journal of the Royal Meteorological Society*, vol. 146, no. 730, pp. 1999–2049, 2020. DOI: 10.1002/qj.3803.

[14] P. Bechtle, M. Schelbergen, R. Schmehl, U. Zillmann, and S. Watson, “Airborne wind energy resource analysis,” *Renewable Energy*, vol. 141, pp. 1103–1116, Oct. 2019. DOI: 10.1016/j.renene.2019.03.118.

[15] H. He et al., “Characteristics of Wind Profiles for Airborne Wind Energy Systems,” *Energies*, vol. 18, no. 9, p. 2373, Jan. 2025. DOI: 10.3390/en18092373.

[16] R. van der Vlugt, A. Bley, M. Noom, and R. Schmehl, “Quasi-steady model of a pumping kite power system,” *Renewable Energy*, vol. 131, pp. 83–99, Feb. 2019. DOI: 10.1016/j.renene.2018.07.023.

[17] G. Gualtieri, “Reliability of ERA5 Reanalysis Data for Wind Resource Assessment: A Comparison against Tall Towers,” *Energies*, vol. 14, no. 14, p. 4169, Jan. 2021. DOI: 10.3390/en14144169.

[18] M. Sommerfeld, M. Dörenkämper, J. De Schutter, and C. Crawford, “Impact of wind profiles on ground-generation airborne wind energy system performance,” *Wind Energy Science*, vol. 8, no. 7, pp. 1153–1178, Jul. 2023. DOI: 10.5194/wes-8-1153-2023.

[19] M. Schelbergen, P. C. Kalverla, R. Schmehl, and S. J. Watson, “Clustering wind profile shapes to estimate airborne wind energy production,” *Wind Energy Science*, vol. 5, no. 3, pp. 1097–1120, Aug. 2020. DOI: 10.5194/wes-5-1097-2020.

[20] F. Pedregosa et al., “Scikit-learn: Machine Learning in Python,” *Journal of Machine Learning Research*, vol. 12, Jan. 2012.

[21] V. van Deursen, “Dynamic Simulation Techniques for Airborne Wind Energy Systems,” MSc Thesis, TU Delft, Aug. 2024.

[22] M. L. Loyd, “Crosswind kite power (for large-scale wind power production),” *Journal of Energy*, vol. 4, no. 3, pp. 106–111, May 1980. DOI: 10.2514/3.48021.

[23] I. Argatov, P. Rautakorpi, and R. Silvennoinen, “Estimation of the mechanical energy output of the kite wind generator,” *Renewable Energy*, vol. 34, no. 6, pp. 1525–1532, Jun. 2009. DOI: 10.1016/j.renene.2008.11.001.

[24] U. Fechner and R. Schmehl, “Model-Based Efficiency Analysis of Wind Power Conversion by a Pumping Kite Power System,” in *Airborne Wind Energy*, U. Ahrens, M. Diehl, and R. Schmehl, Eds., Berlin, Heidelberg: Springer, 2013, pp. 249–269. DOI: 10.1007/978-3-642-39965-7_14.

[25] M. N. Noom, “Theoretical Analysis of Mechanical Power Generation by Pumping Cycle Kite Power Systems,” MSc Thesis, TU Delft, May 2013.

[26] M. J. Venegas, “Path Optimization of a Pumping Kite System,” MSc Thesis, TU Delft, 2017.

[27] H. Johnson, “Kite Farm Simulation,” MSc Thesis, TU Delft, 2019.

[28] M. Ranneberg, D. Wölfle, A. Bormann, P. Rohde, F. Breipohl, and I. Bastigkeit, “Fast Power Curve and Yield Estimation of Pumping Airborne Wind Energy Systems,” in *Airborne Wind Energy: Advances in Technology Development and Research*, R. Schmehl, Ed., Singapore: Springer, 2018, pp. 623–641. DOI: 10.1007/978-981-10-1947-0_25.

[29] M. Talmar, “Quasi-Steady Model Enhancements of Airborne Wind Energy Systems,” MSc Thesis, TU Delft, Jul. 2021.

[30] A. Bertozzi, C. Vergara, S. Reuchlin, and R. Schmehl, “Pumping cycle optimization of soft-wing airborne wind energy systems using quasi-steady simulations,” *Journal of Physics: Conference Series*, vol. 3224, no. 9, p. 092 017, Mar. 2026. DOI: 10.1088/1742-6596/3224/9/092017.

[31] O. Cayon, V. van Deursen, and R. Schmehl, “Translational Dynamics of Bridled Kites: A Reduced-Order Model in the Course Reference Frame,” *Wind Energy Science Discussions*, pp. 1–38, Oct. 2025. DOI: 10.5194/wes-2025-205.

[32] R. C. da Costa, M. C. R. M. Fernandes, L. A. C. Roque, L. T. Paiva, D. B. M. M. Fontes, and F. A. C. C. Fontes, “Airborne wind energy farms layout: Synchronization strategies for power smoothing,” *Optimization Letters*, Jan. 2026. DOI: 10.1007/s11590-025-02273-7.

[33] P. Faggiani, “Pumping Kites Wind Farm,” Ph.D. dissertation, TU Delft, 2015.

[34] P. Faggiani and R. Schmehl, “Design and Economics of a Pumping Kite Wind Park,” in *Airborne Wind Energy: Advances in Technology Development and Research*, Apr. 2018, pp. 391–411. DOI: 10.1007/978-981-10-1947-0_16.

[35] G. Licitra, J. Koenemann, A. Bürger, P. Williams, R. Ruiterkamp, and M. Diehl, “Performance assessment of a rigid wing Airborne Wind Energy pumping system,” *Energy*, vol. 173, pp. 569–585, Apr. 2019. DOI: 10.1016/j.energy.2019.02.064.

[36] IEA Wind TCP, *Task 47: Airborne Wind Energy*, Accessed: 2026-07-09.: https://iea-wind.org/task47/.

[37] ECMWF, *ERA5: Compute pressure and geopotential on model levels*, Accessed on 2026-04-29.: https://confluence.ecmwf.int/display/CKB/ERA5%3A+compute+pressure+and+geopotential+on+model+levels%2C+geopotential+height+and+geometric+height.

[38] ECMWF, *L137 model level definitions*, Accessed on: 2026-04-29.: https://confluence.ecmwf.int/display/UDOC/L137+model+level+definitions.

[39] Vaisala, *WindCube Ground-based vertical profiling lidar Datasheet*, Accessed on: 2026-04-29, Apr. 2026.: https://docs.vaisala.com/api/khub/documents/V4Ktyr0Yr8oCiZtXtLm49A/content?Ft-Calling-App=ft%2Fturnkey-portal&Ft-Calling-App-Version=5.3.3#.

[40] *Awegroup/TUDELFT_V3_KITE: A data repository containing all relevant analysis and specification of the TUDELFT V3 Leading-Edge-Inflatable kite*, Airborne Wind Energy Research Group, 2025.

[41] J. Oehler and R. Schmehl, "Aerodynamic characterization of a soft kite by in situ flow measurement," *Wind Energy Science*, vol. 4, no. 1, pp. 1–21, Jan. 2019. DOI: 10.5194/wes-4-1-2019.

[42] T. Haas, J. De Schutter, M. Diehl, and J. Meyers, "Large-eddy simulation of airborne wind energy farms," *Wind Energy Science*, vol. 7, no. 3, pp. 1093–1135, May 2022. DOI: 10.5194/wes-7-1093-2022.

A
awesIO system input files

A.1. Kitepower V3.25


```yaml
metadata:
  name: TU Delft V3.25
  description: Config for the TU Delft V3.25 soft kite pumping ground gen system
  note: None
  awesIO_version: 0.1.0
  schema: system_schema.yml

assembly:
  connectivity_matrix:
    - [K1, T1]
    - [T2, G1]

components:
  kites:
    - name: TU Delft V3.25
      type: soft_kite_kcu
      version: 1.0
      joints:
        kcu_joint:
          id: K1
      wing:
        name: TU Delft V3.25
        type: LEI_soft_kite
        version: 1.0
        structure:
          projected_surface_area: 19.75 # m2
          mass: 10.6 # kg, wing and bridle mass lumped together
      bridle:
        name: TU Delft V3.25 Bridle
        type: LEI_soft_kite_bridle
        version: 1.0
        structure:
          mass: 0.0
      control_system:
        name: TU Delft V3.25 KCU
        type: kite_control_unit
        version: 1.0
        structure:
          mass: 8.4 # kg

  tethers:
```

```yaml
    - name: Kitepower tether
      type: non_conductive_tether
      version: 1.0
      joints:
        upper_tether_joint:
          id: T1
        lower_tether_joint:
          id: T2
      structure:
        length: 500
        diameter: 0.006
        density: 724.0
        max_tether_force: 8200.0 # N
        conductive: false

  ground_station:
    name: Kitepower Ground Station
    type: pumping_ground_gen_station
    version: 1.0
    joints:
      drum_joint:
        id: G1
    structure:
      mass: 0.0
    drums:
      - type: electric_winch
        max_tether_speed: 10.0
    generators:
      - type: permanent_magnet_synchronous
        max_power: 40000.0 # W
        efficiency: 1.0
    storages:
      - type: battery_bank
        efficiency: 1.0
```


A.2. Meridional 100 kW ground gen system


```yaml
metadata:
  name: Meridional 100 kW ground gen system
  description: None
  note: None
  awesIO_version: 0.1.0
  schema: system_schema.yml
assembly:
  connectivity_matrix:
  - - K1
    - T1
  - - T2
    - G1
components:
  kites:
  - name: 100kW kite
    type: soft_kite_kcu
    version: 1.0
    joints:
      kcu_joint:
        id: K1
    wing:
      name: 160 m^2 wing
      type: LEI_soft_kite
      version: 1.0
```

```yaml
    structure:
      projected_surface_area: 126.4
      mass: 53.0
  bridle:
    name: 100kW bridle
    type: LEI_soft_kite_bridle
    version: 1.0
    structure:
      mass: 0.0
  control_system:
    name: 100 kW system KCU
    type: kite_control_unit
    version: 1.0
    structure:
      mass: 20.0
tethers:
- name: 100 kW Tether
  type: non_conductive_tether
  version: 1.0
  joints:
    upper_tether_joint:
      id: T1
    lower_tether_joint:
      id: T2
  structure:
    length: 600
    diameter: 0.014
    density: 724.0
    max_tether_force: 80000.0
    conductive: false
ground_station:
  name: 100 kW Ground Station
  type: pumping_ground_gen_station
  version: 1.0
  joints:
    drum_joint:
      id: G1
  structure:
    mass: 0.0
  drums:
  - type: electric_winch
    max_tether_speed: 15.0
  generators:
  - type: permanent_magnet_synchronous
    max_power: 170000.0
    efficiency: 0.95
  storages:
  - type: battery_bank
    efficiency: 0.9
```