



Optimal flight pattern debate for airborne wind energy systems: circular or figure of eight?

Dylan Eijkelhof¹, Nicola Rossi^{1,2}, and Roland Schmehl¹

¹Faculty of Aerospace Engineering, Delft University of Technology, 2629 HS Delft, the Netherlands

²Università degli Studi di Trento, Via Sommarive, 9, 38123 Trento, Italy

Correspondence: Dylan Eijkelhof (d.eijkelhof@tudelft.nl)

Received: 31 October 2024 – Discussion started: 18 December 2024

Revised: 28 February 2026 – Accepted: 7 March 2026 – Published: 17 April 2026

Abstract. This computational study compares the performance of circular and figure-of-eight flight patterns for fixed-wing ground-generation airborne wind energy (AWE) systems using a PID-based basic controller that effectively controls the kite during each pattern's pumping cycle in a MATLAB[®] Simulink[®] environment. A simple, adjustable control framework enables a steady analysis within consistent operational parameters, allowing for fair comparisons of power output, power quality, ground surface area requirements, and structural load impacts. Unlike for small lightweight kites, the simulation results reveal that when using the 150 m² MegAWES reference kite, the power is not similar over different patterns. The circular flight pattern achieves the highest cycle-averaged power output in a smaller operational area, making it advantageous for maximising energy within limited spatial constraints. Conversely, the figure-of-eight down-loop pattern demonstrates superior power quality with lower power peaks and lower expected structural fatigue due to a reduced cyclic aerodynamic load frequency and amplitude, supporting greater operational stability and system longevity. The figure-of-eight up-loop does not stand out on any of the metrics considered in this work. This study offers insights into the trade-offs between energy output, efficiency, and structural demands associated with each flight path, providing a foundation for future AWE flight path selection and control strategy optimisations.

1 Introduction

Airborne wind energy (AWE) is an emerging renewable energy technology aiming to harness the power available in higher-altitude winds. Unlike conventional wind turbines, the operating pattern is not fixed. A conventional wind turbine has rotating blades connected to a hub. AWE devices use tether kites, which have a much larger potential operating domain, which allows for better adaptation to the available wind resource but comes with complex controller challenges.

One shared characteristic of the different existing AWE systems is that they all fly in a closed pattern. However, the best shape of this pattern has not yet been decided on. The debate about the optimal flight pattern dates back to the beginning of the technology development in the first decade of the 21st century. The industry appears undecided on a single specific pattern, with circular and figure-of-eight paths both

prevalent. Tether winding is mentioned as one of the key reasons for the latter (Houska and Diehl, 2006; Licitra et al., 2019). Other factors might also influence the choice, as discussed in this study.

In particular, three main patterns are currently flown in AWE. These are divided into two shapes: a circle and a figure of eight. The latter can be flown in two directions, one where the outer segments of the figure are flown upwards (up-loop) and one where they are flown downwards (down-loop). Each pattern offers distinct advantages and challenges in terms of stability, energy capture, and system complexity.

Licitra et al. (2019) conducted an in-depth performance assessment of a small-scale low-mass fixed-wing AWE system in pumping mode using an optimal control framework. Their findings challenge the common debate on the superiority of lemniscate versus circular trajectories by showing that, under equal boundary conditions, both patterns yield

comparable power output. The study further emphasises that practical considerations, such as flight area requirements and mechanical stress, may ultimately drive the choice of flight pattern rather than efficiency alone. This work expands on their conclusions using a much heavier and larger kite, a detailed controller environment, and a user-friendly simulation framework for open-source collaborations.

Specifically, this study focuses on the development of a basic flight controller for a fixed-wing AWE system, designed in a MATLAB[®] Simulink[®] environment, with the purpose of comparing different flight patterns flown by heavy kites. The controller is engineered to replicate the entire pumping cycle of a ground-generation kite system, simulating the critical phases of power generation (traction) and its recovery (retraction). By modelling the kite's flight path and control dynamics within this virtual environment, the model can provide valuable insights into the performance and efficiency of different flight strategies.

Control and navigation of AWE systems often follow a hierarchical control structure, as commonly proposed for ground-generation rigid-wing systems (Rapp et al., 2019; Vermillion et al., 2021). This cascaded approach effectively breaks down the complex control problem into manageable sub-problems, allowing for optimal customisation of each control layer. The higher-level control is typically managed by a finite state machine, which supervises transitions across various flight operations, including launch, traction, retraction, and safety manoeuvres. This control architecture is illustrated in Fig. 1.

For path following, target points are frequently utilised to guide the kite closely along predefined trajectories or to give a looser guidance approach with a planned path. In van der Vlugt et al. (2013), four waypoints were used to steer a soft kite in stable figure-of-eight down-loop manoeuvres. Erhard and Strauch (2013); Fagiano et al. (2014) demonstrate periodic figure-of-eight flight paths by alternating between two target points, also known as the bang–bang control approach used in practice by Skysails at the time. Fechner and Schmehl (2016) added a geometric construction of the planned pumping cycle flight path to the four-waypoint controller of van der Vlugt et al. (2013). Although these methods result in flying less defined path shapes, making it harder to quantify controller performance, their simplicity and computational efficiency make them appealing for real-world applications and sometimes even more superior in power production.

Another method is continuous path parameterisation, where the entire path is discretised into points which are tracked by the navigation system. A widely studied controller logic using this method is given by the L_1 guidance logic. It calculates the required centripetal acceleration to define a reference roll angle, facilitating precise path following and enhanced system performance. This logic was initially proposed by Park et al. (2004) and later refined by Fernandes et al. (2022) to the so-called L_0 guidance logic. Both methods

are widely adopted non-linear path-following strategies in aerial vehicle navigation. Fernandes et al. (2022) also show that the L_0 guidance logic ensures asymptotic stability and is particularly well-suited for AWE systems.

Additionally, controlling the kite's attitude and dynamics is critical in ensuring stable and efficient AWE system operation. For flexible-membrane wings, the turn rate law proposed by Erhard and Strauch (2013) links kite steering inputs to its course rate. However, this law does not apply to fixed-wing systems, where dynamic forces and moments are much more prominent.

An example of a more advanced control strategy can be seen in the flight controller presented by Rüterkamp and Sieberling (2013), where the longitudinal and lateral motions of the kite are decoupled and controlled independently. This approach involves breaking down the figure-of-eight curve into discrete target points and using linear controllers to convert these into reference roll angles. However, challenges remain, particularly in maintaining the tether tension within safe limits.

Other advanced control strategies, which include the use of continuous path parameterisation, exist too. Jehle and Schmehl (2014) employ a cascaded controller using feedback linearisation of an empirical steering-to-yaw rate model combined with a non-linear tracking law on the unit sphere to align the kite's flight path with a prescribed trajectory. Rapp et al. (2019) combine a PID controller for state-ideal path error correction with non-linear dynamic inversion to achieve the desired kite attitude. Lyapunov-based controllers, as developed by Li et al. (2015), focus on controlling kite attitude and rate error dynamics through differential and proportional control laws, addressing both translational and rotational motions independently.

The MegAWES 3 MW fixed-wing reference kite (Eijkelhof and Schmehl, 2022) is used in conjunction with the new improved flight controller developed by Rossi (2023). The earlier implemented controller proposed by Eijkelhof and Schmehl (2022) is complex and requires many controller parameters to be tuned for an optimal trajectory. This makes the controller unsuitable for a comprehensive study of the flight pattern. Therefore, a simpler, easier-to-tune controller is desired.

The main objective of the present work is the development of a simulation framework that allows comparing different flight patterns using the same operational environment, kite design, and flight controller. The path planning is based on the described L_0 logic (Fernandes et al., 2022), and the attitude control of the kite consists of a cascaded PID control loop based on Eijkelhof (2019). To have a consistent system for the analysis of the flight paths, the MegAWES reference kite described by Eijkelhof and Schmehl (2022) is used. A new winch design and controller are implemented as developed in Hummel et al. (2024b). The winch has been better sized for the system size of the MegAWES reference kite,

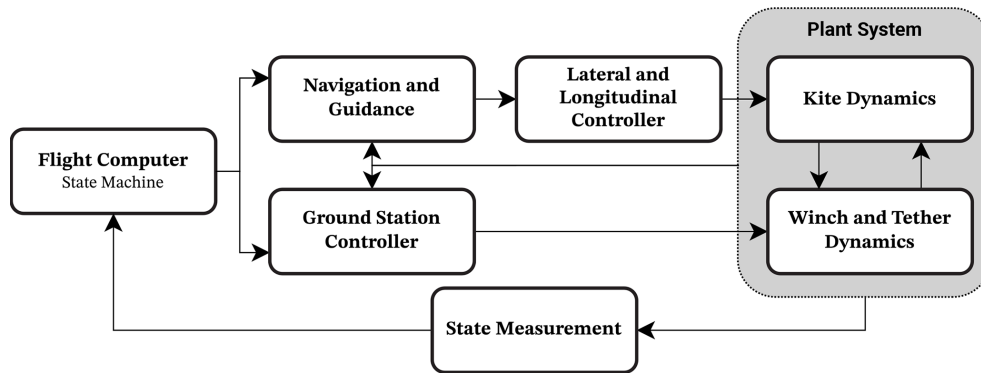


Figure 1. Control architecture representation, featuring outer and inner control loops, for a ground-generation AWE system.

and the winch controller has been changed to follow a power-optimal strategy.

The remainder of the paper is organised as follows. In Sect. 2, a comprehensive overview of the numerical framework is provided, the design and implementation of the flight controller in MATLAB® Simulink® are described, the winch controller is introduced, and the strategy for obtaining comparative flight paths and the reference system are presented. In Sect. 3, a detailed analysis of the simulation results comparing the circular and two figure-of-eight flight paths is presented. Finally, in Sect. 4, a conclusion is drawn based on the results.

2 Methodology

In this section, the numerical framework is described, which is used to compare the three different flight paths. First, the flight controller is discussed, split up into lateral and longitudinal control, and the reference systems are described. Second, the required winch model is discussed. Third, the comparison strategy is explained, and, last, the additional system parameters are provided, including the reference kite model.

2.1 Flight controller

The development of the flight control system for the fixed-wing AWE system is crucial for ensuring reliable and efficient operation. This system is designed to handle the intricate dynamics of tethered flight, which involves managing both the kite's lateral and longitudinal motions. The control architecture is based on a cascaded-control approach, which separates the control tasks into more manageable sub-problems, allowing for precise tuning and optimisation of each component. This section outlines the reference systems used to model the kite's flight dynamics, the methodologies employed in designing the state machine for switching between phases, and the lateral and longitudinal control systems of the kite.

Figure 2 illustrates the control state machine structure designed for the system, detailing the transitions and control objectives for each flight phase. This state machine framework enables smooth transitions between traction, retraction, and other phases, optimising power output and stability under varying operational conditions.

In the traction phase, lateral and pitch control are focused on maintaining the maximum angle of attack for optimal energy capture, while winch control enforces an optimal force set point to manage tether tension. The transition phase from retraction to traction adjusts lateral and pitch controls similarly to the traction phase, driving the angle of attack to its peak to ensure longitudinal stability until the kite reaches the designated pattern target point to go back into traction. During the retraction phase, lateral control reorients to a pre-determined retraction target point, pitch control tracks a descent angle, and winch control reduces the tether tension to a reference force set point, effectively preparing the system for the next cycle.

2.1.1 Coordinate frames

Various coordinate frames are used to model the entire system and effectively represent the many vector quantities that are involved. In most cases, these coordinate frames align with those commonly used in the flight mechanics literature, but some others are specially introduced for the AWE domain. This section presents a description and graphical representation of the reference frames used throughout this study.

Figure 3 presents the inertial (e), wind (w), body (B), and tangential plane (τ) reference frames. In this work, the wind direction is aligned with the inertial x axis ($\varphi_w = 0$). Because the wind reference frame has the x axis aligned with the wind direction, both inertial and wind x axes coincide. The inertial reference frame is defined as a north-east-down (NED) coordinate system with the ground station placed in the origin and the z axis pointing into the ground. This definition is consistent with the standard approach in aerospace literature and is also implemented in the internal aerospace Simulink block.

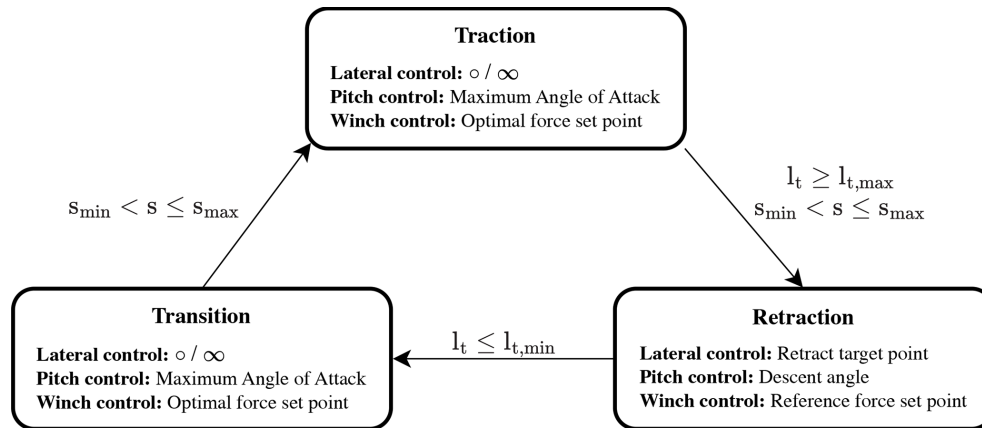


Figure 2. The control state machine describing the transitions and control objectives for each flight phase. The triggers to switch between phases are shown next to the arrows, where l_t is the tether length, and s is the dimensionless path variable resembling the position along the pattern.

To guide the kite through the wind window and to define the desired flight path for the traction phase, it is convenient to define a reference frame centred in the kite's centre of gravity (i.e. kite position) and tangential to the flight sphere. This frame, shown in Fig. 3, can be obtained through a sequence of transformation matrices starting from the wind coordinate system.

The body coordinate system is fixed according to the position and orientation of the kite. In this study, the kite is considered rigid, and any real-time deflection or deformation is not taken into account.

2.1.2 Lateral control

The lateral controller requires a controller objective. This objective is determined by the path-following module. This work utilises an adapted version of the L_1 controller logic, a discretised L_0 controller logic. As both methods are similar in principle, the L_1 method is described first, from which the L_0 method is then derived. The concept is visualised in Fig. 4.

The L_1 controller logic is typically a control strategy used in guidance systems in flight controllers for path following, where a dynamical system must track a desired trajectory. The key principle behind the L_1 logic is that it defines a distance metric, often referred to as the L_1 distance, which measures the separation between the actual position of the system (e.g. kite or aircraft) and a designated target point on the intended path. This distance is often used as a control parameter to ensure that the system moves toward the desired trajectory by looking either closer or further ahead to determine an optimal target position to track. The L_1 controller has only one adjustable design parameter, the L_1 distance, which can be tuned to achieve faster or smoother convergence to the desired path. This makes the strategy simple to tune and ideal in the context of this work. Park et al. (2004) demonstrate

the effectiveness of this tracking method and explain how to optimise the parameter across different scenarios, including straight lines, perturbed paths, and circular trajectories. They further employ an approach based on the Lyapunov invariant set theorem to guarantee asymptotic stability. This makes the L_1 control strategy particularly well-suited for non-linear systems, offering an advantage over traditional linear control methods.

Rather than specifying the L_1 distance as a control metric, a discretised interpretation of the L_0 distance is used (Fernandes et al., 2022). In this approach, the distance to the target point is calculated from the nearest point along the trajectory projected from the kite's position.

In contrast to conventional methods, this formulation broadens the field of attraction toward the desired course, guaranteeing that the guidance system consistently retains a target point on the path, even when the kite drifts further from the desired trajectory. When the L_1 distance is used, the controller becomes confused when the kite is further away from the trajectory than the specified L_1 distance. There are several ways to solve this, for example, by widening the distance when it cannot converge on a point on the trajectory. However, using the L_0 distance instead was deemed to be the most effective.

Due to the closed and intersecting geometry of certain flight paths, accurately selecting the target point becomes more complex. When dealing with a continuous curve, iterative techniques, such as Newton–Euler methods, are typically employed to determine both the nearest and the target points with a manageable number of iterations. However, for the discrete representation of the path as a finite set of points, an alternative approach is necessary.

In the studied scenario, the set of path points is discretised using the dimensionless path variable s over the interval from 0 to 2π . This ensures unique tracking points, independent of path shape. The path is discretised on the tangent plane and




Figure 3. Inertial (e), wind (w), body (B), and tangent plane (τ) reference systems. Visualisation of azimuthal (φ) and elevation (β) angles along with wind direction (φ_w).




Figure 4. Principle of L_1 and L_0 guidance logic. The closest point to the kite and the target point are highlighted in red. V_k is the kite velocity magnitude, a_c the required acceleration, and η the heading.

then projected onto a spherical surface. Using a path parameterisation dependent on this variable s , Cartesian path coordinates can be calculated and stored.

The discretisation of the circular path is straightforward given the definition of a circle. The figure of eight can be defined in several different ways, a Lissajous or (hyperbolic) lemniscate of Booth, for example. Lissajous figures are looping patterns formed by the combination of two perpendicular

harmonic oscillations. By choosing the right parameters, a shape similar to a figure of eight can be obtained. The lemniscate of Booth is a generalised version of the lemniscate of Bernoulli, which was first introduced by the Bernoulli brothers in 1694. Bernoulli’s lemniscate was then generalised in the 19th century by James Booth. The lemniscate of Booth can be described as the set of points where the product of distances to two foci of the hyperbolic function is constant.

When the hyperbolic function is rectangular, the lemniscate of Bernoulli is obtained (Gray, 1997). The lemniscate of Booth is used in this framework, and given the path variable s , the shape is defined in the tangent plane by the following coordinate expressions and then projected on the flight sphere:

$$x_{\text{fig8}}(s) = \frac{b_{\text{Booth}} \sin s}{1 + \left(\frac{a_{\text{Booth}}}{b_{\text{Booth}}}\right)^2 \cos^2 s}, \quad (1)$$

$$y_{\text{fig8}}(s) = \frac{a_{\text{Booth}} \sin s \cos s}{1 + \left(\frac{a_{\text{Booth}}}{b_{\text{Booth}}}\right)^2 \cos^2 s}, \quad (2)$$

where a_{Booth} and b_{Booth} define the height and width of the figure of eight, respectively. For this study, a discretisation of 50 points was chosen. This leads to a smooth path discretisation but is still computationally efficient.

Additionally, a distance array is introduced, containing the distances between the kite's current position and each point on the path. This array is organised cyclically, allowing the algorithm to continually search for the closest point as the kite follows the closed-loop trajectory. To select the correct nearest point, special attention is needed for complex paths like the figure of eight. In this shape, the kite crosses over multiple intersections, which can lead to wrong tracking points. To overcome this problem, the selection process uses a smaller subset of the distance array, starting from the previous time step. This proved to be a reliable method to choose the correct projected kite location on the trajectory.

Instead of choosing a distance value as the controller parameter, the tracking point is determined using a set number of elements ahead. After identifying the closest point, the target point is set as the n th element ahead in the sequence. This guarantees that the target point remains consistently defined as n points ahead of the nearest one, where n serves as a tunable discrete control parameter.

The lateral controller determines the actuator demands based on the tracking point given by the path-following module. The target vector is computed using the coordinates of both the kite and the target point, within the wind reference frame described in Sect. 2.1.1. Using the appropriate transformation matrices, the target vector and the kite's velocity are projected onto the tangent plane. This tangent plane is visible in Fig. 3 and is tangent to the sphere surface, with a radius equal to the straight tether length and the centre at the ground station.

The resulting projections are then used to derive the heading angle η . To avoid impractically sharp turns, which could lead to hazardous manoeuvres, the calculated heading angle η is restricted to $\pm \frac{\pi}{4}$. Finally, by considering all forces acting on the kite, the desired roll angle $\phi_{\tau, \text{des}}$, relative to the tangent plane, is determined as

$$\phi_{\tau, \text{des}} = \arcsin \left(\frac{a_c m + F_{g,y} \cos \phi_{\tau, \text{act}} + F_{g,z} \sin \phi_{\tau, \text{act}}}{F_L} \right), \quad (3)$$

where m is the kite mass, F_g is the gravitational force acting on the kite, $\phi_{\tau, \text{act}}$ is the actual roll angle, F_L is the lift force, and a_c is the lateral acceleration needed for the vehicle to trace a curved trajectory. This acceleration can be calculated as

$$a_c = 2 \frac{V_k^2}{L_1} \sin \eta, \quad (4)$$

where V_k is the kite velocity magnitude, and η is the angle between the kite velocity and the direction to the reference point.

This desired roll angle is then converted into a demand for the aileron deflection using multiple PI controllers, illustrated in Fig. 5.

2.1.3 Longitudinal control

One of the key factors influencing power extraction from the wind during crosswind flight of AWE systems is the tether force, which is directly linked to the kite's lift. To maintain proper tether tension, the kite's angle of attack must be controlled via the longitudinal or pitch controller.

Since the system's primary goal is to maximise wind energy extraction, the control of longitudinal flight dynamics aims to sustain a high C_L^3/C_D^2 ratio to optimise power generation. Loyd (1980) showed that the maximum theoretical power extraction is proportional to this ratio. Considering tether drag, the highest feasible ratio is practically lower than for untethered aircraft. Hence, it is essential to consider the total system's combined C_L^3/C_D^2 ratio, including the kite and tether drag. In general, a higher C_L^3/C_D^2 ratio corresponds to a higher angle of attack when including the tether drag, yielding more power. Hence, a simple approach is used, choosing a constant high angle of attack as a reference target during the power production phase.

As developing a new reference kite and aerodynamic model is beyond the scope of this paper, the existing aerodynamic data of the MegAWES kite are used (Eijkelhof and Schmehl, 2022). Using this, a target angle of attack of approximately 4° was chosen for the traction phase to keep the kite operating within the linear region. This corresponds to an approximate C_L of 1.8. Hence, the stall region is entirely avoided. While possible in numerical studies to neglect tether force control, in practical applications, it is critical to ensure that tether tension does not exceed its limits. During winch controller saturation, it may require changing the desired angle of attack to lower the aerodynamic performance, as done in Hummel et al. (2024a). This work does not include such a tether force controller, but the limit is taken into account in the optimisation objective.

The desired angle of attack α_{des} is then used as an objective for the pitch controller, which adjusts the kite's pitch by means of elevator deflection δ_e . This conversion is done by the PI cascaded loops visible in Fig. 5.




Figure 5. Flow diagram of the attitude and rate controller for both lateral and longitudinal control. The lateral controller needs a PID to accurately follow the desired input, and for the longitudinal controller, a PI controller suffices. The gains are tuned separately for traction and retraction. The integral part of the PI(D) controller is reset at the start of the specific flight phase to prevent windup.

While following circular or figure-eight paths, the longitudinal controller must compensate for external disturbances, including winch and tether forces and kite dynamics. Wind plays a significant role, particularly as the relative wind direction changes during ascent and descent. When the kite flies against the wind, the relative wind speed increases, raising α_a , while descending in the same direction as the wind decreases α_a . These oscillations, dependent on the velocity along the path, must be actively managed.

2.2 Winch controller

The winch, which includes the drum, drivetrain, and generator, is modelled as a single rotating body with inertia J and a constant radius R (Hummel et al., 2024b) positioned in the origin of the wind and inertial reference frames. The modelling approach simplifies the equations by ignoring viscous friction, as the friction torque is much smaller than the forces exerted by the tether on the winch and the control motor torque. Additionally, the radius is assumed constant, despite the spooling of the tether, because the changes in radius are small and slow, making their dynamic impact negligible. Two primary moments act on the winch: the tether force multiplied by the winch radius and the motor torque, both of which are used in the control strategy.

An optimal control curve is derived to define a control target. The ground tether force, the force at the winch, is calculated using a quasi-steady model, neglecting cosine losses and the kite’s mass (Schmehl et al., 2013)

$$F_t = C v_{a,r}^2, \tag{5}$$

where

$$C = \frac{1}{2} \rho S C_L E_{eq}^2 \left(1 + \frac{1}{E_{eq}} \right)^{\frac{3}{2}}, \tag{6}$$

where F_t is the ground tether force, $v_{a,r}$ the radial component of the apparent wind velocity, ρ the air density, S the reference wing area, C_L the lift coefficient, and E_{eq} the equivalent lift-to-drag ratio (including tether drag).

Instead of modelling individual forces like gravity or wind speed variation, an equivalent wind speed is used to account for these effects. Choosing an optimal reel-out factor f^* of 1/3 (Loyd, 1980), the motor torque on the winch aims to

balance out the torque from the tether force at steady state. The reel-out factor is defined as reeling velocity divided by the wind speed. This results in a winch control curve, often called the “optimal force-squared speed manifold”, where the motor torque T_m is the product of the optimal tether force and winch radius

$$T_m = 4C v_t^2 R, \tag{7}$$

where v_t is the tether reeling speed.

In Hummel et al. (2024b), the winch is sized for the MegAWES kite, as shown in Table 2. The given size parameters ensure a robust and steady response of the winch when trying to stay on the optimal control curve defined by Eqs. (5), (6), and (7).

2.3 Comparative flight path strategy

To evaluate the flight patterns, the following key criteria are investigated: cycle-averaged power, power oscillations (power quality), and projected ground area.

Cycle power is the average power output over a complete pumping cycle. It serves as a primary measure of system efficiency, representing the overall energy production capability of the system. Higher cycle-averaged power indicates more efficient wind energy extraction. The cycle power can be calculated using only the numerator of Eq. (9).

Power oscillations (power quality) reflect the stability and consistency of power generation. A system with lower power oscillations provides more stable power output, which is crucial for integration into the electrical grid. Hence, minimising power fluctuations is a desirable characteristic in AWE systems. The power quality is quantified using the peak-to-average-power ratio (PAPR) during the traction phase. It is defined as the ratio between the maximum instantaneous mechanical power and the average power over the traction interval:

$$PAPR = \frac{\max_{t \in T_{traction}} P_{mech}(t)}{\frac{1}{T_{traction}} \int_{T_{traction}} P_{mech}(t) dt}, \tag{8}$$

where $P_{mech}(t)$ is the instantaneous power at time t , and $T_{traction}$ denotes the time interval corresponding to the traction phase.

The projected ground area provides an estimate of the space required for each system when deployed in a farm configuration. It is crucial to determine how dense systems can be arranged within a given surface area, directly influencing the potential energy density of the farm. Additionally, it offers a preliminary assessment of the necessary safety buffers, which are vital for mitigating the risk of system interference and ensuring operational safety. The projected ground area is defined as the area of the two-dimensional projection of the flight trajectory onto the ground plane. Additionally, the surface area power density is used to evaluate the average power output normalised by the projected ground area:

$$P_d = \frac{\frac{1}{T_{\text{cycle}}} \int_{T_{\text{cycle}}} P_{\text{mech}}(t) dt}{\text{PGA}}, \quad (9)$$

where P_d is the surface area power density, PGA is the projected ground area, and T_{cycle} denotes the time interval corresponding to the entire pumping phase.

To support the comparative analysis of circular and figure-of-eight flight patterns, this study employs the covariance matrix adaptation evolution strategy (CMA-ES) (Hansen, 2006). The CMA-ES method iteratively refines parameters by minimising a specified objective function, allowing for targeted adjustments in both controller settings and flight path characteristics.

Rather than aiming for a global optimum, the optimisation serves as a strategic exploration of the design space, focusing on identifying representative configurations for each flight pattern. This suboptimisation approach thus provides a balanced set of configurations that enable a fair and insightful comparison of the performance, stability, and structural implications of different flight patterns, aligning directly with the study's goal of understanding each pattern's relative advantages and limitations. This means the solution should be in no way interpreted as an optimal condition. This method works well for different kite designs and scales where the initial condition is hard and time-consuming to determine. The numerical exploration using CMA-ES allows for finding a parameter set which converges, which is key to proceeding with any optimisation process. Gradient-based optimisation methods can effectively yield an optimised solution when an appropriate initial condition is available. However, despite the existence of techniques aimed at reducing the likelihood of suboptimal convergence, it remains inherently uncertain whether the obtained local optimum corresponds to the global optimum. Gradient-based methods are considered more suitable for a more detailed analysis of a single pattern.

Table 1 provides a set of system parameters which could be varied by the optimiser. The main objective is to maximise the average cycle power (mean(P_{mech})). The imposed soft constraints are formulated based on exceeded maximum thresholds, which are carefully selected to restrict the design space. Their implementation is detailed in Appendix A.

Table 1. Operational and controller parameters varied during optimisations. K_p , K_i , and K_d are the PID gains.

Controller	Flight path	Tether
K_p , K_i winch (retraction)	Figs a_{Booth} *	Tether length max
K_p roll rate (traction)	Figs b_{Booth} *	Tether length min
K_p roll rate (retraction)	Circle radius*	Retraction force
K_p pitch rate (traction)	Elevation angle	
K_p pitch rate (retraction)		
K_p , K_i , K_d roll (traction)		
K_p , K_i , K_d roll (retraction)		
K_p , K_i pitch (traction)		
K_p , K_i pitch (retraction)		

* Whether the parameter is needed depends on the chosen flight path.

The soft constraints are defined as follows:

- maximum cross-track error (traction phase) ≤ 75 m
- maximum kite velocity (traction phase) ≤ 100 m s⁻¹
- maximum tether force (excluding transitions) $\leq 1\,664\,910$ N
- maximum peak-to-average-power ratio (traction phase) ≤ 2.5 .

As PAPR turned out to be a difficult constraint to satisfy, optimisations are also run without the PAPR constraint to evaluate the influence of this constraint on the overall optimisation process and conclusions.

2.4 Additional system parameters

As this work focuses on comparing flight paths, it is of key importance to keep as many factors as possible the same throughout the analysis. Therefore, for each case study, the same kite, ground station, and tether are used. Similarly, one wind shear model without turbulence is chosen and kept constant throughout the analysis. The wind shear is determined from the averaged Ijmuiden measurement data presented in Schelbergen et al. (2020). The 3 MW MegAWES kite is chosen as the ground-generation fixed-wing kite (Eijkelhof and Schmehl, 2022). Since the original MegAWES simulation framework used a winch which was not sized for megawatt-scale power, a better-sized winch is used instead (Hummel et al., 2024b).

The kite is modelled as a rigid body (6 degrees of freedom). The total sum of forces acting on the kite is composed of the aerodynamic force, the tether force, and the gravitational force. The aerodynamic force is a resultant force combining lift, drag, and side force.

The tether is modelled similarly to the original MegAWES simulation environment. The tether is modelled as quasi-steady, discretised by lumped masses which are connected by straight elements and have a stiffness based on the Dyneema[®] material.

Table 2. General planform parameters of the wing, tail, and fuselage and characteristics of the winch and tether, adapted from Eijkelhof and Schmehl (2022).

Parameter	Value	Unit
Inertia J_{xx}	5.7680×10^5	kg m^2
Inertia J_{yy}	0.8107×10^5	kg m^2
Inertia J_{zz}	6.5002×10^5	kg m^2
Mass	6.8852×10^3	kg
Wing/RevE _H C airfoil		
Span	42.47	m
Chord _{root}	4.46	m
Chord _{tip}	2.11	m
LE sweep	2	°
Aspect ratio	12.0	–
Surface area	150.45	m^2
Twist _{root→tip}	5 → 0	°
Horizontal tail / NACA0012 airfoil		
Span	7.6	m
Chord	2.8	m
Vertical tail / NACA0012 airfoil		
Span	3	m
Chord	2.8	m
Winch		
Total inertia	1.0×10^4	m
Radius	1.5	m
Tether/Dyneema [®]		
Density	0.97×10^3	kg m^{-3}
Modulus	1.16×10^{11}	Pa
Ultimate strength	3.6×10^9	Pa
Cylindrical C_D	1.2	–

Table 2 shows a summary of the specifications of the entire system. This includes kite design, winch size, and tether specifications.

3 Results

This section presents a detailed comparison of the circular and figure-of-eight flight patterns, evaluating their performance based on key metrics such as cycle power and peak-to-average-power ratio. The simulations are performed in a consistent operational environment at a wind speed of 12 and 15 ms^{-1} across all tests. These wind speeds are below the rated wind speed. Above the rated wind speed, a real-world kite system is limited to not exceeding the rated power. This limit is not included in this simulation framework. Thus, a below-rated wind speed was chosen as a good compromise between high power output and a non-limited condition. For

both wind speeds, the simulations are performed with and without the PAPR constraint. Often, the constraint cannot be satisfied, and the solutions without the constraint can provide some insight into whether the relative differences are not based on the constraint handling but rather a characteristic of the pattern. The simulations with and without an active PAPR constraint are started from the same initial guess.

This section is structured as follows. First, the optimal solution for a fair comparison is chosen. Second, the full pumping cycle is compared between patterns. Even though complete pumping cycles are compared, the focus has been on the traction phase, where the pattern makes a big difference; the retraction strategy is relatively similar in all cases and therefore discussed in less detail.

3.1 Pattern exploration

In this section, only two of the earlier-mentioned three performance metrics are considered. The objective function focused on average cycle power, making it of key importance in selecting the cycle for detailed discussion in the next section. As the peak-to-average-power ratio (PAPR) constraint was not satisfied, this is included as well. First, the cycle power is compared to the circle radius in Fig. 6. As the algorithm is free to choose any radius, many combinations of parameters are evaluated, hence multiple power results per radius. The optimal radius for this kite at the given wind speed, with the operation conditions explained throughout this paper, is around 220 m. This means that the diameter of the flight pattern is around 10.5 times the wingspan.

From theory, a smaller turning radius should allow more of the kite's path to pass through regions of higher wind speed (after accounting for cosine losses), which suggests that smaller circles could extract more power. In practice, however, the radius cannot be made arbitrarily small because other physical and operational limits start to dominate.

Several factors explain why optimisation tends to favour larger radii instead:

- *Roll angle effects.* Smaller radii require the kite to roll more steeply. As shown in Makani's study (Tucker, 2020), this reduces power production efficiency.
- *Velocity limits.* Tighter turns increase the tangential velocity needed to follow the path along a turn. This follows from the definition of the required centripetal force. When the kite is already operating around its maximum aerodynamic capability, the only option is larger velocities to follow the sharper turns. If this velocity exceeds the maximum allowed limit, the optimiser penalises the solution, effectively discouraging small radii.
- *Wake effects.* In reality, smaller radii would also amplify wake losses (i.e. stronger wind speed reduction upstream). Although this is not included in the present




Figure 6. Circular pattern performance at 15 m s^{-1} wind speed with different circle radii. Left is the average cycle power. Right is the peak-to-average-power ratio.

analysis, this effect would reinforce the preference for larger radii.

Another important metric that needs to be considered is the PAPR. The peak power of the traction phase is divided by the average traction phase power. The large concentration of data points at low PAPR shows that the penalisation of the objective is working as expected. While optimising for more cycle power, the power peaks are minimised. In the given conditions, the optimal PAPR is a little over 4. Even though some results show that it could obtain a slightly lower PAPR than 4, those solutions are deemed less optimal, with either lower average power or exceeding any of the other constraints. Accepting a lower cycle power would potentially lead to a lower PAPR. However, the same can be said for the other patterns, likely leading to similar differences.

A similar analysis is performed for both variations of the figure of eight. Figures 7 and 8 show the optimisation results, illustrating the geometry (width and height) of the figure of eight with respect to the cycle power on the left and the PAPR on the right.

Both down-loop and up-loop flight patterns converge to an optimal combination of width and height. The up-loop seems to converge to smaller-sized figure of eights compared to the down-loop; however, as shown in the next section, the converged down-loop is slightly smaller. To reduce the PAPR, the up-loop pattern width needs to shrink. This is the expected behaviour since the smaller-width figure-of-eight patterns in particular remain more in the effective wind window. This benefits the upward flight. The down-loop uses the inner part to counteract gravity and seems less affected by the total width.

The left image in Fig. 7 shows that the optimal solution of the down-loop is limited at a lower cycle power than the circle pattern. The PAPR, however, converges to the lower

value of below 4. Similar to the circular pattern, lower PAPR solutions exist but are less optimal.

The left image in Fig. 8 shows that the most optimal solution of the up-loop is similarly limited at a lower cycle power than the circle pattern. Unlike the down-loop, the up-loop PAPR, however, converges to the highest PAPR value of the three patterns.

3.2 Pumping cycle comparison

From the previous results, the best-performing path is selected for a more detailed comparison. Here, the flight patterns are compared based on the power oscillations (power quality), projected ground surface area, and cycle average power.

Figure 9 shows the 15 m s^{-1} wind speed circular three-dimensional flight path coloured by the mechanical power, which is shown with respect to time on the right. The projected ground area is shown with the dashed red box on the ground plane. The centreline of the cycle trajectory is slightly shifted on the y axis, aligning the upward flight more with the centre of the wind window. As the target path is centred at zero azimuth, this result is solely a consequence of precise gain tuning and changing tracking accuracy along the pattern. Tracking accuracy can change as different velocities occur throughout the pattern, making it easier and harder to track a given target (Diehl et al., 2004; Ilzhöfer et al., 2007; Gros et al., 2013).

It can be noted that at this wind speed, most of the mechanical power is positive; however, during each loop, the winch needs to retract slightly to insert energy into the system to get it to the top of the circle. This extra required energy is due to the gravitational losses, which the kite alone cannot compensate for. The cycle average is at 2.254 MW, which is much higher than the figure-of-eight flight patterns.




Figure 7. Figure-of-eight down-loop pattern at 15 m s^{-1} wind speed with different pattern sizes at minimum tether length. Left is the average cycle power. Right is the peak-to-average-power ratio. Note: for visibility, the colour bar on the right has different limits compared to the up-loop pattern.




Figure 8. Figure-of-eight up-loop pattern at 15 m s^{-1} wind speed with different pattern sizes at minimum tether length. Left is the average cycle power. Right is the peak-to-average-power ratio. Note: for visibility, the colour bar on the right has different limits compared to the down-loop pattern.

To illustrate the ground area that the system covers, the projected ground area is shown under the flight path. The converged circular trajectory under these conditions projected on the ground covers an area of 0.31 km^2 ($0.53 \text{ km} \times 0.59 \text{ km}$). This results in a surface area power density of 7.24 MW km^{-2} .

Figure 10 presents the velocity and lift force with respect to cycle time for the circle. As expected, the velocity and lift force more or less follow the same behaviour.

The maximum velocity is at the start of the cycle. The relatively bad transition from retraction to the traction phase causes this spike in velocity and is, therefore, not considered a limiting factor in the validity of the result for this paper. Transition phases are highly complex and not relevant to the scope of this work. A higher velocity is reached at the end of the downward flight, slowing down on the way back to the upper part of the circle.

An additional performance metric that can be derived from these results is the dominant load frequency associated with the aerodynamic forces. Among these, the lift force emerges

as the most significant. Specifically, when considering only the traction phase of operation, the lift force exhibits a cyclic load on the structure with a frequency of approximately 0.041 Hz . In comparison to the down-loop phase, this higher cycle frequency, as well as a higher amplitude, suggests an increased likelihood of fatigue-related damage over time, as the system is subjected to more frequent load cycles.

Figure 11 shows the 15 m s^{-1} wind speed figure-of-eight down-loop three-dimensional flight path coloured by the mechanical power, which is shown with respect to time on the right. The projected ground area is shown with the dashed red box on the ground plane.

It can immediately be noted that at this wind speed, almost all of the mechanical power is positive. Out of the three patterns, this delivers a smoother power output, requiring less power-smoothing equipment. At the centre of the figure of eight, the kite is right in the optimal part of the wind window. Using this part to fly up, the forces can use this beneficial wind to counteract gravity, reducing overall power oscillations.

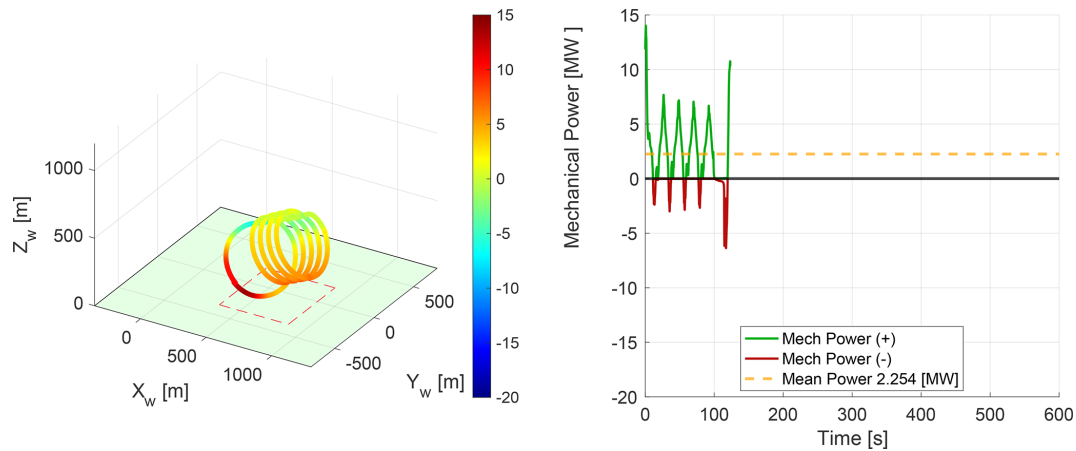


Figure 9. Circular flight path of the most optimised solution set at 15 m s^{-1} wind speed with active PAPR constraint. Left is the three-dimensional flight trajectory, which is coloured by mechanical power. Right is the mechanical power with respect to cycle time.




Figure 10. Circular flight path of the most optimised solution set at 15 m s^{-1} wind speed with active PAPR constraint. Left is the kite velocity with respect to cycle time. Right is the lift force with respect to cycle time.

tion. The cycle average is at 1.645 MW , which is lower than the circular path but higher than the up-loop figure of eight.

The converged down-loop trajectory under these conditions projected on the ground covers an area of 1.06 km^2 ($0.95 \text{ km} \times 1.10 \text{ km}$). This results in a surface area power density of 1.56 MW km^{-2} .

Figure 12 presents the velocity and lift force with respect to cycle time for the down-loop. Similarly to the circular path, the velocity and lift force follow the same behaviour.

The higher velocity occurs during downward flight at the outer curves, while the kite slows down in the centre, flying up.

The right of Fig. 12 shows that the dominant load frequency of the lift force is approximately 0.034 Hz . Out of the three studied patterns, this is the lowest, having favourable effects on kite lifetime.

Figure 13 shows the 15 m s^{-1} wind speed figure-of-eight up-loop three-dimensional flight path coloured by the mechanical power, which is shown with respect to time on the right. The projected ground area is shown with the dashed red box on the ground plane.

It can be noted that at this wind speed, the mechanical power is the most oscillatory of all patterns, requiring more power smoothing than the others. Opposite to the down-loop, the gravitational force needs to be counteracted at the outer edges of the figure of eight, where the kite is at the least optimal part of the wind window. Using this part to fly up, the winch has to assist using retraction, increasing overall power oscillation. The cycle average is at 0.877 MW , which is lower than the down-loop figure of eight and with worse power quality.

The converged up-loop trajectory under these conditions projected on the ground covers an area of 1.31 km^2

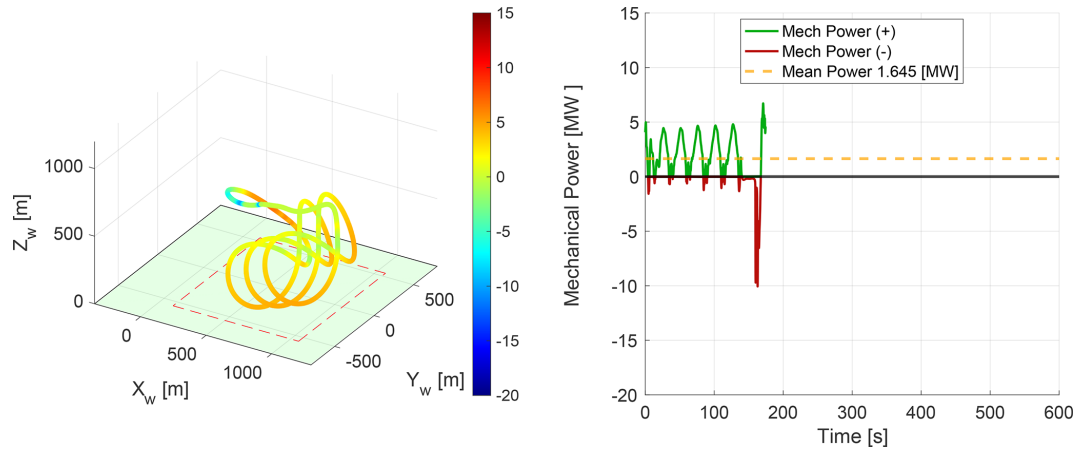


Figure 11. Figure-of-eight down-loop flight path of the most optimised solution set at 15 m s^{-1} wind speed with active PAPR constraint. Left is the three-dimensional flight trajectory, which is coloured by mechanical power. Right is the mechanical power with respect to cycle time.




Figure 12. Figure-of-eight down-loop flight path of the most optimised solution set at 15 m s^{-1} wind speed with active PAPR constraint. Left is the kite velocity with respect to the cycle time. Right is the lift force with respect to cycle time.

($1.16 \text{ km} \times 1.12 \text{ km}$). This results in a surface area power density of 0.67 MW km^{-2} .

Figure 14 presents the velocity and lift force with respect to cycle time for the up-loop. Similarly to the other paths, the velocity and lift force follow the same behaviour.

The maximum velocity is at the end of the cycle; looking at the three-dimensional flight path in Fig. 13, this might not be specific to this flight path. It seems to be a result of the sub-optimal transition phase, which causes instability in the flight direction and velocity. Unlike the down-loop, the higher velocity occurs in the centre, slowing down at the outer edge.

The right of Fig. 14 shows that the dominant load frequency of the lift force is approximately 0.045 Hz . Out of the three studied patterns, this is the highest, having unfavourable effects on kite lifetime.

An overview of the results discussed before is shown in Table 3.

Apart from the down-loop figure of eight, the optimiser cannot find solutions that satisfy the PAPR constraint. This raises the question whether the implemented constraint handling itself could have an influence on the most optimal solution and thus the pattern comparison. To validate the relative differences between patterns, the comparative study is also performed on the most optimal solutions when turning off the PAPR constraints. This allows for solutions where all constraints are satisfied.

For clarity only the overview of the results are shown in Table 4. Comparing this table to Table 3, it can be noted that the relative differences stay the same, even though the optimal power of the figure-of-eight paths differ more from their PAPR-limited counterpart. As expected, the power and PAPR of the circle and up-loop patterns are higher when not constraining the PAPR. The down-loop does show an unexpected result – the power and PAPR are actually lower. This




Figure 13. Figure-of-eight up-loop flight path of the most optimised solution set at 15 m s^{-1} wind speed with active PAPR constraint. Left is the three-dimensional flight trajectory, which is coloured by mechanical power. Right is the mechanical power with respect to cycle time.




Figure 14. Figure-of-eight up-loop flight path of the most optimised solution set at 15 m s^{-1} wind speed with active PAPR constraint. Left is the kite velocity with respect to cycle time. Right is the lift force with respect to cycle time.

illustrates the difficulty of optimisation of non-linear problems; it is very challenging to avoid local optima. A possible solution would be to start from several different initial conditions, but this lies beyond the scope of the present work. The down-loop figure of eight remains beneficial when looking at power quality. With respect to the ground surface area power density, the figure-of-eight patterns did improve, although the circle still strongly outperforms the down-loop and up-loop patterns.

Finally, the simulations for both with (Table 5) and without (Table 6) PAPR constraint are repeated at a lower wind speed of 12 m s^{-1} . This is done to evaluate whether the relative results are wind speed (operating position on the power curve) dependent. This wind speed was chosen such that the heavy MegAWES kite would still be able to fly in all patterns. Lower than 12 m s^{-1} , the kite has difficulties performing a full power cycle, especially for the figure-of-eight up-loop.

Using the PAPR constraint, the cycle power, PAPR, and surface power density show similar characteristics between patterns. The main difference seems to be for the PAPR, where the figure-of-eight up-loop outperforms the circle. The cause of this seems to be the pattern size. The minimum size of the circle is limited, connected to the kite size. To reduce the PAPR, mainly a consequence of gravity, the pattern should reduce in size to reduce the difference between the upper and lower part of the pattern. As the minimum is indirectly constrained, the PAPR cannot be reduced to a value lower than the up-loop. Looking at the unconstrained results, the power indeed increases for the circle and down-loop patterns. However, this also causes a drastic increase in PAPR. At lower wind speeds, gravity has a much bigger influence, and thus the difference becomes much larger. The up-loop, however, became stuck in a local optimum, and the PAPR constraint actually helped to steer away from this. Of

Table 3. Summary of pattern performance of the three studied flight patterns using the MegAWES kite at 15 m s^{-1} wind speed with active PAPR constraint.

	Circle	Figure-of-eight down-loop	Figure-of-eight up-loop
Cycle power	2.25 MW	1.64 MW	0.88 MW
PAPR	5.07	2.37	7.27
Ground area	0.31 km^2	1.06 km^2	1.31 km^2
Surface area power density	7.24 MW km^{-2}	1.56 MW km^{-2}	0.67 MW km^{-2}
Cycle duration	123.3 s	174.5 s	397.6 s
– Traction	98.2 s (79.7 %)	137.9 s (79.0 %)	353.2 s (88.8 %)
– Retraction and transition	25.1 s (20.3 %)	36.6 s (21.0 %)	44.4 s (11.2 %)
Maximum kite velocity (traction)	95.55 m s^{-1}	80.05 m s^{-1}	91.30 m s^{-1}
Average kite velocity (traction)	66.98 m s^{-1}	63.27 m s^{-1}	68.88 m s^{-1}
Dominant load frequency (lift)	0.041 Hz	0.034 Hz	0.045 Hz

Table 4. Summary of pattern performance of the three studied flight patterns using the MegAWES kite at 15 m s^{-1} wind speed without PAPR constraint.

	Circle	Figure-of-eight down-loop	Figure-of-eight up-loop
Cycle power	2.27 MW	1.55 MW	1.50 MW
PAPR	5.69	2.29	9.07
Ground area	0.31 km^2	0.91 km^2	1.14 km^2
Surface area power density	7.24 MW km^{-2}	1.71 MW km^{-2}	1.32 MW km^{-2}
Cycle duration	121.1 s	122.4 s	387.4 s
– Traction	96.0 s (79.3 %)	89.5 s (73.1 %)	364.3 s (94.0 %)
– Retraction and transition	25.1 s (20.7 %)	32.9 s (26.9 %)	23.1 s (6.0 %)
Maximum kite velocity (traction)	99.79 m s^{-1}	80.19 m s^{-1}	95.63 m s^{-1}
Average kite velocity (traction)	66.80 m s^{-1}	63.95 m s^{-1}	73.41 m s^{-1}
Dominant load frequency (lift)	0.041 Hz	0.041 Hz	0.059 Hz

course, this is random behaviour and cannot be considered a common phenomenon. For the sake of this comparison, the simulations were started from the same initial conditions. To help kick-start the no-PAPR constraint problem, one could have used the solution of the constrained problem as a starting condition instead. Nonetheless, finding the global optimum is not the purpose of this study.

4 Conclusions

This study systematically evaluated the performance of circular and figure-of-eight flight patterns for fixed-wing AWE systems under consistent operational conditions. Through the development of a MATLAB[®] Simulink[®]-based controller, the study achieved effective path tracking during the traction phase and demonstrated clear distinctions between these flight patterns in terms of energy output, power quality, spatial efficiency, and structural load effects.

Unlike the results from the analysis on a small, lightweight kite, the heavy kite studied here had different power production levels over different patterns. The circular flight pattern yielded the highest cycle-averaged power output, outperforming both figure-of-eight patterns in energy capture efficiency. Additionally, it required a smaller operational area, offering a beneficial power density, which is substantially higher than the down-loop and up-loop figure-of-eight paths. These advantages make the circular path attractive for applications prioritising maximal energy capture within limited ground areas. It must be noted that there is some room to find solutions which reduce the average power output and improve the power quality.

Conversely, the figure-of-eight down-loop configuration demonstrated superior power quality, producing lower power oscillations with a peak-to-average-power ratio that could actually get below the 2.5 constraint. This pattern also mitigated gravitational losses more effectively by leveraging optimal wind window positions during ascent. The up-loop variant, exhibiting the lowest cycle power and highest peak-

Table 5. Summary of pattern performance of the three studied flight patterns using the MegAWES kite at 12 m s^{-1} wind speed with active PAPR constraint.

	Circle	Figure-of-eight down-loop	Figure-of-eight up-loop
Cycle power	0.91 MW	0.77 MW	0.49 MW
PAPR	9.77	3.99	5.26
Ground area	0.36 km^2	0.97 km^2	0.77 km^2
Surface area power density	2.50 MW km^{-2}	0.80 MW km^{-2}	0.63 MW km^{-2}
Cycle duration	291.0 s	339.7 s	351.2 s
– Traction	271.1 s (93.1 %)	300.8 s (88.5 %)	330.6 s (94.1 %)
– Retraction and transition	19.9 s (6.9 %)	38.9 s (11.5 %)	20.5 s (5.9 %)
Maximum kite velocity (traction)	99.64 m s^{-1}	76.50 m s^{-1}	84.08 m s^{-1}
Average kite velocity (traction)	62.67 m s^{-1}	61.10 m s^{-1}	61.56 m s^{-1}
Dominant load frequency (lift)	0.045 Hz	0.044 Hz	0.037 Hz

Table 6. Summary of pattern performance of the three studied flight patterns using the MegAWES kite at 12 m s^{-1} wind speed without PAPR constraint.

	Circle	Figure-of-eight down-loop	Figure-of-eight up-loop
Cycle power	0.99 MW	0.82 MW	0.28 MW
PAPR	11.01	12.42	10.20
Ground area	0.28 km^2	0.87 km^2	0.82 km^2
Surface area power density	3.49 MW km^{-2}	0.94 MW km^{-2}	0.34 MW km^{-2}
Cycle duration	273.1 s	372.6 s	714.5 s
– Traction	254.4 s (93.2 %)	339.6 s (91.1 %)	698.5 s (97.8 %)
– Retraction and transition	18.7 s (6.8 %)	33.0 s (8.9 %)	16.0 s (2.2 %)
Maximum kite velocity (traction)	107.88 m s^{-1}	91.71 m s^{-1}	92.12 m s^{-1}
Average kite velocity (traction)	63.21 m s^{-1}	61.35 m s^{-1}	60.85 m s^{-1}
Dominant load frequency (lift)	0.048 Hz	0.046 Hz	0.042 Hz

to-average-power ratio, suggests it is the least optimal flight path.

Additionally, each pattern displayed different factors which could impact structural fatigue and kite lifetime, driven by differences in load amplitude and frequency, as well as velocity profiles. The figure-of-eight down-loop configuration generally exhibited a lower load frequency, beneficial for mitigating structural fatigue and potentially extending kite life. However, at the lower wind speed, the up-loop showed some potential to reduce the frequency below the down-loop. Furthermore, the amplitude of the force is generally lower for the down-loop compared to the other two patterns. Nevertheless, fatigue should be studied separately, in a more detailed study, to ensure a well-defined relation between pattern and cyclic loads.

In conclusion, the choice of flight pattern for AWE systems involves a trade-off between power density, power quality, and operational stability. Circular patterns are preferable for maximising energy output and spatial efficiency, while the down-loop figure of eight offers a promising solution for

applications where system longevity and smooth power output are critical for grid integration.

Future research should expand the optimisation domain, including adjustments to azimuth angles and more refined phase transitions, to further optimise power extraction and load management. Also, more than two wind speeds should be considered to evaluate if conclusions change throughout the complete wind speed domain.

Appendix A: Constraint handling

As this paper is trying to find a comparable flight path rather than an innovative and fast optimisation strategy, a simple constraint handling process is used to implement constraints in the CMA-ES optimisation. The power objective is only penalised when maxima are exceeded. The objective is calculated by

$$C_{\text{obj}} = -\text{mean}(P_{\text{mech}}) + p_{\text{cte}} + p_{V_k} + p_{F_t} + p_{\text{papr}}, \quad (\text{A1})$$

with the penalties

$$p_{\text{cte}} = \begin{cases} 10^2 \cdot \left(\frac{\max(\text{cross-track error})}{75\text{m}} - 1 \right), & \text{if } \max(\text{cross-track error}) > 75 \text{ m} \\ 0, & \text{otherwise} \end{cases}$$

$$p_{V_k} = \begin{cases} 10^4 \cdot \left(\frac{\max(V_k)}{100 \text{ m s}^{-1}} - 1 \right), & \text{if } \max(V_k) > 100 \text{ m s}^{-1} \\ 0, & \text{otherwise} \end{cases}$$

$$p_{F_t} = \begin{cases} 10^4 \cdot \left(\frac{\max(F_t)}{1664910 \text{ N}} - 1 \right), & \text{if } \max(F_t) > 1664910 \text{ N} \\ 0, & \text{otherwise} \end{cases}$$

$$p_{\text{papr}} = \begin{cases} 10^5 \cdot \left(\frac{\max(\text{PAPR})}{2.5} - 1 \right), & \text{if } \max(\text{PAPR}) > 2.5 \\ 0, & \text{otherwise} \end{cases}$$

The weight factors were chosen based on importance. For example, the cross-track error, the error on position with respect to the planned trajectory, is considered less important than velocity, tether force, and peak-to-average-power ratio.

Code and data availability. The reference design of the kite and the simulation framework is available on <https://github.com/awegroup/MegAWES> (last access: 20 February 2026) in open-access format (release version 3.0). The specific environment and parameter inputs used in this work are available on the 4TU database (<https://doi.org/10.4121/70679dc6-30ba-4f75-a6ba-94144f456da7>, Eijkelhof, 2025).

Author contributions. Conceptualisation, DE, NR, and RS; methodology, DE, NR, and RS; software, DE and NR; writing – original draft preparation, DE; writing – review and editing, DE; supervision, RS; funding acquisition, RS. All authors have read and agreed to the published version of the paper.

Competing interests. At least one of the (co-)authors is a member of the editorial board of *Wind Energy Science*. The peer-review process was guided by an independent editor, and the authors also have no other competing interests to declare.

Disclaimer. Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims made in the text, pub-

lished maps, institutional affiliations, or any other geographical representation in this paper. The authors bear the ultimate responsibility for providing appropriate place names. Views expressed in the text are those of the authors and do not necessarily reflect the views of the publisher.

Financial support. This research has been supported by the Nederlandse Organisatie voor Wetenschappelijk Onderzoek (grant no. 17628).

Review statement. This paper was edited by Alessandro Croce and reviewed by four anonymous referees.

References

- Diehl, M., Magni, L., and De Nicolao, G.: Efficient NMPC of unstable periodic systems using approximate infinite horizon closed loop costing, *Annu. Rev. Control*, 28, 37–45, <https://doi.org/10.1016/j.arcontrol.2004.01.011>, 2004.
- Eijkelhof, D.: Design and Optimisation Framework of a Multi-MW Airborne Wind Energy Reference System, Master's thesis, Delft University of Technology & Technical University of Denmark, <http://resolver.tudelft.nl/uuid:e759f9ad-ab67-43b3-97e0-75558ecf222d> (last access: 20 February 2026), 2019.
- Eijkelhof, D.: MegAWES (Version 3), Six-Degrees-of-Freedom Simulation Model for Future Multi-Megawatt Airborne Wind Energy Systems, Delft University of Technology, <https://doi.org/10.4121/70679dc6-30ba-4f75-a6ba-94144f456da7>, 2025.
- Eijkelhof, D. and Schmehl, R.: Six-degrees-of-freedom simulation model for future multi-megawatt airborne wind energy systems, *Renew. Energ.*, 196, 137–150, <https://doi.org/10.1016/j.renene.2022.06.094>, 2022.
- Erhard, M. and Strauch, H.: Theory and Experimental Validation of a Simple Comprehensible Model of Tethered Kite Dynamics Used for Controller Design, in: *Airborne Wind Energy*, edited by: Ahrens, U., Diehl, M., and Schmehl, R., Green Energy and Technology, chap. 8, Springer, Berlin Heidelberg, 141–165, https://doi.org/10.1007/978-3-642-39965-7_8, 2013.
- Fagiano, L., Zraggen, A. U., Morari, M., and Khammash, M.: Automatic Crosswind Flight of Tethered Wings for Airborne Wind Energy: Modeling, Control Design, and Experimental Results, *IEEE T. Contr. Syst. T.*, 22, 1433–1447, <https://doi.org/10.1109/TCST.2013.2279592>, 2014.
- Fechner, U. and Schmehl, R.: Flight path control of kite power systems in a turbulent wind environment, in: 2016 American Control Conference (ACC), 4083–4088, <https://doi.org/10.1109/ACC.2016.7525563>, 2016.
- Fernandes, M. C. R. M., Vinha, S., Paiva, L. T., and Fontes, F. A. C. C.: L0 and L1 Guidance and Path-Following Control for Airborne Wind Energy Systems, *Energies*, 15, <https://doi.org/10.3390/en15041390>, 2022.
- Gray, M.: *Modern Differential Geometry of Curves and Surfaces with Mathematica*, Second Edition, Textbooks in Mathematics, Taylor & Francis, ISBN 9780849371646, 1997.

- Gros, S., Zanon, M., and Diehl, M.: Control of Airborne Wind Energy systems based on Nonlinear Model Predictive Control & Moving Horizon Estimation, in: 2013 European Control Conference (ECC), 1017–1022, <https://doi.org/10.23919/ECC.2013.6669713>, 2013.
- Hansen, N.: The CMA Evolution Strategy: A Comparing Review, in: *Towards a New Evolutionary Computation: Advances in the Estimation of Distribution Algorithms*, edited by: Lozano, J. A., Larrañaga, P., Inza, I., and Bengoetxea, E., Springer Berlin Heidelberg, Berlin, Heidelberg, 75–102, https://doi.org/10.1007/3-540-32494-1_4, 2006.
- Houska, B. and Diehl, M.: Optimal Control of Towing Kites, in: *Proceedings of the 45th IEEE Conference on Decision and Control*, 2693–2697, <https://doi.org/10.1109/CDC.2006.377210>, 2006.
- Hummel, J., Pollack, T., Eijkelhof, D., Van Kampen, E., and Schmehl, R.: Power smoothing by kite tether force control for megawatt-scale airborne wind energy systems, *J. Phys. Conf. Ser.*, 2767, 072019, <https://doi.org/10.1088/1742-6596/2767/7/072019>, 2024a.
- Hummel, J., Pollack, T. S., Eijkelhof, D., van Kampen, E.-J., and Schmehl, R.: Winch Sizing for Ground-Generation Airborne Wind Energy Systems, in: 2024 European Control Conference (ECC), 675–680, <https://doi.org/10.5281/zenodo.13384380>, 2024b.
- Ilzhöfer, A., Houska, B., and Diehl, M.: Nonlinear MPC of kites under varying wind conditions for a new class of large-scale wind power generators, *Int. J. Robust Nonlin.*, 17, 1590–1599, <https://doi.org/10.1002/rnc.1210>, 2007.
- Jehle, C. and Schmehl, R.: Applied Tracking Control for Kite Power Systems, *J. Guid. Control Dynam.*, 37, <https://doi.org/10.2514/1.62380>, 2014.
- Li, H., Olinger, D. J., and Demetriou, M. A.: Control of an airborne wind energy system using an aircraft dynamics model, in: 2015 American Control Conference (ACC), 2389–2394, <https://doi.org/10.1109/ACC.2015.7171090>, 2015.
- Licitra, G., Koenemann, J., Bürger, A., Williams, P., Ruitkamp, R., and Diehl, M.: Performance assessment of a rigid wing Airborne Wind Energy pumping system, *Energy*, 173, 569–585, <https://doi.org/10.1016/j.energy.2019.02.064>, 2019.
- Loyd, M. L.: Crosswind kite power (for large-scale wind power production), *J. Energy*, 4, 106–111, <https://doi.org/10.2514/3.48021>, 1980.
- Park, S., Deyst, J., and How, J.: A new nonlinear guidance logic for trajectory tracking, in: *AIAA Guidance, Navigation, and Control Conference and Exhibit*, <https://doi.org/10.2514/6.2004-4900>, 2004.
- Rapp, S., Schmehl, R., Oland, E., and Haas, T.: Cascaded Pumping Cycle Control for Rigid Wing Airborne Wind Energy Systems, *J. Guid. Control Dynam.*, 42, 1–18, <https://doi.org/10.2514/1.G004246>, 2019.
- Rossi, N.: Performance comparison and flight controller of circular and figure-of-eight paths for fixed-wing airborne wind energy systems, Master's thesis, University of Trento, Zenodo [thesis], <https://doi.org/10.5281/zenodo.10160420>, 2023.
- Ruitkamp, R. and Sieberling, S.: Description and Preliminary Test Results of a Six Degrees of Freedom Rigid Wing Pumping System, in: *Airborne Wind Energy*, edited by Ahrens, U., Diehl, M., and Schmehl, R., Green Energy and Technology, chap. 26, Springer, Berlin Heidelberg, 443–458, https://doi.org/10.1007/978-3-642-39965-7_26, 2013.
- Schelbergen, M., Kalverla, P. C., Schmehl, R., and Watson, S. J.: Clustering wind profile shapes to estimate airborne wind energy production, *Wind Energ. Sci.*, 5, 1097–1120, <https://doi.org/10.5194/wes-5-1097-2020>, 2020.
- Schmehl, R., Noom, M., and van der Vlugt, R.: Traction Power Generation with Tethered Wings, in: *Airborne Wind Energy*, edited by Ahrens, U., Diehl, M., and Schmehl, R., Green Energy and Technology, chap. 2, Springer, Berlin Heidelberg, 23–45, ISBN 978-3-642-39964-0, https://doi.org/10.1007/978-3-642-39965-7_2, 2013.
- Tucker, N.: Airborne Wind Turbine Performance, in: *The Energy Kite: Selected Results From the Design, Development and Testing of Makani's Airborne Wind Turbines, Part I*, Makani Technologies, Alameda, CA, USA, 101–224, <https://x.company/projects/makani/> (last access: 20 February 2026), 2020.
- van der Vlugt, R., Peschel, J., and Schmehl, R.: Design and Experimental Characterization of a Pumping Kite Power System, in: *Airborne Wind Energy*, edited by Ahrens, U., Diehl, M., and Schmehl, R., Green Energy and Technology, chap. 23, Springer, Berlin Heidelberg, 403–425, ISBN 978-3-642-39964-0, https://doi.org/10.1007/978-3-642-39965-7_23, 2013.
- Vermillion, C., Cobb, M., Fagiano, L., Leuthold, R., Diehl, M., Smith, R. S., Wood, T. A., Rapp, S., Schmehl, R., Olinger, D., and Demetriou, M.: Electricity in the air: Insights from two decades of advanced control research and experimental flight testing of airborne wind energy systems, *Annu. Rev. Control*, 52, 330–357, <https://doi.org/10.1016/j.arcontrol.2021.03.002>, 2021.