https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

A rigidly-convected, lifting-line vortex model in the optimal control of airborne wind energy systems

Rachel Leuthold[1], Jochem De Schutter[2], Curran Crawford[3], Sébastien Gros[4], and Moritz Diehl[1]

[1]Systems Control and Optimization Laboratory, Department of Microsystems Engineering (IMTEK), Albert-Ludwigs-University, Freiburg, Germany
[2]TransnetBW, Wendlingen am Neckar, Baden-Württemberg, Germany
[3]Institute for Integrated Energy Systems, University of Victoria, Victoria, Canada
[4]Department of Engineering Cybernetics, Norwegian University of Science and Technology, Gløshaugen, Norway

Correspondence: Rachel Leuthold (rcleuthold@gmail.com)

Abstract. Solving optimal control problems (OCPs) to estimate the performance of airborne wind energy (AWE) systems can raise the question: "How should one include a vortex-filament model into an OCP on an irregular flight path, specifically where there are also a substantial number of inequality constraints?" Since a rigidly-convected, lifting-line vortex (RLL) model might balance computational expense against sensitivity to geometric assumptions, we construct a solvable problem by judiciously lifting certain algebraic variables into an AWE OCP's decision variables. This RLL model makes predictions that are consistent considering modeling assumptions with independently-published Large Eddy Simulations, during rotationally-steady as well as dynamic situations. Further, demonstrating the importance of including wake-modeling in airborne wind energy applications, solving an AWE OCP including this RLL model finds a substantial difference in power compared to solving an AWE OCP assuming negligible induced velocity; this paper's demonstration problem, with a two-kite (Ampyx-AP2) system and a single-winding trajectory, predicts a 20%-30% loss in harvestable power (depending on initialization and objective). Unfortunately, solving such an RLL AWE OCP is also quite expensive: in the problem instances and solver configurations tested within this work, the required computational memory scales with the total number of vortex elements instantiated in the entire problem, while the processor time scales with more than that number squared. In conclusion, it is possible to include a RLL model into a solvable AWE OCP with an irregular trajectory and many inequality constraints, though the solution is quite expensive and is heavily influenced by the homotopy strategy used for initialization.

1 Introduction

With Hansen et al. (2026) confirming that 2026 is on-track to be yet another 'hottest year' on record, we're again reminded of the importance of transitioning to a diverse portfolio of clean, renewable energy sources. Among these options is airborne wind energy (AWE). In lift-mode AWE energy systems, tethered drones - also called 'kites' - fly variable-lift trajectories, so that a tether wound around a generator drum can be unrolled under high tension and re-rolled again under low tension. In the particular system variant considered here (see Figure 1), multiple kites are held by secondary tethers that branch off of a main

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

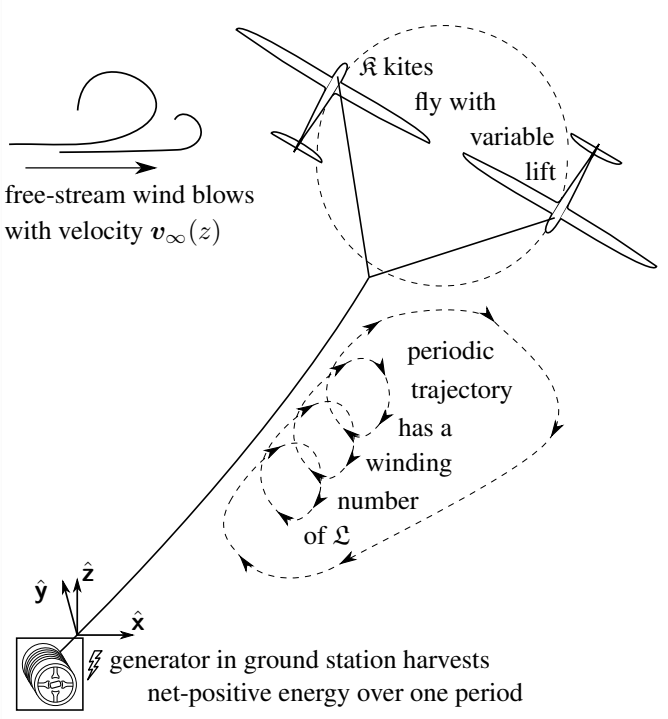


Figure 1. A sketch of an airborne wind energy system. In this sketch, the number of kites $\mathfrak{K} = 2$, and the winding number of the trajectory is $\mathfrak{L} = 4$. The example problem used in Sections 4 and 5 similarly uses two $\mathfrak{K} = 2$ kites, but only one winding $\mathfrak{L} = 1$.

tether, in a manner that roughly balances the forces perpendicular to the tether so that the losses due to the drag on the main tether can be kept small (Payne and McCutchen, 1976; Houska and Diehl, 2007; Zanon et al., 2013).

As with all energy sources, the contribution of AWE to an energy portfolio depends on the ability of system designers and purchasers to make accurate predictions of the performance of any particular system at any particular site. As reported in Vermillion et al. (2021), one of the main methods of estimating this performance uses the offline solution of optimal control problems (OCPs), wherein parameters, states, and controls are found for the system by maximizing - minimizing the negative of - a relevant figure-of-merit, while satisfying some long list of constraints, including a model of the system's dynamical behavior and additional constraints that ensure safety in the system's operation. This approach was introduced to AWE by Houska and Diehl (2007).

To date, there exists at least one open-source, user-friendly, and modular toolbox that will formulate and solve these AWE OCPs, the `awebox` (De Schutter et al., 2023). However, the solution of an AWE OCP can only be as good as the models included in its constraints. The gap being considered in this paper specifically concerns wake modeling. Betz optimality suggests that representing the conservation of momentum within the flow among the physical-model constrained by an AWE OCP will impact that OCP's solution, especially in AWE systems with multiple kites. Indeed, previous work - eg., Zanon et al. (2014)'s two-kite OCP and Leuthold et al. (2017)'s three-kite rotationally-steady optimization - has shown that there can be a large

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

difference in the predicted system power output using a low-fidelity (actuator-based) wake model compared to no wake model at all. This, of course, naturally leads to the question: what happens to this difference if we use a higher-fidelity wake model?

There has been a decent amount of work in the last years including wake models into OCPs, such as including wake steering in wind turbine farms and adding realistic induction effects to formation flight. Much of this work uses either actuator-surface models (see, for example, Zanon et al. (2014), Leuthold et al. (2018), Trevisi et al. (2022), and Sterle et al. (2024)) or high-fidelity CFD models such as Large Eddy Simulations (see, for example, Goit and Meyers (2015), Goit et al. (2016), Munters and Meyers (2017), Munters and Meyers (2018), Vali et al. (2018), and Janssens and Meyers (2024)) - not including the cases following partitioned approaches where the determination by optimal control of a trajectory or control input has been performed separately from the high-fidelity flow simulation (see, for example, Haas et al. (2019), Haas et al. (2022), Crismer et al. (2024), and Crismer et al. (2026)).

The difficulty with using high-fidelity computational fluid dynamics (CFD) models in AWE OCPs generally has to do with computational cost. As an example of this high cost: Munters (2017)'s optimal control problem maximizing the power harvested by 16 wind turbines - each controlled by the rate-of-change of its thrust coefficient as well as its yaw rate - over an optimization period of 300 s is reported as requiring "35,000 corehours per optimization case" once a Large Eddy Simulation (LES) wake model is included among the constraints. Computational costs on the order of multiple-thousands of core-hours would be appropriate during the later phases of a detailed design, but not during the rough, initial sorting through conceptual designs.

Unfortunately, the actuator-surface models on the other side of the literature have historically been developed for rotors, with the downside that they are built on assumptions of regular actuator geometries. Even the most general sorts of actuator-based models, such as Paraschivoiu (1988)'s double-multiple streamtube model which was developed for vertical-axis wind turbine geometries, require some underlying geometrical definitions. In an AWE context, where OCP-solutions generate trajectories (from the butterfly-shaped trajectories of Houska and Diehl (2007), to the stacked figure-eights of Licitra et al. (2019), to the expanding helices of Haas et al. (2022)) that are generally non-circular, non-planar, and non-symmetrical, it is not immediately obvious how one should define the sorts of geometric-parameters that these actuator-surface models require.

So, if the high- and low-fidelity ends of the modeling spectrum are both unsatisfactory, this leaves mid-fidelity modeling: for example, vorticity-based modeling. Within the field of vorticity-based methods, point-vortices require too many elements, while three-dimensional, structural vortices like cylinders appear to be vulnerable to the same difficulty with geometric assumptions that limit actuator-surface models. Narrowing our view to two-dimensional vortex objects, vortex-ring models - as used in van den Broek et al. (2022) and van den Broek et al. (2023)'s wake-steering OCPs or proposed for simulation by Trevisi et al. (2023) - and engineering models derived from vortex-ring models like Leuthold et al. (2019), Gaunaa et al. (2020), and Trevisi et al. (2025) again suffer from the aforementioned problematic geometric over-specificity. This suggests trying to represent wake effects within an AWE OCP, using a vortex filament model.

Encouragingly, there has been research into modeling the wakes of systems similar to the multi-kite AWE system using vortex filament models: see, eg., Rancourt and Mavris (2018)'s simulation of two tethered-drones in vertical take-off and landing (VTOL), which explicitly concerns itself with reducing the necessary computational memory by consolidating the

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

trailing vortices of a lifting-line model. However, the OCPs presented in the associated thesis (Rancourt, 2016, secs. 5.2, 5.4) turns out neither to use Rancourt and Mavris (2018)'s wake model (instead, relying on a pre-computed rigid, straight, steady lifting-line wake for their OCPs, then performing higher-fidelity free-wake lifting-line simulations on trajectories that have been perturbed from the OCP solution) nor to include the irregular trajectory used in Rancourt and Mavris (2018) (instead, reducing the trajectory to a phase-shifted ellipse). That particular presentation suggests that the authors of this work are not the only ones who have found it difficult[1] to include classical vortex filament simulation techniques into OCPs with irregular trajectories and inequality constraints.

To the authors' knowledge, there have been few successful solutions of OCPs that have included vortex models. The most relevant example may be the work of van den Broek et al. (2024), wherein free-wake vortex filaments - shed in a ring-layout from the blades' tips - are used to represent the wake in an inequality-constrained OCP that maximizes the mean power of a farm of wind turbines. Other examples in the literature are DeVries and Paley (2013)'s and Zheng and Lyu (2023)'s inclusions of quasi-steady lifting-line effects into tracking problems related to formation flight; Lee et al. (2022)'s use of an unsteady vortex-lattice method to find an optimal trajectory over the 4-second transition of an insect-like, flapping-wing micro air vehicle from hovering to forwards flight; van den Broek et al. (2022)'s and van den Broek et al. (2023)'s use of unsteady (free-wake) vortex rings to describe the vorticity in the flow behind two two-dimensional turbines, one downstream of the other, in economic nonlinear model predictive control problems; Tagg et al. (2024)'s use of a vortex-particle method in their energy-minimizing trajectory-optimization of an electric VTOL aircraft's transition maneuver; and De Schutter et al. (2026)'s use of a vortex-patch near-wake and trailed-dipole far-wake on a half-periodic AWE trajectory optimization. We might make the brief observation that while DeVries and Paley (2013), Zheng and Lyu (2023), Tagg et al. (2024) and De Schutter et al. (2026) include inequality constraints into their OCPs, neither van den Broek et al. (2022) nor van den Broek et al. (2023) do.

The above summary leaves a gap in the state-of-the-art, not just as relates to the question, 'How should one include a vortex-filament model into an AWE OCP?' but also in relation to the broader question of 'How should one include a vortex-filament model into an OCP on an irregular flight path, specifically where there are also a substantial number of inequality constraints?' This gap is what this paper will attempt to address.

2 Model and optimal control problem

2.1 Optimal control

The basic idea behind optimal control is that one doesn't determine the desired controls alone, but rather the entirety of a system's planned trajectory - the history in time $t \in [0, t_{\text{final}}]$ of the states $\boldsymbol{x}(t)$, the algebraic variables $\boldsymbol{y}(t)$, and the controls $\boldsymbol{u}(t)$, as well as any relevant finite-dimensional system parameters $\boldsymbol{\theta}$ - by setting the contents of this trajectory as the decision variables $\boldsymbol{w} = [\boldsymbol{x}(\cdot)^\top, \boldsymbol{y}(\cdot)^\top, \boldsymbol{u}(\cdot)^\top, \boldsymbol{\theta}^\top]^\top$ of an optimization problem that maximizes or minimizes some figure-of-merit, called the objective function ψ. It falls to the enforced equality constraints $\boldsymbol{g}$ and inequality constraints $\boldsymbol{h}$ to ensure that only

[1]The reader is kindly referred to the footnotes of this paper for brief summaries of some of the authors' previous unsuccessful approaches at solving the problem of including a vortex-filament model into an AWE OCP.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

those trajectories that are physically-possible and acceptable-to-the-user will actually be selected.

$$\begin{aligned} \boldsymbol{w}^\star = \quad & \underset{\boldsymbol{w}}{\arg.\min.} \quad && \psi\left(\boldsymbol{x}(\cdot),\boldsymbol{y}(\cdot),\boldsymbol{u}(\cdot),\boldsymbol{\theta}\right) && && (1) \\ & \text{subj. to} && \boldsymbol{g}\left(\dot{\boldsymbol{x}}(t),\boldsymbol{x}(t),\boldsymbol{y}(t),\boldsymbol{u}(t),\boldsymbol{\theta}\right) && = \boldsymbol{0} && \forall t \in [0,t_{\text{final}}] \\ & && \boldsymbol{h}\left(\boldsymbol{x}(t),\boldsymbol{y}(t),\boldsymbol{u}(t),\boldsymbol{\theta}\right) && \leq \boldsymbol{0} && \forall t \in [0,t_{\text{final}}]. \end{aligned}$$

These equality constraints always include the requirement that the selected trajectory must satisfy some physics-derived model of the system dynamics $\boldsymbol{d}$, frequently given as an Index-1 differential-algebraic-equation (DAE) that has the right number of independent entries to uniquely determine all of the states' derivatives $\dot{\boldsymbol{x}}$ and all of the algebraic variables $\boldsymbol{y}$, at any modeled instance. However, since numerical solvers don't typically accept functionals as arguments, so-called direct optimal control methods will discretize this continuous-time problem prior to solution. This means that we divide the optimization period $[0,t_{\text{final}}]$ into $\mathfrak{I}$ control intervals, each between $[t_{\mathfrak{i}},t_{\mathfrak{i}+1}]$, so that $0=t_0<t_1<...<t_{\mathfrak{I}-1}<t_{\mathfrak{I}}=t_{\text{final}}$. In the case of discretization by direct collocation - as we will use in this paper - we further augment this time grid by introducing the intermediate times of the integration-stages (the 'collocation nodes') $t_{\mathfrak{i},\mathfrak{d}}$ within each control interval, as $t_{\mathfrak{i}}<t_{\mathfrak{i},0}<t_{\mathfrak{i},1}<...<t_{\mathfrak{i},\mathfrak{D}-1}=t_{\mathfrak{i}+1}$. In this way, the continuous time variables $\boldsymbol{w}$ are replaced by their discrete-time counterparts[2] $\boldsymbol{W}=[X_0^\top,X_{0,0}^\top,Y_{0,0}^\top,X_{0,1}^\top,Y_{0,1}^\top,...,X_{0,\mathfrak{D}-1}^\top,Y_{0,\mathfrak{D}-1}^\top,U_0^\top,X_1^\top,X_{1,0}^\top,Y_{1,0}^\top,...X_{\mathfrak{I}-1,\mathfrak{D}-1}^\top,Y_{\mathfrak{I}-1,\mathfrak{D}-1}^\top,U_{\mathfrak{I}-1}^\top,X_{\mathfrak{I}}^\top,\boldsymbol{\theta}^\top]^\top$.

Then, the continuous-form equality $\boldsymbol{g}=\boldsymbol{0}$ and inequality $\boldsymbol{h}\leq\boldsymbol{0}$ constraints are replaced by the discrete-form equality $\boldsymbol{G}=\boldsymbol{0}$ and inequality $\boldsymbol{H}\leq\boldsymbol{0}$ constraints, a large set of equations that are applied to the discrete-form variables $\boldsymbol{W}$. Further replacing any integrals in the objective function with numerical integrals transforms the continuous $\psi(\boldsymbol{w})$ to a discrete $\Psi(\boldsymbol{W})$, after which we have a discrete-form OCP that can be solved with any desired optimization solver.

$$\begin{aligned} \boldsymbol{W}^\star = \quad & \underset{\boldsymbol{W}}{\arg.\min.} \quad && \Psi(\boldsymbol{W}) && && (2) \\ & \text{subj. to} && \boldsymbol{G}(X_{\mathfrak{i}+1},X_{\mathfrak{i}},X_{\mathfrak{i},0},Y_{\mathfrak{i},0},X_{\mathfrak{i},1},Y_{\mathfrak{i},1},...,X_{\mathfrak{i},\mathfrak{D}-1},Y_{\mathfrak{i},\mathfrak{D}-1},U_{\mathfrak{i}},\boldsymbol{\theta}) && = \boldsymbol{0} && \forall \mathfrak{i} \in \{0,1,...,\mathfrak{I}-1\} \\ & && \boldsymbol{H}(X_{\mathfrak{i}},Y_{\mathfrak{i}},U_{\mathfrak{i}},\boldsymbol{\theta}) && \leq \boldsymbol{0} && \forall \mathfrak{i} \in \{0,1,...,\mathfrak{I}\}. \end{aligned}$$

2.2 Baseline optimal control problem

Having given that brief background, this paper's story starts with the continuous-time average-power-maximization AWE OCP described in De Schutter et al. (2023); implemented in the open-source software tool, the `awebox`; and solved (off-line) to generate the sorts of reference trajectories that one might later attempt to track (on-line) with a simpler controller. This baseline

[2]The discrete variables are also re-scaled to keep the condition number of the resulting problem's Karush–Kuhn–Tucker matrix down. But, we're presently leaving that portion out of the story for clarity's sake.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

problem includes no wake model - either in the dynamics $\boldsymbol{d}_{\text{bsln}}$ or otherwise in the OCP - and reads as:

$$
\begin{aligned}
\boldsymbol{w}^{\star} = \quad & \underset{\boldsymbol{w}}{\text{arg. min.}} \quad && \psi\left(\boldsymbol{x}(\cdot), \boldsymbol{y}(\cdot), \boldsymbol{u}(\cdot), \boldsymbol{\theta}, \boldsymbol{\varsigma}\right) && && && (3)\\
& \text{subj. to} && \boldsymbol{d}_{\text{bsln}}\left(\dot{\boldsymbol{x}}(t), \boldsymbol{x}(t), \boldsymbol{y}(t), \boldsymbol{u}(t), \boldsymbol{\theta}, \boldsymbol{\varsigma}\right) && = && \mathbf{0} && \forall t \in [0, t_{\text{final}}]\\
& && \boldsymbol{h}_{\text{path}}\left(\boldsymbol{x}(t), \boldsymbol{y}(t), \boldsymbol{u}(t), \boldsymbol{\theta}, \boldsymbol{\varsigma}\right) && \leq && \mathbf{0} && \forall t \in [0, t_{\text{final}}]\\
& && \boldsymbol{g}_{\boldsymbol{\theta}}\left(\boldsymbol{x}(t_{\text{switch}}), \boldsymbol{x}(t_{\text{final}}), \boldsymbol{\theta}, \boldsymbol{\varsigma}\right) && = && \mathbf{0} && \\
& && \boldsymbol{h}_{\boldsymbol{\theta}}\left(\boldsymbol{\theta}, \boldsymbol{\varsigma}\right) && \leq && \mathbf{0} && \\
& && \boldsymbol{\xi}\left(\boldsymbol{x}(0), \boldsymbol{x}(t_{\text{final}})\right) && = && \mathbf{0} &&
\end{aligned}
$$

where we introduce a new variable-type, the homotopy parameters $\boldsymbol{\varsigma}$ (pronounced 'sigma'), and enforce as constraints a set of path equalities that include the dynamics $\boldsymbol{d}_{\text{bsln}}$, path inequalities $\boldsymbol{h}_{\text{path}}$, parameter equalities $\boldsymbol{g}_{\boldsymbol{\theta}}$, parameter inequalities $\boldsymbol{h}_{\boldsymbol{\theta}}$, and a set of periodic boundary conditions $\boldsymbol{\xi}$. We also include a switching time $t_{\text{switch}} \in (0, t_{\text{final}})$ into the system parameters, which describes the moment when the system switches from reeling-out motion to reeling-in motion.

As mentioned above, the components of this baseline problem are discretized then scaled, steps that are not printed here for the sake of legibility; we refer the reader to De Schutter et al. (2023) for expanded definitions of the problem components. Reproducibility and the story-to-follow do require, however, a bit of expansion about the nature of the objective function and the path constraints.

First, the objective function of the above AWE OCP is constructed so that the problem can be deformed from a 'tracking problem' whose objective is dominated by an L2-penalization (with weight $\underline{\underline{\boldsymbol{k}}}_x$) of the difference between the states ($\boldsymbol{x}$) and a given reference ($\boldsymbol{x}_{\text{Ref}}$), to a 'power problem' whose objective should allow us to maximize the system's average power $\overline{P}$ (weighted with a factor k_{P}), by decreasing the power homotopy parameter ς_{P} from 1 to 0.

$$
\psi \quad = \quad \varsigma_{\text{P}} \left(\frac{\int_0^{t_{\text{final}}} \left(\left(\boldsymbol{x}(t) - \boldsymbol{x}_{\text{Ref}}(t)\right)^{\top} \underline{\underline{\boldsymbol{k}}}_x \left(\boldsymbol{x}(t) - \boldsymbol{x}_{\text{Ref}}(t)\right) \right) \mathrm{d}t}{t_{\text{final}}} \right) + \left(1 - \varsigma_{\text{P}}\right) \left(-k_{\text{P}} \overline{P}\right) + \psi_{\text{reg}} + \psi_{\text{hmtp}}. \tag{4}
$$

In both the tracking- and power- problems, we additionally penalize extreme inputs and system designs using a regularization cost ψ_{reg}. We also have a homotopy cost ψ_{hmtp} that we'll use to push our problem through the predictor-corrector steps of the penalty-based homotopy.

The average power $\overline{\text{P}}$ mentioned in (4) is the temporal-average of the instantaneous mechanical power $P(t)$ over the entire optimization period t_{final}. This instantaneous power is itself found - following the approach of Gros and Diehl (2013) - as the product of the main tether's reel-out speed $\dot{\ell}_{\text{t}}$, the main tether length ℓ_{t} and the tension-per-unit-length in the main tether τ_1, which happens in Gros and Diehl (2013)'s specific approach to be the multiplier of the Lagrangian dynamics' holonomic constraint on the main tether length. This reads as:

$$
\overline{P} = \frac{\int_0^{t_{\text{final}}} P(t)\, \mathrm{d}t}{t_{\text{final}}}, \qquad P(t) = (\tau_1 \ell_{\text{t}}) \dot{\ell}_{\text{t}}. \tag{5}
$$

One point that will become relevant in the final section of this paper, is that these three values - $\dot{\ell}_{\text{t}}$, ℓ_{t}, and τ_1 - are all directly represented among the AWE OCP's model variables, $\dot{\ell}_{\text{t}}$ and ℓ_{t} among the problem's states, and τ_1 among its algebraic variables.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

Second, the `awebox` allows a user to modularly select path-inequalities. As such, the example problems of Sections 4 and 5 will use a 'tether stress' inequality, which prevents the tether elements from experiencing stress above a safety factor $\Phi_{\text{safety}_\sigma}$ of the tether material's yield stress σ_{ub}; an 'airspeed' inequality, which acts as a substitute for a stall- and compressibility-model by keeping the kite's effective airspeed (norm of the effective velocity) within the bounds $v_{\text{eff},\mathfrak{k}\text{lb}}$ and $v_{\text{eff},\mathfrak{k}\text{ub}}$; an 'aerodynamic validity' inequality, which prevents each kite's effective velocity from leaving the linearization-range of the stability derivatives that are used to relate the aerodynamic forces on the kite to the effective velocity, applying bounds on the angle of attack $\alpha_{\mathfrak{k}\text{lb}}$ and $\alpha_{\mathfrak{k}\text{ub}}$, and sideslip angle $\beta_{\mathfrak{k}\text{lb}}$ and $\beta_{\mathfrak{k}\text{ub}}$; an 'anti-collision' inequality, which prevents collisions by ensuring that the kites remain farther apart than $\Phi_{\text{safety}_{\text{acoll}}}$-multiples of the wingspan; and a 'yaw-rotation' inequality, which limits each kite's maximum yaw-angle to remain between $\varphi_{\text{yaw}_{\mathfrak{k}\text{lb}}}$ and $\varphi_{\text{yaw}_{\mathfrak{k}\text{ub}}}$.[3] We also apply a number of simple bounds on various decision variables including the tether length and the tether acceleration, and impose a schedule on the tether's reeling-out and -in motion to give the optimizer a preference in the temporal-orientation of the trajectory. Some specific values for these bounds, which we'll use in this paper unless otherwise specified, are given in Table B3.

Following the discretization procedure summarized in Section 2.1 and some cosmetic rearrangement, the continuous-form AWE OCP (3) can be converted into a discrete-form AWE OCP (6):

$$
\begin{aligned}
\boldsymbol{W}_{\text{bsln}}{}^{\star} = \quad & \underset{\boldsymbol{W}_{\text{bsln}}}{\text{arg. min.}} \quad & & \varsigma_{\text{P}}\Psi_{\text{track}} + (1-\varsigma_{\text{P}})\left(-k_{\text{P}}\overline{P}\right) + \Psi_{\text{reg}} + \Psi_{\text{hmtp}} & & & & \text{from (4)} \qquad (6)\\
& \text{subj. to} & & \boldsymbol{G}_{\text{bsln}}(\boldsymbol{W}_{\text{bsln}}) & = & \mathbf{0} & & \text{from (3)}\\
& & & \boldsymbol{H}_{\text{bsln}}(\boldsymbol{W}_{\text{bsln}}) & \leq & \mathbf{0} & & \text{from (3)}
\end{aligned}
$$

where $\boldsymbol{W}_{\text{bsln}}$ are the discrete-form baseline AWE OCP decision variables, Ψ_{track}, Ψ_{reg}, and Ψ_{hmtp} are the discrete versions of the tracking, regularization, and homotopy costs from (4); and $\boldsymbol{G}_{\text{bsln}}$ and $\boldsymbol{H}_{\text{bsln}}$ are the discrete versions of the concatenated equality and inequality constraints, respectively, from (3).

2.3 Model concept

The goal of this paper is to introduce a vortex filament model into the baseline AWE OCP (6), with the problem remaining tractable. To do this, we will model only a certain portion of the vorticity trailed and shed behind the kites: no matter at what time t one considers, only the vorticity shed during the last $\mathfrak{A}$ discrete time-steps into the past will be represented. Conceptually, this is similar to the forward-simulation practice (see Katz and Plotkin (2010)) of either forgetting or consolidating the oldest vortex element each time a new vortex element is shed.

More specifically, this wake model will be built from vortex filaments, with constraints that will fix the filaments' positions and strengths over time. We will specifically discretize the wake using $3\mathfrak{A}\mathfrak{K}$ vortex filaments, comprising: one (per kite) bound filament; $(\mathfrak{A}-1)$ (per kite) shed filaments of originally-spanwise shed vorticity that have now been convected downstream; $2(\mathfrak{A}-1)$ (per kite) trailing filaments that connect the wingtips' bound- and previously-shed- wake nodes into two irregular

[3] The letter "$\mathfrak{k}$" is pronounced as "k" and corresponds to that letter in LaTeX's old-Germanic font `mathfrak`. For help recognizing other unusual characters, the reader is kindly referred to this work's nomenclature, which begins on page 55.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

helices; plus two (per kite) semi-infinite filaments parallel to the freestream, leaving from the most-downstream wake-nodes as a stand-in for the vorticity shed over the history of the system's periodic motion. This structure is sketched in Figure 2, with a detailed explanation of the various filaments' strengths, starting positions, and ending positions listed in Table A1.

Naturally, there are many ways that someone could fix the positions and strengths of those wake nodes shed prior to the actual start of the optimization period $t < 0\,\mathrm{s}$. But since the baseline AWE OCP given in (3) already includes periodic boundary conditions, we will simply recycle our current trajectory to give the necessary 'historical' information. Returning the method to a non-periodic trajectory for either launch and landing or non-AWE applications would hopefully only require a small adjustment from the method laid out here.

2.4 Wake structure

In order to limit computational memory, we will separate the wake into three regions: the bound vorticity along the lifting-line of the kite; a near-wake region where the vorticity in the flow is resolved; and a far-wake region where the only modeling of vorticity serves to close open horseshoe-vortices, conserving vorticity in this assumed externally-potential flow and representing the periodic history of the system's motion. This free-wake region, specifically, is intended to help avoid spurious vortex roll-up in the future step after progressing from rigid convection (this paper) to free-wake convection (future work).

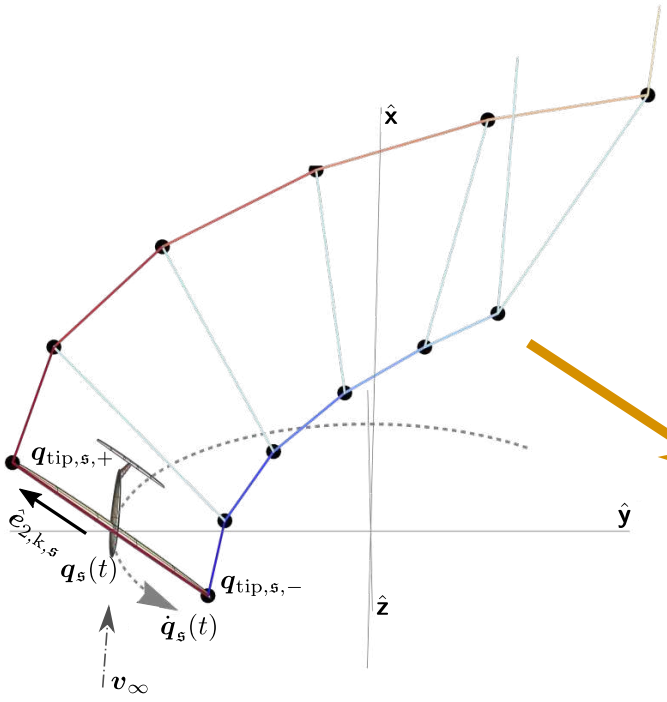
Since this method is intended primarily for applications where little information is known about the exact kite geometry and load distribution - the kite is modeled only as a rigid, 6 degree-of-freedom (DOF) object with a known mass, wingspan, moment of inertia, and stability derivatives, rather than a schematic layout of the geometries of all lifting and control surfaces - describing the bound vorticity in any more specificity[4] than a single lifting-line vortex of constant circulation would not just require significantly more computational effort because the vorticity shed from additional spanwise-positions would also need to be tracked, but also risks producing misleading[5] results.

Therefore, the near-wake region is formed by connecting a collection of convected wake nodes - each shed from one end of a lifting-line - with straight, constant-strength vortex filaments. At any moment in the (discrete) time within the modeled universe, this discretization includes $\mathfrak{A}$ wake nodes per wing-tip, per kite. That is, when $\mathfrak{A} = 1$, the only wake nodes to be modeled would be the end-points of the bound-vortices at each of the kites' wingtips; for every additional increment to $\mathfrak{A}$, the vorticity shed during one previous control interval is included in the tracked wake. We'll use the wake node index $\mathfrak{a}$ to

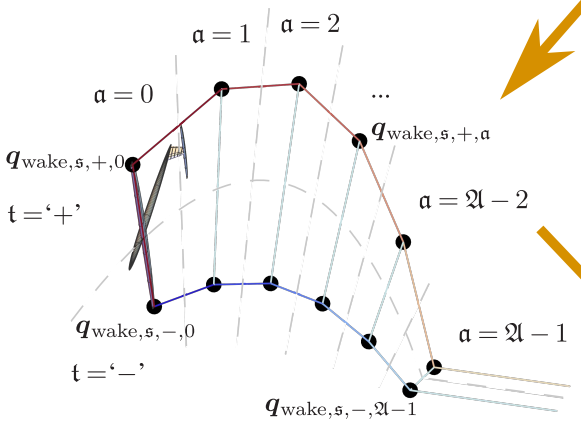
[4]The idea is that we don't know what circulation distribution is correct, but we anyways need to choose one. And, since a 'trajectory optimization problem based on very-little information about the kite's geometry but nevertheless solved for the purpose of performance optimization' is a problem that is likely to be especially relevant at the beginning of an iterative design process, we specifically want to choose a circulation distribution that will introduce the least amount of bias into the problem's solution. As such, this question seems conceptually similar to the choice of 'prior' for a Bayesian update, for which - we observe - uniform priors are the classical (Bayes, 1763) recommendation.

[5]Using any more-specific circulation distribution risks accidentally suggesting to readers, `awebox` users, and the authors themselves, that the additional specificity represents additional knowledge about the actual circulation distribution on the kites, whereas using a uniform distribution should hopefully make it clear - immediately, from a single glance at the two-trailing-vortices-per-kite in the plots - that no such distribution knowledge exists. This visual-demonstration of uncertainty would not happen if we assumed, instead, an elliptical circulation distribution. Furthermore, considering the fact that elliptic circulation wings are known (Raymer, 1992) to perform badly in stall conditions and that AWE OCPs typically return solutions with saturated angle-of-attack upper bounds, we consider it unlikely that an assumption of elliptical circulation would improve the physicality of the method.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

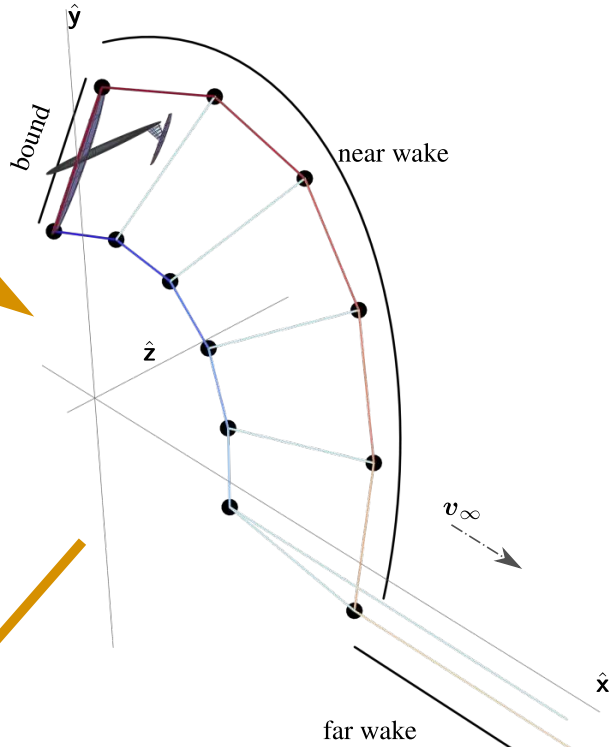


(a, top left). Each kite sheds vorticity at the current position of its wingtips. That vorticity is then convected with the freestream. The vorticity distribution is discretized into $3\mathfrak{A}$ (per kite) vortex filaments (colored lines), joining $2\mathfrak{A}$ (per kite) wake nodes (black dots).



(c, middle left). The wake nodes are labelled according to their shedding kite $\mathfrak{s}$, their shedding wingtip $\mathfrak{t} \in \{+,-\}$, and the number of control intervals into the past when they were shed $\mathfrak{a} \in \{0,1,...,2\mathfrak{A}-1\}$.

(b, middle right). The vortex filaments belong to 3 regions: bound filaments do not induce velocity on their shedding kite; the near wake discretizes the resolved ('tracked') portion of the wake; and the far wake is a placeholder for the truncated vorticity.



(d, bottom right). Visualizations show the vortex filaments with color darkness indicating the vortex strength (vorticity magnitude). The colors are red-toned if the net vorticity runs parallel to the downstream- and spanwise-scaffold (black arrows), and blue-toned if antiparallel. In this example: the kite's circulation has increased over time.

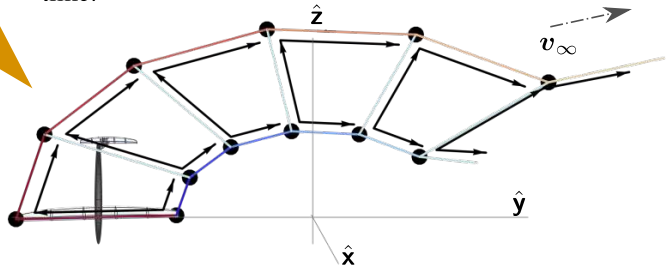
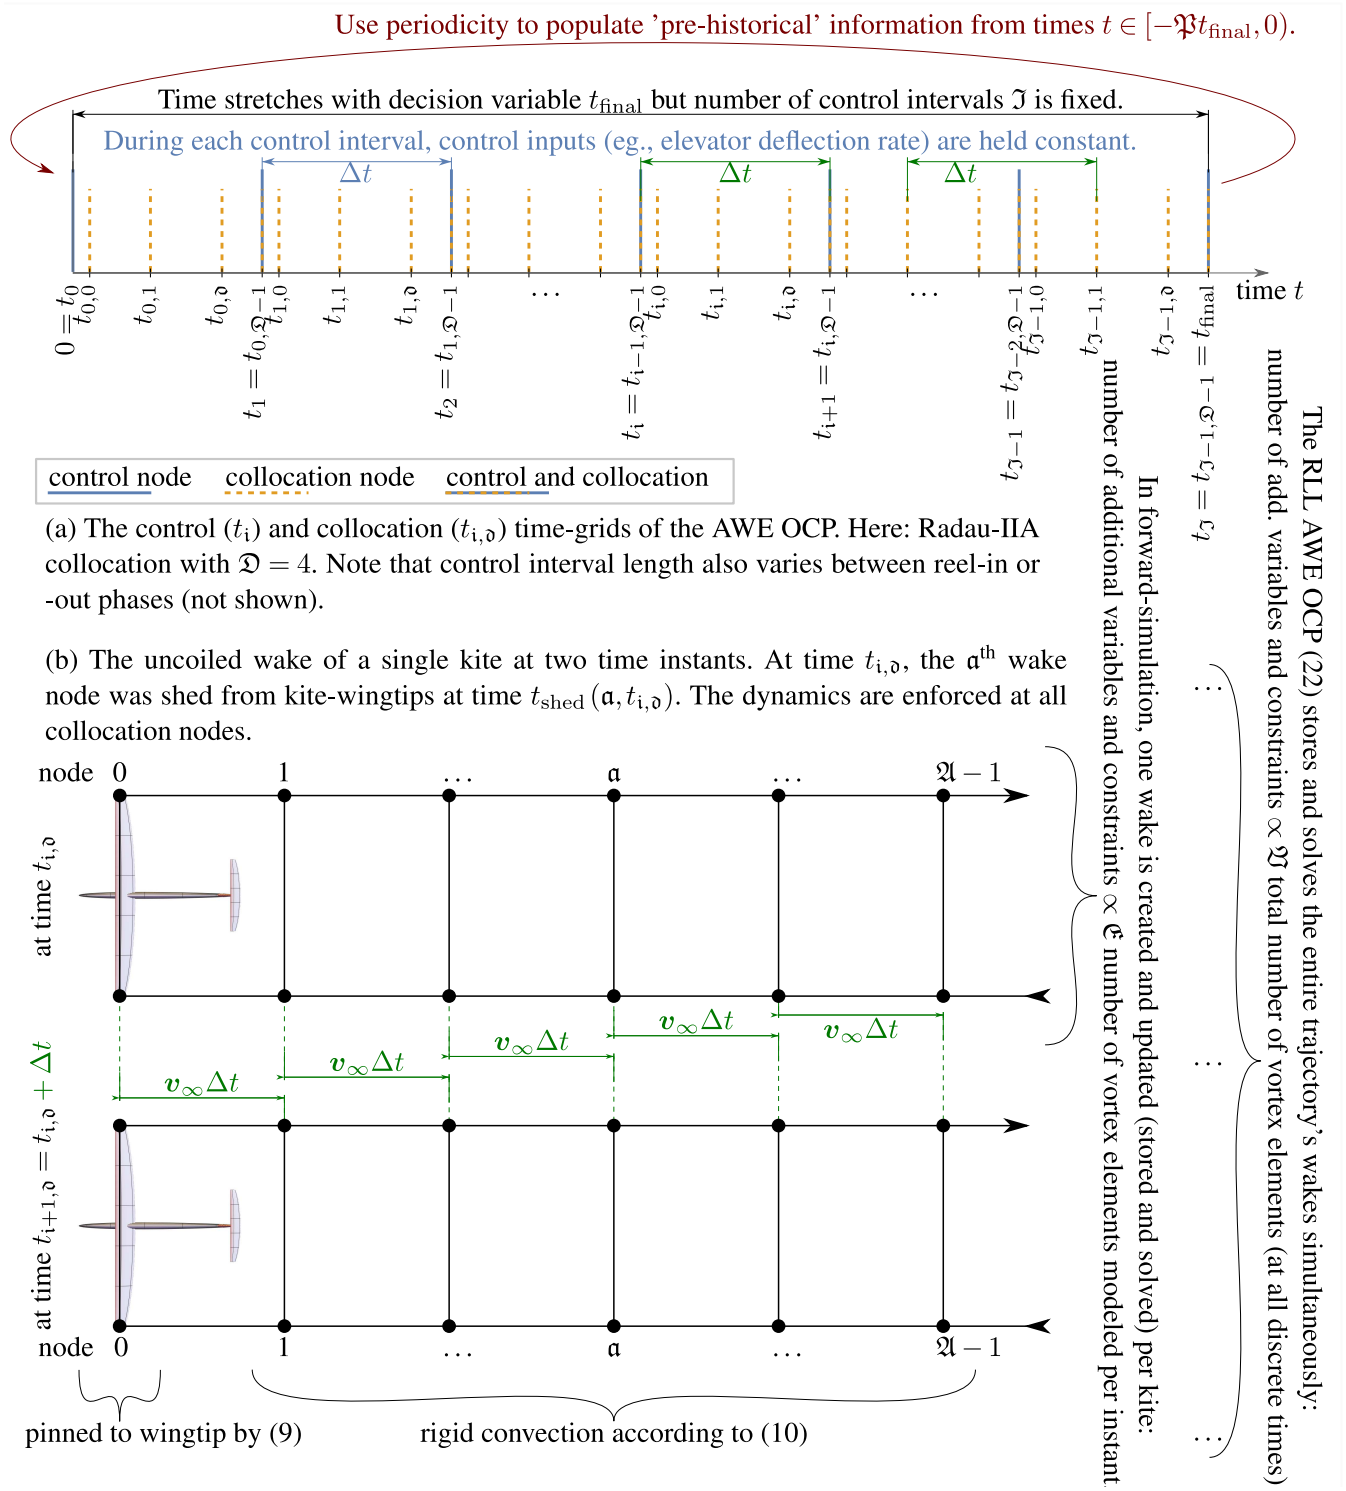


Figure 2. A sketch (in four viewpoints) to illustrate the vorticity in the wake behind a single kite according to the RLL model with $\mathfrak{A}=6$.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



(a) The control (t_{i}) and collocation ($t_{\mathrm{i},\mathfrak{d}}$) time-grids of the AWE OCP. Here: Radau-IIA collocation with $\mathfrak{D} = 4$. Note that control interval length also varies between reel-in or -out phases (not shown).

(b) The uncoiled wake of a single kite at two time instants. At time $t_{\mathrm{i},\mathfrak{d}}$, the $\mathfrak{a}^{\text{th}}$ wake node was shed from kite-wingtips at time $t_{\text{shed}}(\mathfrak{a}, t_{\mathrm{i},\mathfrak{d}})$. The dynamics are enforced at all collocation nodes.

Figure 3. A sketch showing the time-grids of the AWE OCP as well as the rigid convection of the wake nodes - also explaining why the computational cost of the RLL AWE OCP will be so large.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

count through the downstream locations of these wake nodes. Then, the wake node described by $\mathfrak{a} = 0$ is located at the kite's wingtip on the lifting-line, where it is - in that specific instance - being shed; whereas the wake node described by $\mathfrak{a} = \mathfrak{A} - 1$ is positioned where the downstream edge of the near wake meets the far wake.

These wake nodes are shed from the kite's wingtips, with 'past shedding' times being the associated number of control-lengths[6] prior to the 'present shedding' time. The shedding-time-grid is coupled to the time-grid of the OCP's discretization, because the chosen method of discretization means that the dynamics and inequality constraints are only enforced on the collocation nodes[7], and the optimizer is judged to have less room to artificially improve the objective by manipulating the discretization error - that is, by cheating physics[8] - if vortex shedding occurs from time-points where the constraints are enforced. Following this strategy, we specify the shedding time $t_{\text{shed}}\left(\mathfrak{a}, t_{\mathfrak{i},\mathfrak{d}}\right)$ of the wake node that is the $\mathfrak{a}^{\text{th}}$ wake node at the discretized time $t_{\mathfrak{i},\mathfrak{d}}$ at control index $\mathfrak{i}$ and collocation index $\mathfrak{d}$. (This is sketched in Figure 3.) If this $\mathfrak{a}^{\text{th}}$ wake node was shed within this optimization period, then the shedding time must have been $\mathfrak{a}$ control intervals before $t_{\mathfrak{i},\mathfrak{d}}$; since the trajectory is periodic, any wake node shed prior to the optimization period, needs its shedding time to be corrected according to the number of integer whole-optimization-periods into the past, that the node was shed. This reads as:

$$t_{\text{shed}}\left(\mathfrak{a}, t_{\mathfrak{i},\mathfrak{d}}\right) = t_{\mathfrak{i}-\text{mod}(\mathfrak{a},\mathfrak{I}),\mathfrak{d}} - \text{floor}\left(\frac{\mathfrak{a}}{\mathfrak{I}}\right) t_{\text{final}}, \tag{7}$$

where $\mathfrak{I}$ is the number of control intervals, t_{final} is the optimization period (the period of the AWE system's entire periodic orbit, during which the trajectory may make more than one winding or looping $\mathfrak{L}$), and $\text{floor}\left(\right)$ is the floor function.

We can similarly relate the number of wake nodes $\mathfrak{A}$ and the number of control intervals $\mathfrak{I}$ to describe how many periods 'worth' of the wake are resolved $\mathfrak{P}$:

$$\mathfrak{P} = \frac{\mathfrak{A} - 1}{\mathfrak{I}}. \tag{8}$$

Those wake nodes that are immediately being shed ($\mathfrak{a} = 0$) can be found on the wingtip $\boldsymbol{q}_{\text{tip},\mathfrak{s},\mathfrak{t}}$ from which they are being shed:

$$\boldsymbol{r}_{\text{tip}} = \left(\boldsymbol{q}_{\text{wake},\mathfrak{s},\mathfrak{t},0}\left(t\right) - \boldsymbol{q}_{\text{tip},\mathfrak{s},\mathfrak{t}}\left(t\right) \quad \text{for } \mathfrak{s} \in \mathcal{K}, \mathfrak{t} \in \mathcal{T}\right) = \boldsymbol{0}, \tag{9}$$

$$\boldsymbol{q}_{\text{tip},\mathfrak{s},\mathfrak{t}}\left(t\right) = \boldsymbol{q}_{\mathfrak{s}}\left(t\right) + \frac{b}{2} \begin{cases} (+1)\, \hat{\boldsymbol{e}}_{2,\text{k},\mathfrak{s}}(t) & \text{if } \mathfrak{t} = \text{`}+\text{'} \\ (-1)\, \hat{\boldsymbol{e}}_{2,\text{k},\mathfrak{s}}(t) & \text{if } \mathfrak{t} = \text{`}-\text{'} \end{cases}$$

when $\mathfrak{s}$ is the index of the architectural-node containing the vortex-shedding kite within the set of all kites $\mathcal{K}$, $\mathfrak{t}$ is the index of the wingtip within the set of wingtips $\mathcal{T} = \{\text{`}+\text{'}, \text{`}-\text{'}\}$, b is the kite's wingspan, and $\hat{\boldsymbol{e}}_{2,\text{k},\mathfrak{s}}$ is the unit vector pointing along the

[6]The collocation nodes are unevenly spaced in time; shedding vorticity at time-intervals determined by the collocation grid (instead of the more-regular control grid) would lead to a substantial difference in filament length, rendering the numerical work done in relation to the vortex-model variables less cost-effective for the short filaments shed at the end of a control interval than for the long filaments shed at the beginning of a control interval.

[7]More precisely: the inequality constraints are enforced on all collocation nodes, except for the simple bounds on the states that are the first anti-derivatives of piecewise constant controls, which would otherwise cause linear-constraint-dependence when active ($\dot{\ell}_{\text{tlb}} \leq \dot{\ell}_{\text{t}} \leq \dot{\ell}_{\text{tub}}$ and $\delta_{\text{cslb}} \leq \delta_{\text{cs}} \leq \delta_{\text{csub}}$). The approach of avoiding this mode of LICQ-violation while also reducing computational cost, by only enforcing the inequalities on the control nodes is not recommended if the inequality constraints include the validity bounds of sub-models, such as bounds on the angle-of-attack and airspeed; practical experience suggests that the optimizer will violate these particular constraints as far as the enforcement-spacing allows.

[8]The famous 'Theerum' of Betts (2010) applies: "If there is a flaw in the problem formulation, the optimization algorithm will find it."

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

$\mathfrak{s}^{\text{th}}$ kite's span from the 'negative' to 'positive' wingtip. Meanwhile, all other wake nodes are 'fixed' at the positions to which they have been convected during the preceding control interval. Under the assumption that the induced velocities are relatively low, the convected wake nodes will be convected rigidly by a freestream velocity $\boldsymbol{v}_\infty$, which is constant in time and in the downstream and horizontal directions $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$, but which may vary in the vertical direction $\hat{\mathbf{z}}$.

Since this position-fixing procedure relates the algebraic variables between different time instances, we can't include the convection constraint into the dynamics[9]. Instead, we'll lift the positions $\boldsymbol{q}_{\text{wake},\mathfrak{s},\mathfrak{t},\mathfrak{a}}$ of the $\mathfrak{a}^{\text{th}}$ wake node, shed from the $\mathfrak{s}^{\text{th}}$ kite's $\mathfrak{t}^{\text{th}}$ wingtip, at all of the discretized collocation nodes into the OCP's algebraic variables. These additional variables ($\boldsymbol{q}_{\text{wake},\mathfrak{s},\mathfrak{t},\mathfrak{a}_{\mathfrak{i},\mathfrak{d}}} \in \mathbb{R}^3$, for $\mathfrak{s} \in \mathcal{K}, \mathfrak{t} \in \mathcal{T}, \mathfrak{a} \in \mathcal{A}, \mathfrak{i} \in \{0,1,...,\mathfrak{I}-1\}, \mathfrak{d} \in \{0,1,...,\mathfrak{D}-1\}$) are pinned to the appropriate positions at the collocation nodes by including (9) into the dynamics, and separately introducing the following convection relation directly as an equality constraint into the discrete-form of the problem:

$$\boldsymbol{G}_{\text{RLL},\boldsymbol{q}} = \begin{pmatrix} \boldsymbol{q}_{\text{wake},\mathfrak{s},\mathfrak{t},\mathfrak{a}_{\mathfrak{i},\mathfrak{d}}} - \left(\boldsymbol{q}_{\text{wake},\mathfrak{s},\mathfrak{t},\mathfrak{a}_{\mathfrak{i}-1,\mathfrak{d}}} \cdots \right. & \text{for } \mathfrak{s} \in \mathcal{K}, \mathfrak{t} \in \mathcal{T}, \mathfrak{a} \in \{1,2,...,\mathfrak{A}-1\}, \ldots \\ \left. + (t_{\mathfrak{i},\mathfrak{d}} - t_{\mathfrak{i}-1,\mathfrak{d}})\, \boldsymbol{v}_\infty \left(\boldsymbol{q}_{\text{wake},\mathfrak{s},\mathfrak{t},\mathfrak{a}_{\mathfrak{i}-1,\mathfrak{d}}}\right)\right), & \mathfrak{i} \in \{0,1,...,\mathfrak{I}-1\}, \mathfrak{d} \in \{0,1,...,\mathfrak{D}-1\} \end{pmatrix} = \mathbf{0}, \tag{10}$$

where $\mathfrak{D}$ is the number of collocation nodes.

2.5 Filament strengths

The AWE OCP described in (3) and formulated in the `awebox` assumes that the user knows the relationship between the inflow experienced on-board the kite and the aerodynamic force on the kite:

$$\boldsymbol{F}_{\text{aero},\mathfrak{s}} = (C_{\text{D,eff},\mathfrak{s}}\, \hat{\boldsymbol{e}}_{\text{D,eff},\mathfrak{s}} + C_{\text{S,eff},\mathfrak{s}}\, \hat{\boldsymbol{e}}_{\text{S,eff},\mathfrak{s}} + C_{\text{L,eff},\mathfrak{s}}\, \hat{\boldsymbol{e}}_{\text{L,eff},\mathfrak{s}})\, s \left(\frac{1}{2}\rho_{\text{air}}(\boldsymbol{q}_\mathfrak{s}) \|\boldsymbol{v}_{\text{eff},\mathfrak{s}}\|_2^2\right) = \boldsymbol{D}_{\text{eff},\mathfrak{s}} + \boldsymbol{S}_{\text{eff},\mathfrak{s}} + \boldsymbol{L}_{\text{eff},\mathfrak{s}}, \tag{11}$$

where $\boldsymbol{F}_{\text{aero},\mathfrak{s}}$ is the aerodynamic force on the $\mathfrak{s}^{\text{th}}$ kite; $\hat{\boldsymbol{e}}_{\text{D,eff},\mathfrak{s}}$, $\hat{\boldsymbol{e}}_{\text{S,eff},\mathfrak{s}}$, and $\hat{\boldsymbol{e}}_{\text{L,eff},\mathfrak{s}}$ are the unit vectors for projecting the aerodynamic force in the drag-, side-, and lift-force directions according to the effective velocity; and $C_{\text{D,eff},\mathfrak{s}}$, $C_{\text{S,eff},\mathfrak{s}}$, and $C_{\text{L,eff},\mathfrak{s}}$, and $\boldsymbol{D}_{\text{eff},\mathfrak{s}}$, $\boldsymbol{S}_{\text{eff},\mathfrak{s}}$, and $\boldsymbol{L}_{\text{eff},\mathfrak{s}}$ are the respective coefficients and aerodynamic forces in those directions, as determined by projecting the total aerodynamic force in the effective-velocity-determined coordinate frame. We can further specify that s is the kite's planform area; $\boldsymbol{q}_\mathfrak{s}$ is the position of the $\mathfrak{s}^{\text{th}}$ kite; $\rho_{\text{air}}(\boldsymbol{q}_\mathfrak{s})$ is the air density at the $\mathfrak{s}^{\text{th}}$ kite's position; and $\boldsymbol{v}_{\text{eff},\mathfrak{s}}$ is the effective velocity observed by the $\mathfrak{s}^{\text{th}}$ kite. This effective velocity is the sum of the freestream velocity at the kite's position $\boldsymbol{v}_\infty(\boldsymbol{q}_\mathfrak{s})$, the induced velocity at the kite's position $\boldsymbol{v}_{\text{ind},\mathfrak{s}}$, and the relative inflow due to the kite's own motion $\dot{\boldsymbol{q}}_\mathfrak{s}$. Put together, this reads as:

$$\boldsymbol{v}_{\text{eff},\mathfrak{s}} = \boldsymbol{v}_\infty(\boldsymbol{q}_\mathfrak{s}) + \boldsymbol{v}_{\text{ind},\mathfrak{s}} - \dot{\boldsymbol{q}}_\mathfrak{s}. \tag{12}$$

[9]Formulating this problem by representing the wake node positions as states - whose convection could be easily handled by adding the convection PDE to the OCP's dynamics - and shedding the circulation directly from the kites onto the vortex elements moving past them, unfortunately requires many more vortex elements than can be meaningfully solved with a personal computer. First, because the tracked wake needs to be long enough to give a reasonable representation of the historically shed vorticity already at the start of the optimization period and not just at the end. And, second, because any insufficiently-fine resolution of the wake passing over the kites will encourage the optimizer to manipulate the integration error to create a non-physical advantage (Harzer et al., 2025), such as by 'flapping' the control surfaces.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

The decision to use the effective velocity as both the reference velocity for the linearized stability derivatives and the projection direction for the aerodynamic forces - $\boldsymbol{F}_{\text{aero},\mathfrak{s}} = \boldsymbol{F}_{\text{aero},\mathfrak{s}}(\boldsymbol{v}_{\text{eff},\mathfrak{s}}) = \boldsymbol{F}_{\text{aero},\mathfrak{s}}(\boldsymbol{v}_\infty(\boldsymbol{q}_\mathfrak{s}) - \dot{\boldsymbol{q}}_\mathfrak{s} + \boldsymbol{v}_{\text{ind}}(\boldsymbol{q}_\mathfrak{s}))$, and $\boldsymbol{T}_{\text{aero},\mathfrak{s}} = \boldsymbol{T}_{\text{aero},\mathfrak{s}}(\boldsymbol{v}_{\text{eff},\mathfrak{s}}) = \boldsymbol{T}_{\text{aero},\mathfrak{s}}(\boldsymbol{v}_\infty(\boldsymbol{q}_\mathfrak{s}) - \dot{\boldsymbol{q}}_\mathfrak{s} + \boldsymbol{v}_{\text{ind}}(\boldsymbol{q}_\mathfrak{s}))$ - was made because it renders the process of including any desired wake-model into the AWE OCP very straightforward without risking 'double-counting' the induction effect, as rightly cautioned by Gaunaa et al. (2020) and Trevisi et al. (2023). (And, having decided to represent the aerodynamic force as a function of the effective velocity, we then decide to use the Ampyx AP2 as a test-kite in the example problems of Section 4 and 5, because textual clues in Malz et al. (2019), Malz (2020, pp.30-31), and Licitra (2018, Ch.3 & p.134) suggest that the stability derivatives of the Ampyx AP2 reported in Malz et al. (2019) (see Table B2) might have been given in relation to the inflow velocities measured on-board[10] the kite.) However, for the case of stability derivatives that are not based on the effective velocity, the here-described method could still be used, though it would require a local correction to find the approximate apparent velocity $(\boldsymbol{v}_\infty - \dot{\boldsymbol{q}}_\mathfrak{s})$ that a kite must have had to experience the effective velocity described in (12) but with a wake of whatever-geometry was flown-or-simulated when the stability derivatives were determined.

Then, the aerodynamic forces are carried into the $\mathfrak{s}^{\text{th}}$ kite's translational dynamics using a Lagrangian formulation:

$$\frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial \mathcal{L}_{\text{mech}}}{\partial \dot{\boldsymbol{q}}_\mathfrak{s}} - \frac{\partial \mathcal{L}_{\text{mech}}}{\partial \boldsymbol{q}_\mathfrak{s}} = \boldsymbol{F}_{\text{aero},\mathfrak{s}} + \boldsymbol{F}_{\text{LMC}\mathfrak{s}} + \boldsymbol{D}_{\text{t},\mathfrak{s}} + \varsigma_{\text{fict}}\boldsymbol{F}_{\text{fict}\mathfrak{s}}, \tag{13}$$

where $\boldsymbol{F}_{\text{LMC}\mathfrak{s}}$ is a momentum correction for open systems, such as lift-mode AWE systems whose main tether is being rolled or unrolled from an unmodeled generator drum; $\boldsymbol{D}_{\text{t},\mathfrak{s}}$ is the portion of the tether drag that has been assigned to the architectural node $\mathfrak{s}$; and $\boldsymbol{F}_{\text{fict}\mathfrak{s}}$ is an imaginary 'fictitious' force acting on the $\mathfrak{s}^{\text{th}}$ kite that is used to stabilize the problem early in the solution homotopy but which is later eliminated during the specific 'fictitious-forces' predictor-corrector homotopy step, when the corresponding homotopy parameter ς_{fict} is driven to zero, in a process similar to that developed by Horn et al. (2013). For more detail about the various components of the translational dynamics (13), as well as an analogous expression of the rotational dynamics, the reader is referred to De Schutter et al. (2023).

In any case, having a known aerodynamic force means that we can determine the lift force due to the effective velocity $\boldsymbol{L}_{\text{eff},\mathfrak{s}}$ by projecting the aerodynamic force in this lift-due-to-effective-velocity direction $\boldsymbol{L}_{\text{eff},\mathfrak{s}} = (\boldsymbol{F}_{\text{aero},\mathfrak{s}}^\top \hat{\boldsymbol{e}}_{\text{L,eff},\mathfrak{s}})\hat{\boldsymbol{e}}_{\text{L,eff},\mathfrak{s}}$. Then, we can use generalized Kutta-Joukowski (Katz and Plotkin, 2010; Saffman, 1995) to determine that the circulation experienced by the $\mathfrak{s}^{\text{th}}$ kite, at any instant in time, must read as:

$$\frac{\boldsymbol{L}_{\text{eff},\mathfrak{s}}}{b} = \rho_{\text{air}}\Gamma_{\text{k},\mathfrak{s}}(\boldsymbol{v}_{\text{eff},\mathfrak{s}} \times \hat{\boldsymbol{e}}_{2,\text{k},\mathfrak{s}}) \quad \Rightarrow \quad \Gamma_{\text{k},\mathfrak{s}} = \frac{1}{b\rho_{\text{air}}}\frac{(\boldsymbol{F}_{\text{aero},\mathfrak{s}}^\top \hat{\boldsymbol{e}}_{\text{L,eff},\mathfrak{s}})}{\hat{\boldsymbol{e}}_{\text{L,eff},\mathfrak{s}}^\top(\boldsymbol{v}_{\text{eff},\mathfrak{s}} \times \hat{\boldsymbol{e}}_{2,\text{k},\mathfrak{s}})}. \tag{14}$$

(Recall from (12), that we've evaluated the effective velocity $\boldsymbol{v}_{\text{eff},\mathfrak{s}}$ at the $\mathfrak{s}^{\text{th}}$ kite's position, which serves here both as center of mass and center of lift.) Notice that the latter force terms of (13) have no impact on the bound - and, later, shed - circulation: the drag on the tether's non-rotating cylinder not being associated with circulation in the highly-simplified externally-potential flow, the momentum correction not being produced aerodynamically, and the fictitious forces being entirely nonphysical.

Since the length of the optimization period in seconds stretches with the decision variable t_{final}, while the homotopy solution method tends to maintain the number of windings $\mathfrak{L}$ from the initial guess into the final trajectory, it makes most sense to

[10] most probably, at the Pitot tube. But, as the method specifically precludes knowledge of the kite's planform layout, we will here use the induced velocity evaluated at the center of mass.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

consider the shedding frequency in terms of 'shedding-instances-per-winding'. At roughly 20-30 control nodes (shedding instances) per winding, the vortex elements in the RLL AWE OCP are likely to be shed at relatively long intervals. For comparison, the above number of shedding-instances-per-winding is roughly half that used[11] in the two-kite vortex filament simulation of Rancourt and Mavris (2018). Due to the resulting relatively-large vortex rings, the bound vortex on the $\mathfrak{s}^{\text{th}}$ kite, has its strength set as the time-averaged[12] circulation $\overline{\Gamma}_{\mathfrak{s},\mathfrak{a}}$ during the time interval between shedding the $\mathfrak{a}^{\text{th}}$ wake node and the $(\mathfrak{a}+1)^{\text{th}}$ wake node. This averaging interval changes very little[13] over the optimization period, and has the same length as a full control interval.

One of the major advantages of elementary solutions of the Biot-Savart integral with constant vorticity distributions, is that the resulting induced velocity is linearly dependent on the vortex strength. To preserve this feature, as well as isolate the complication due to the averaging, we'll lift[14] the individual average-circulation $\overline{\Gamma}_{\mathfrak{s},\mathfrak{a}}$ terms into the AWE OCP's decision variables. On the other hand, we won't lift the filament strengths of the vortex elements $\Gamma_{\mathfrak{e}}$ because there are three times as many elements $\mathfrak{e}$ in the set of all elements $\mathcal{E}$ as there are paired shedding kite indices $\mathfrak{s}$ and wake nodes $\mathfrak{a}$ in, respectively, the set of all kites $\mathcal{K}$ and the set of all wake nodes $\mathcal{A}$; and, anyways, the relationship between $\overline{\Gamma}_{\mathfrak{s},\mathfrak{a}}$ and the shed horseshoe-vortices' elements' strengths $\Gamma_{\mathfrak{e}}$ is linear.

Lifting the average circulation of the shed kite at the time of shedding each wake node at all of the discrete collocation nodes into the OCP's decision space, introduces the following algebraic variables: $\overline{\Gamma}_{\mathfrak{s},\mathfrak{a}_{\mathfrak{i},\mathfrak{d}}} \in \mathbb{R}^1$, for $\mathfrak{s} \in \mathcal{K}, \mathfrak{a} \in \mathcal{A}, \mathfrak{i} \in \{0,1,...,\mathfrak{I}-1\}, \mathfrak{d} \in \{0,1,...,\mathfrak{D}-1\}$. The additional degrees of freedom introduced by these newly lifted decision variables are removed by adding a circulation-fixing constraint to the discrete-form OCP:

$$\boldsymbol{G}_{\text{LL},\Gamma} = \begin{pmatrix} \overline{\Gamma}_{\mathfrak{s},0_{\mathfrak{i},\mathfrak{d}}} - \frac{\int_{t_{\text{shed}}\left(1,t_{\mathfrak{i},\mathfrak{d}}\right)}^{t_{\text{shed}}\left(0,t_{\mathfrak{i},\mathfrak{d}}\right)} \Gamma_{\text{k},\mathfrak{s}}(\phi)\text{d}\phi}{\int_{t_{\text{shed}}\left(1,t_{\mathfrak{i},\mathfrak{d}}\right)}^{t_{\text{shed}}\left(0,t_{\mathfrak{i},\mathfrak{d}}\right)} \text{d}\phi}, & \text{for } \mathfrak{s} \in \mathcal{K}, \mathfrak{i} \in \{0,1,...,\mathfrak{I}-1\}, \mathfrak{d} \in \{0,1,...,\mathfrak{D}-1\} \\ \overline{\Gamma}_{\mathfrak{s},\mathfrak{a}_{\mathfrak{i},\mathfrak{d}}} - \overline{\Gamma}_{\mathfrak{s},0_{\mathfrak{i}-\text{mod}(\mathfrak{a},\mathfrak{I}),\mathfrak{d}}}, & \text{for } \mathfrak{s} \in \mathcal{K}, \mathfrak{a} \in \{1,2,...,\mathfrak{A}-1\}, \mathfrak{i} \in \{0,1,...,\mathfrak{I}-1\}, \\ & \mathfrak{d} \in \{0,1,...,\mathfrak{D}-1\} \end{pmatrix} = \boldsymbol{0}, \tag{15}$$

where $\Gamma_{\text{k},\mathfrak{s}}(t_{\text{shed}}(\mathfrak{a},t_{\mathfrak{i},\mathfrak{d}}))$ is the instantaneous circulation on the $\mathfrak{s}^{\text{th}}$ kite at the time when the $\mathfrak{a}^{\text{th}}$ wake node was shed, as determined with (14), and $\overline{\Gamma}_{\mathfrak{s},0_{\mathfrak{i}-\text{mod}(\mathfrak{a},\mathfrak{I}),\mathfrak{d}}}$ refers back to (7) as the average circulation on the $\mathfrak{s}^{\text{th}}$ kite during the control interval

[11] Flying "a 24-m-radius flight path... at 30m/s" implies a roughly 151 m-circumference flight path, traversed in roughly 5 seconds-per-winding. Since "[t]he simulation lasts for 20s, with 200 equally spaced time steps", this implies 0.1 seconds between shedding instance, a.k.a. 50 shedding instances per winding.(Rancourt and Mavris, 2018, p. 9)

[12] Experience shows that using the typical 'Forwards Euler'-style method of representing the strength of a bound vortex (and its following vortex ring) with the circulation experienced by the lifting surface at the exact instant of shedding, enables the optimizer to choose strange flight patterns that happen to result in beneficial circulation values at only the control nodes.

[13] Only 'very little' as opposed to 'not at all', because the lengths of the control intervals may be deliberately designed on a nonuniform temporal-grid, as occurs when using a phase-fixing strategy.

[14] To 'lift', here, refers to the reformulation strategy of representing a term of the larger model as an algebraic variable among the OCP's decision variables - with an additional equality constraint of the same dimension introduced simultaneously to ensure that the solution has not actually gained any additional degrees of freedom - in order to trade an increased problem size for 'less-nonlinear' constraints that will require fewer Newton-iterations during root-finding.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

when the $\mathfrak{a}^{\text{th}}$ wake node was shed. The above integral has to be computed with whatever numerical integration approach is being used to discretize the OCP (here, Radau-IIA direct collocation).

In the form given in (15), the assumption of periodicity is used to find the average circulation for even those wake nodes shed 'prior' to the start of the optimization period. If the trajectory is not assumed to be periodic, some other method has to be used to determine the average circulation values for $\mathfrak{a} > \mathfrak{i}$. Luckily, once the vortex elements have been shed, the inviscid assumption means that the associated average circulation values do not need to be adjusted. This means that we do not need to include the computation of the circulation's integral $\mathfrak{AKID}$ times in the OCP constraints, but we can instead include the integral in (15) only $\mathfrak{KID}$ times and rely on linear equality constraints to fix the remaining average circulation values.

2.6 Induced velocities

Next, we need to include the induction predicted at each kite into the OCP, by including the Biot-Savart relationship into the constraints. The difficulty, when using Newton-type solvers, is that analytical solutions of the Biot-Savart integral necessarily have very nonlinear behavior[15] when the observer is close to the vortex core, while the region far from the core has asymptotic behavior that is much more linear. This nonlinearity imposes the danger of producing a numerical instability as the optimizer pushes one or another observation point into or through the Biot-Savart's nonlinearity. To avoid this, we will adjust our solution process to introduce some separation between the induced velocity applied in (12) and the sum of the Biot-Savart contributions. We specifically want to create this separation when we're early on the homotopy path, and the simplified circular-trajectory initial guess leads to poor estimates of the constraint multipliers and rapid changes to the decision variables. Therefore, we'll introduce a new homotopy predictor-corrector step into the sequence of predictor-corrector steps described in De Schutter et al. (2023): specifically, rather than applying the total superposed induced velocities at the kite's mid-span (12) already from the beginning of the homotopy path, we will instead wait to introduce that total superposed induced velocity into the calculation of the aerodynamic forces on the kite during a new homotopy step after the fictitious-forces $\boldsymbol{F}_{\text{fict}}$ have been eliminated but before[16] the trajectory leaves the circular flight path within the homotopy's power step. In this new homotopy step, we'll require that the induced velocity of the $\mathfrak{k}^{\text{th}}$ observer kite $\boldsymbol{v}_{\text{ind},\mathfrak{k}}$ that is applied in the effective velocity (12) and determines the aerodynamic forces in (11), must be equal to the total induced velocity $\boldsymbol{v}_{\text{ind},\sum,\mathfrak{k}}$ calculated from the superposition of the induced velocities of all of the vortex elements.

To do this, we'll lift both the applied induced velocity ($\boldsymbol{v}_{\text{ind},\mathfrak{k}}(t) \in \mathbb{R}^3$ for $\mathfrak{k} \in \mathcal{K}$) and the total calculated induced velocity ($\boldsymbol{v}_{\text{ind},\sum,\mathfrak{k}}(t) \in \mathbb{R}^3$ for $\mathfrak{k} \in \mathcal{K}$) into the (continuous-form) problem as algebraic decision variables, as well as introduce a new induction homotopy parameter $\varsigma_{\text{ind}} \in \mathbb{R}$ that will be driven from 1 to 0 in a predictor-corrector step.

These lifted variables are related by adding the following induction-homotopy residual function to the dynamics:

$$\boldsymbol{r}_{\text{ind,hmtp},\mathfrak{k}} = (\boldsymbol{v}_{\text{ind},\mathfrak{k}} - \boldsymbol{0})\varsigma_{\text{ind}} + \left(\boldsymbol{v}_{\text{ind},\mathfrak{k}} - \boldsymbol{v}_{\text{ind},\sum,\mathfrak{k}}\right)(1 - \varsigma_{\text{ind}}) = \boldsymbol{0}, \tag{16}$$

[15] in the non-regularized forms, even a discontinuity

[16] This specific ordering is carried over from the inclusion of actuator models into the `awebox`, since actuator models require the determination of geometric properties like a radius and a normal-vector.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

with the above-mentioned total calculated induced velocity $\boldsymbol{v}_{\mathrm{ind},\sum,\mathfrak{k}}$ constrained by a superposition residual[17] that equates each $\mathfrak{k}^{\mathrm{th}}$ total induced velocity value to the sum of the Biot-Savart contributions from all of the vortex elements:

$$\boldsymbol{r}_{\mathrm{ind,sprptn},\mathfrak{k}} = \boldsymbol{v}_{\mathrm{ind},\sum,\mathfrak{k}}(t) - \sum_{\mathfrak{e}}^{\mathcal{E}_{\mathrm{Ffil}}} \boldsymbol{v}_{\mathrm{ind:Ffil}}\left(\boldsymbol{q}_{\mathfrak{k}}(t); \Gamma_{\mathfrak{e}}(t), \boldsymbol{q}_{\mathrm{start},\mathfrak{e}}(t), \boldsymbol{q}_{\mathrm{end},\mathfrak{e}}(t), R_{\mathrm{c}}\right) \ldots \tag{17}$$

$$\ldots - \sum_{\mathfrak{e}}^{\mathcal{E}_{\mathrm{SIfil}}} \boldsymbol{v}_{\mathrm{ind:SIfil}}\left(\boldsymbol{q}_{\mathfrak{k}}(t); \Gamma_{\mathfrak{e}}(t), \boldsymbol{q}_{\mathrm{start},\mathfrak{e}}(t), \hat{\mathbf{x}}, R_{\mathrm{c}}\right) = \mathbf{0},$$

$$\mathcal{E}_{\mathrm{Ffil}} = \{0, 1, \ldots, \mathfrak{E} - (2\mathfrak{K} + 1)\}, \qquad \mathcal{E}_{\mathrm{SIfil}} = \{\mathfrak{E} - 2\mathfrak{K}, \ldots, \mathfrak{E} - 1\}, \tag{18}$$

where $\mathcal{E}_{\mathrm{Ffil}}$ and $\mathcal{E}_{\mathrm{SIfil}}$ are the set of all, respectively, finite- and semi-infinite- filament elements modeled at that instance in time; $\boldsymbol{q}_{\mathfrak{k}}$ is the position of the $\mathfrak{k}^{\mathrm{th}}$ kite's center of mass, where we observe the induced velocity given the lack of any definite information about the locations of the kite's Pitot tube and aerodynamic center; and $\mathfrak{e}$ is the index of one of the various vortex elements in the set of all vortex elements $\mathcal{E}$. Meanwhile, $\Gamma_{\mathfrak{e}}(t), \boldsymbol{q}_{\mathrm{start},\mathfrak{e}}(t), \boldsymbol{q}_{\mathrm{end},\mathfrak{e}}(t), R_{\mathrm{c}}$ are the $\mathfrak{e}^{\mathrm{th}}$ element's strength, starting position, ending position, and core radius, as listed in Table A1. These values are the inputs of the simplest common (see (Branlard, 2017)) regularization (the 'cut-off' model) of the analytical solutions to the Biot-Savart integral for a finite-length vortex filament $\boldsymbol{v}_{\mathrm{ind:Ffil}}$ and for a semi-infinite[18] vortex filament $\boldsymbol{v}_{\mathrm{ind:SIfil}}$:

$$\boldsymbol{v}_{\mathrm{ind:Ffil}}(\boldsymbol{q}; \Gamma, \boldsymbol{q}_0, \boldsymbol{q}_1, R_{\mathrm{c}}) = \frac{\Gamma}{4\pi} \cdots \tag{19}$$

$$\left(\frac{\left(\|\boldsymbol{q}_0 - \boldsymbol{q}\|_2 + \|\boldsymbol{q}_1 - \boldsymbol{q}\|_2\right)\left((\boldsymbol{q}_0 - \boldsymbol{q}) \times (\boldsymbol{q}_1 - \boldsymbol{q})\right)}{\left(\|\boldsymbol{q}_0 - \boldsymbol{q}\|_2^2 \|\boldsymbol{q}_1 - \boldsymbol{q}\|_2^2 + \|\boldsymbol{q}_0 - \boldsymbol{q}\|_2 \|\boldsymbol{q}_1 - \boldsymbol{q}\|_2 (\boldsymbol{q}_0 - \boldsymbol{q}) \cdot (\boldsymbol{q}_1 - \boldsymbol{q})\right) + R_{\mathrm{c}}^2 \|\boldsymbol{q}_1 - \boldsymbol{q}_0\|_2^2} \right),$$

$$\boldsymbol{v}_{\mathrm{ind:SIfil}}(\boldsymbol{q}; \Gamma, \boldsymbol{q}_0, \hat{\boldsymbol{\ell}}, R_{\mathrm{c}}) = \frac{\Gamma}{4\pi} \frac{(\boldsymbol{q}_0 - \boldsymbol{q}) \times \hat{\boldsymbol{\ell}}}{\|\boldsymbol{q}_0 - \boldsymbol{q}\|_2^2 + \|\boldsymbol{q}_0 - \boldsymbol{q}\|_2 \left(\hat{\boldsymbol{\ell}}^{\top} (\boldsymbol{q}_0 - \boldsymbol{q})\right) + R_{\mathrm{c}}^2} \tag{20}$$

where the core radius R_{c} is a positive, real value with units of meters. The term containing the core radius is negligible unless the observer $\boldsymbol{q}$ is roughly on the filament.

Since we are here already relying on an assumption of a rigidly-convected wake and using relatively large vortex elements, we choose the vortex core radius in relation to the generating airfoil's thickness (in the style of Leishman (2006) and Sant et al. (2006)) rather than according to the size of the vortex elements (as discovered by Beale and Majda (1982) and recommended by Branlard (2017, Ch. 41)). Therefore, we'll assume that the vortex core radius can be meaningfully fixed at 5 percent of the kite mean aerodynamic chord c_{MA}, as suggested by van der Wall and Richard (2006) and Mula et al. (2013). Beyond this, an assumption of externally potential flow removes any viscous core growth, so that the core radius can remain constant even as the vortex elements convect.

[17] Notice that attempting to lift or rearrange the various components of (19) within (17) will tend to make that constraint more non-linear rather than less.

[18] We find this cut-off regularized semi-infinite-filament analogue to the well known finite-filament expression of (19) by taking the limit $\lim_{\phi \to \infty} \left(\boldsymbol{v}_{\mathrm{ind:Ffil}}(\boldsymbol{q}; \Gamma, \boldsymbol{q}_0, \boldsymbol{q}_0 + \phi\hat{\boldsymbol{\ell}}, 0)\right)$ and re-introducing an additive, unit-appropriate smoothing term based on the core radius back into the denominator.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

The expression for the (continuous-form) dynamics that constrains the position of the 0^{th} wake nodes, the computed induced velocities, and the applied induced velocities, now reads as:

$$\boldsymbol{d}_{\text{LL}} = \begin{pmatrix} \boldsymbol{r}_{\text{tip}}\left(\boldsymbol{x}\left(t\right),\boldsymbol{y}\left(t\right)\right) & & \text{from (9)} \\ \boldsymbol{r}_{\text{ind,sprptn,}\mathfrak{k}}\left(\boldsymbol{x}\left(t\right),\boldsymbol{y}\left(t\right)\right) & \text{for } \mathfrak{k} \in \mathcal{K} & \text{from (17)} \\ \boldsymbol{r}_{\text{ind,hmtp,}\mathfrak{k}}\left(\boldsymbol{y}\left(t\right),\varsigma_{\text{ind}}\right) & \text{for } \mathfrak{k} \in \mathcal{K} & \text{from (16)} \end{pmatrix} = \boldsymbol{0}. \tag{21}$$

These RLL dynamics - which allow a relaxed relationship between the induced velocities on the kites and the vortex element positions and strengths until the induction homotopy parameter reaches its final, Biot-Savart-enforcing value of $\varsigma_{\text{ind}} = 0$ - can be discretized according to the process described in Section 2.1 to give a set of discrete equality constraints $\boldsymbol{G}_{\text{LL,dyn}}$.

2.7 The AWE OCP with a rigidly-convected, lifting-line vortex model

In addition to the discrete-form baseline AWE OCP decision variables $\boldsymbol{W}_{\text{bsln}}$, we have now introduced several lifted (algebraic) wake variables which we will group together under the symbol $\boldsymbol{W}_{\text{wake}}$: the wake node positions $\boldsymbol{q}_{\text{wake},\mathfrak{k},\mathfrak{t},\mathfrak{a}_{\mathfrak{i},\mathfrak{d}}}$, the average circulations $\overline{\Gamma}_{\mathfrak{k},\mathfrak{a}_{\mathfrak{i},\mathfrak{d}}}$, the computed induced velocities $\boldsymbol{v}_{\text{ind},\sum,\mathfrak{k}_{\mathfrak{i},\mathfrak{d}}}$, and the applied induced velocities $\boldsymbol{v}_{\text{ind},\mathfrak{k}_{\mathfrak{i},\mathfrak{d}}}$, instantiated according to the subscripts for $\mathfrak{i} \in \{0,1,...,\mathfrak{I}-1\}$, $\mathfrak{d} \in \{0,1,...,\mathfrak{D}-1\}$, $\mathfrak{k} \in \mathcal{K}$, $\mathfrak{t} \in \mathcal{T}$, and $\mathfrak{a} \in \{0,1,...,\mathfrak{A}-1\}$. When concatenated, the above combination of the discrete variables - $\boldsymbol{W}_{\text{bsln}}$ and $\boldsymbol{W}_{\text{wake}}$, as well as the induction homotopy parameter ς_{ind} - together give a set of decision variables $\boldsymbol{W}_{\text{RLL}}$ for an AWE OCP with a rigidly-convected, lifting-line vortex model (RLL AWE OCP). Notice that the number of wake-model-related variables and constraints is substantially higher in this optimal control formulation than in a typical forward-simulation approach because the wake's state at every discretized time instance is simultaneously represented- and solved- using one monolithic body of equations.

We can construct this RLL AWE OCP by starting with the discrete-form baseline AWE OCP of (6) and appending all of the constraints that we've developed to uniquely determine the discrete lifted wake variables.

$$\begin{aligned} \boldsymbol{W}_{\text{RLL}}{}^{\star} = \quad & \underset{\boldsymbol{W}_{\text{RLL}}}{\text{arg. min.}} && \varsigma_{\text{P}}\Psi_{\text{track}} + \left(1-\varsigma_{\text{P}}\right)\left(-k_{\text{P}}\overline{P}\right) + \Psi_{\text{reg}} + \Psi_{\text{hmtp}} && && \text{from (6)} \\ & \text{subj. to} && \boldsymbol{G}_{\text{bsln}}(\boldsymbol{W}_{\text{bsln}}) && = \ \boldsymbol{0} && \text{from (6)} \\ & && \boldsymbol{H}_{\text{bsln}}(\boldsymbol{W}_{\text{bsln}}) && \leq \ \boldsymbol{0} && \text{from (6)} \\ & && \boldsymbol{G}_{\text{LL,dyn}}(\boldsymbol{W}_{\text{RLL}}) && = \ \boldsymbol{0} && \text{from (21)} \\ & && \boldsymbol{G}_{\text{RLL},\boldsymbol{q}}\left(\boldsymbol{W}_{\text{wake}},\boldsymbol{\theta}\right) && = \ \boldsymbol{0} && \text{from (10)} \\ & && \boldsymbol{G}_{\text{LL},\Gamma}\left(\boldsymbol{W}_{\text{RLL}}\right) && = \ \boldsymbol{0}, && \text{from (15)} \end{aligned} \tag{22}$$

where the concatenation of $\boldsymbol{G}_{\text{LL,dyn}}$, $\boldsymbol{G}_{\text{RLL},\boldsymbol{q}}$, and $\boldsymbol{G}_{\text{LL},\Gamma}$ has an invertible Jacobian with respect to the discrete lifted wake variables $\boldsymbol{W}_{\text{wake}}$.

As already mentioned, the formulation given above relies on an assumption of periodicity to build the wake's history prior to the initial time $t_0 = 0\,\text{s}$. But, the RLL AWE OCP can be adjusted fairly straightforwardly for non-periodic trajectories by correcting (7) and (15) to reflect the desired system history.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

In all of the following problems, the AWE OCPs are formulated in the `CasADi`-based(Andersson et al., 2019) `awebox`, using Radau-IIA direct collocation with $\mathfrak{D} = 4$ collocation nodes per control interval[19] for discretization and the `awebox` default options, unless otherwise noted. These default inputs - including the selection of an Ampyx AP2 kite - were chosen (De Schutter et al., 2023) after considering many combined factors, including the availability of modeling data, the presence of independently-generated AWE OCP solutions for comparison, and experience building problems that are numerically tractable. After formulation, the problems are solved using the interior point solver `IPOPT` (Wächter and Biegler, 2006), exact-Hessian Newton-steps[20], the linear-algebra pre-conditioner `Metis` (Karypis and Kumar, 1997) and the highly-parallelizable but, unfortunately, not necessarily repeatable[21] linear solver `MA86` (SCD Computational Mathematics Theme, 2024)[22].

2.8 Modeled wake resolution

Having now defined our RLL model and the RLL AWE OCP, the next logical steps are to demonstrate that the implemented model makes reasonable predictions, to explore the tradeoffs of including the RLL model into the AWE OCP, and to use a test problem to determine whether there's actually any practical value gained by solving the RLL AWE OCP over the baseline AWE OCP. But before we move on to those questions, we first need to briefly consolidate the definitions we've already seen about the wake resolution.

The resolution of the RLL model in the RLL AWE OCP can be adjusted by tuning two values: the number of control intervals dividing the optimization period $\mathfrak{I}$ and the number of periods' worth of wake being tracked $\mathfrak{P}$, as defined in (8). Increasing the number of periods tracked will lengthen the amount of the wake that's resolved, whereas increasing the number of control intervals will make the resolution within the resolved-region finer. The impact of these two parameters is shown by presenting the solution to (22) with two different resolutions in Figure 4: the solution $\mathbf{A}_3$, which is shown above and boxed in orange, has a small $\mathfrak{I}$ and a long $\mathfrak{P}$, compared to the solution $\mathbf{A}_1$, shown below and boxed in black, with a larger $\mathfrak{I}$ and a short $\mathfrak{P}$.

[19] The integration error in a Radau-IIA collocation scheme is proportional to the control interval timestep to the $(2\mathfrak{D}-1)^{\text{th}}$ power (Biegler, 2010). So, reducing the number of collocation nodes risks giving the optimizer an incentive to artificially improve its objective by increasing the integration error. In addition, a low collocation order also means that the number of control intervals $\mathfrak{I}$ needs to be increased in order to sufficiently enforce the inequality constraints. And, unfortunately, increasing $\mathfrak{I}$ means increasing the number of wake nodes needed to resolve the same length of wake. The result is that: experience suggests that the superficial appeal of decreasing $\mathfrak{D}$ to 2 (or below) does not ultimately lead to reducing computational expense.

[20] Attempts at reducing the necessary memory using the L-BFGS Hessian approximation, as available via `IPOPT` (Kawajir et al.) and suggested by van den Broek et al. (2024) (whose inequality-constraint-free problem doesn't have complementarity conditions that need to be handled by, e.g., a barrier-parameter-tightening homotopy), tend to get stuck in `IPOPT`'s barrier-parameter-tightening homotopy ($\varsigma_{\text{comp}} \to 0$) for the RLL AWE OCP. This may be related to the fact that the tuning of the problem is still imperfect (see the non-constant relationship between iterations and problem-size in the exact-Newton solution, in Figure 11, bottom-left) and L-BFGS is known to perform badly on ill-conditioned problems "where the Hessian matrix contains a wide distribution of eigenvalues" (Nocedal and Wright, 2000).

[21] While `Metis` (Karypis, 2013) can use a random-seed during sparse-matrix reordering, the 'non-repeatability' mentioned here is most-likely the result of the heuristics and solution techniques used by `IPOPT` (automatic re-scaling and re-regularization) and the `Coin-HSL` solvers (automatic pivot selection, see Hogg and Scott (2013)) to allow solution of even ill-conditioned problems. The result is that floating-point differences can lead different paths of Newton iterates to possibly-different local minima, a fact that becomes especially relevant when the linear algebra is broken up for parallel solution.

[22] Practical experience finds that the non-repeatable `MA86` still performs better than the guaranteed-to-be-repeatable-on-equivalent-linear-systems-with-equivalent-computational-resources `Coin-HSL` solvers `MA57` (requires more clock-time) and `MA97` (frequently gets caught in infeasible pockets during the initial homotopy step).

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

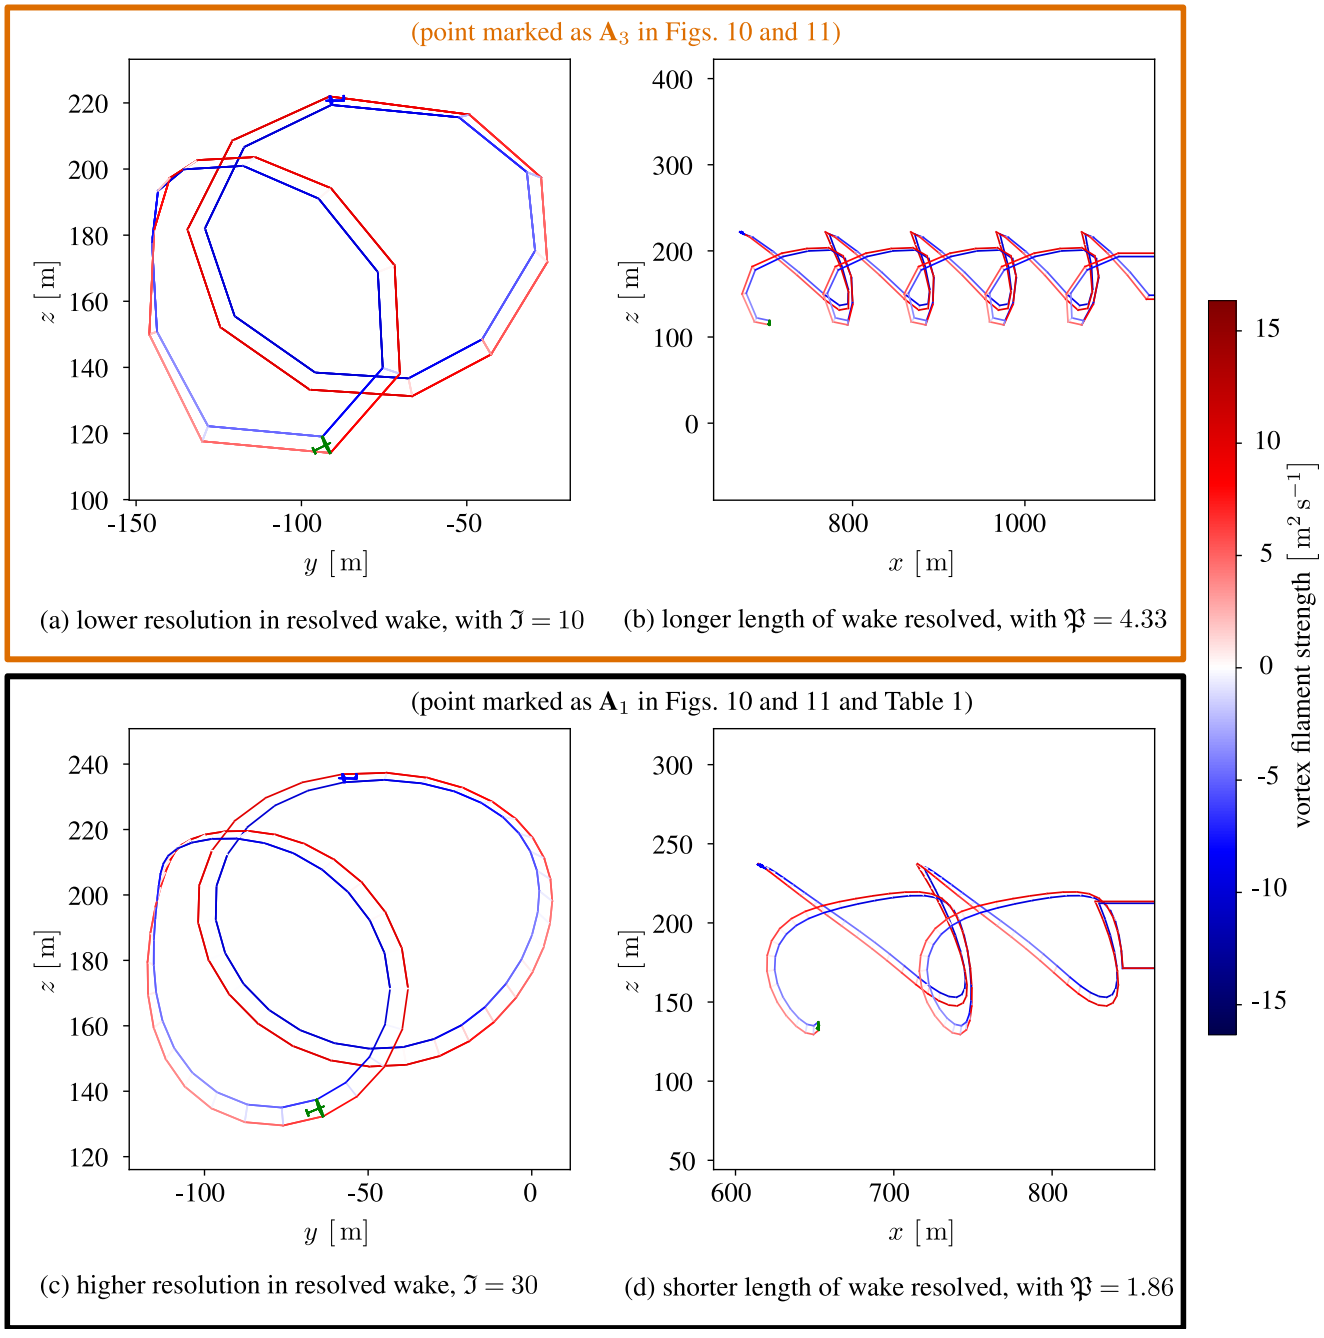


Figure 4. The vorticity distribution found in the solution of a two-kite RLL AWE OCP: top-subfigures (a and b) with $\mathfrak{I} = 10$ control intervals and $\mathfrak{P} = 4.33$ periods tracked; bottom-subfigures (c and d) with $\mathfrak{I} = 30$ control intervals and $\mathfrak{P} = 1.86$ periods tracked. Each of these snapshots is 'taken' when the kites (blue and green) have just finished the reel-in (low lift) phase and are about to begin reeling-out again (at high lift, to harvest power). This can be seen by the comparatively weaker vorticity (pale red-and-blue filaments) behind the kite's shown locations, as compared to the stronger vorticity (darker red-and-blue filaments) in the remainder of the shed vorticity.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

Many of the following figures within this paper will show solutions of the RLL AWE OCP for various combinations of control intervals $\mathfrak{I}$ and periods tracked $\mathfrak{P}$. To allow trends to be identified in these plots, the constant-$\mathfrak{I}$ series are identified by color, with lighter colors corresponding to coarser resolutions and smaller values of $\mathfrak{I}$, and darker colors corresponding to finer resolutions and larger values of $\mathfrak{I}$. Where knowledge of the length of the tracked-wake ($\mathfrak{P}$) is necessary to understanding the plot but not otherwise visible, the $\mathfrak{P}$ values are labeled next to their respective data-points.

3 Verification of the wake model against higher-fidelity simulations

In equation (4), we described an objective function that can be shifted via the power homotopy parameter ς_P from a tracking problem to a power problem. The advantage of the homotopy approach is, most generally, that it can control the solver's path from uninformed initial guess to final solution, rendering the problem more tractable. But, this homotopy approach also has the advantage that we can decide to terminate the homotopy path early, for example at the solution of the tracking problem. This control over the homotopy-path, effectively allows us to simulate the models included in the OCP constraints exactly as they will behave when formulated in the OCP and solved with the chosen OCP solvers. This means that it is fairly simple to modify the RLL AWE OCP (22), so that we can use it as a simulator of the rigidly-convected, lifting-line vortex model.

We create this RLL Simulation OCP by: following only the homotopy steps that drive the problem towards the desired goal - in the case of simulation, this means progressing through the induction- and interior point- homotopy steps (sending parameters ς_{ind} and ς_{comp} to zero) and skipping the homotopy steps to eliminate the fictitious forces and switch into the power problem (parameters $\varsigma_{\text{fict}} = \varsigma_{\text{P}} = 1$); disabling as many of the inequality constraints as possible, to reduce expense and ensure feasibility; choosing regularization weights that penalize deviation from the reference in relation to the importance and scaling of that variable, without damaging the problem's second order sufficiency; applying whatever trajectory we intend to simulate as the initial guess and reference of what is now a pure tracking problem; and finally, using whatever custom determination of aerodynamic coefficients is needed to produce in (11) the specified aerodynamic forces.

After these modifications, the RLL Simulation OCP reads as:

$$
\begin{aligned}
\boldsymbol{W}^{\star}_{\text{sim}} = \quad & \underset{\boldsymbol{W}_{\text{RLL}}}{\text{arg. min.}} && \Psi_{\text{track}} + \Psi_{\text{reg}} + \Psi_{\text{hmtp}} && && \text{from (6) \& (24)} \\
& \text{subj. to} && \boldsymbol{G}_{\text{bsln}}(\boldsymbol{W}_{\text{bsln}}) && = \quad \boldsymbol{0} && \text{from (6) \& (24)} \\
& && \boldsymbol{H}_{\text{sim}}(\boldsymbol{W}_{\text{bsln}}) && \leq \quad \boldsymbol{0} && \text{from (24)} \\
& && \boldsymbol{G}_{\text{LL,dyn}}(\boldsymbol{W}_{\text{RLL}}) && = \quad \boldsymbol{0} && \text{from (21)} \\
& && \boldsymbol{G}_{\text{RLL},\boldsymbol{q}}(\boldsymbol{W}_{\text{wake}}, \boldsymbol{\theta}) && = \quad \boldsymbol{0} && \text{from (10)} \\
& && \boldsymbol{G}_{\text{LL},\Gamma}(\boldsymbol{W}_{\text{RLL}}) && = \quad \boldsymbol{0}, && \text{from (15)}
\end{aligned}
\tag{23}
$$

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

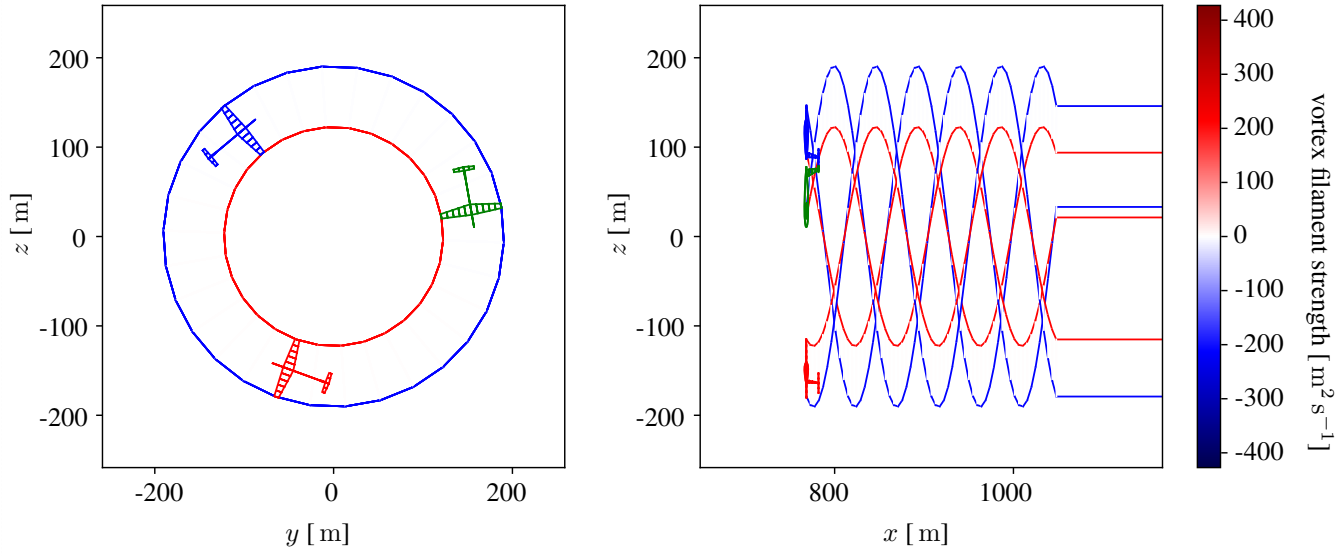


Figure 5. The vorticity distribution predicted with the `awebox`, when simulating with the RLL model for $\mathfrak{I} = 30$ and $\mathfrak{P} = 2$, under the steady-flow situation described by Haas and Meyers (2017). Left shows a back view into the wind; right shows a side view perpendicular to the wind.

using an extremely reduced 'simulation and tracking version' of the path inequalities $\boldsymbol{H}_{\text{sim}}$ that serves to avoid divide-by-zero errors in the time derivatives. The parameter equalities $\boldsymbol{g}_{\boldsymbol{\theta}}$ (3) that enter into (6)'s $\boldsymbol{G}_{\text{bsln}}$ are also adjusted to the below $\boldsymbol{g}_{\boldsymbol{\theta},\text{sim}}$:

$$\boldsymbol{H}_{\text{sim}}(\boldsymbol{W}_{\text{RLL}}) = \begin{pmatrix} 0 - t_{\text{switch}} \\ 0 - (t_{\text{final}} - t_{\text{switch}}) \end{pmatrix} \leq \boldsymbol{0}, \qquad \boldsymbol{g}_{\boldsymbol{\theta},\text{sim}}(\boldsymbol{\theta}, \varsigma) = \begin{pmatrix} \boldsymbol{\theta} - \boldsymbol{\theta}_{\text{Ref}} \\ \varsigma - \varsigma_{\text{Ref}} \end{pmatrix} = \boldsymbol{0}, \tag{24}$$

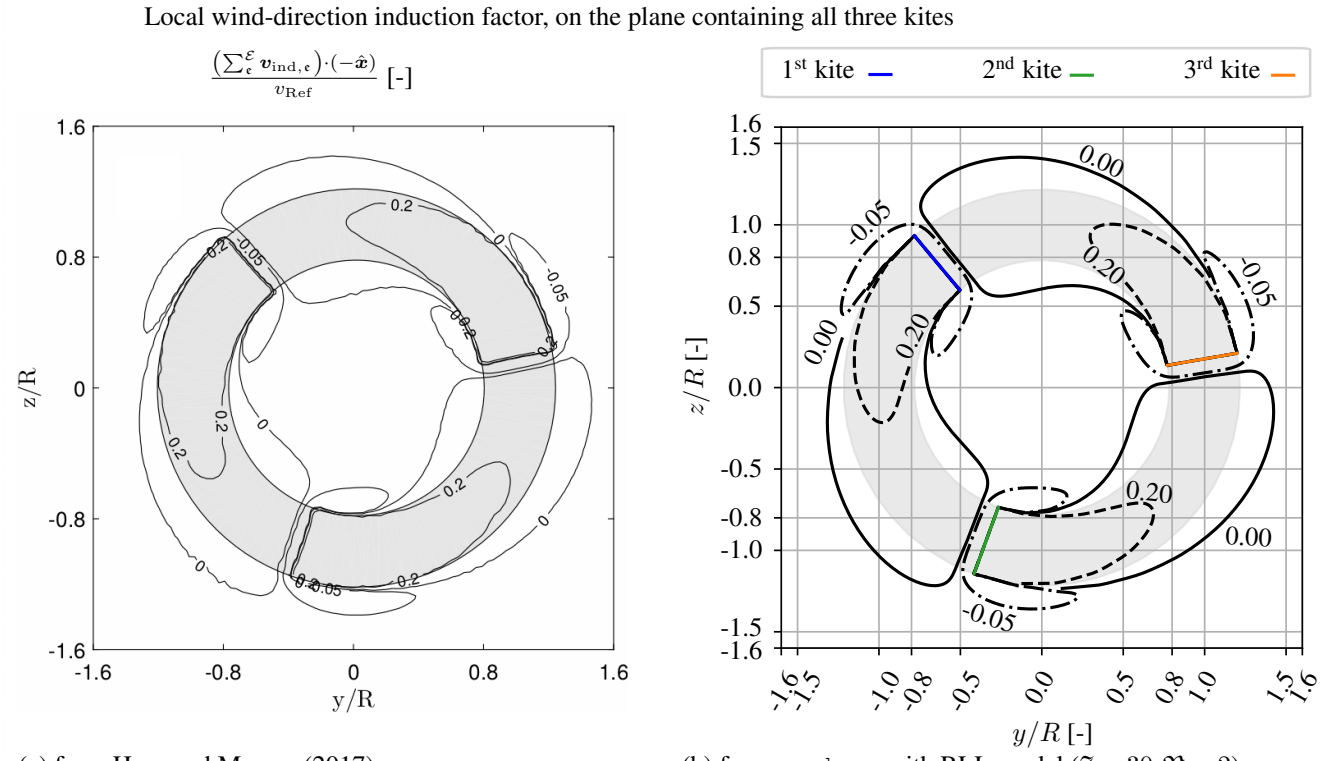
where the reference homotopy parameters are chosen for simulation to be: $\varsigma_{\text{ind Ref}} = \varsigma_{\text{comp}_{\text{Ref}}} = 0$, and $\varsigma_{\text{fict Ref}} = \varsigma_{\text{P Ref}} = 1$.

There are some significant downsides to using a Simulation OCP to verify the wake model - substantially higher memory cost, decreased flexibility in the situations that can be simulated, and a greater reliance on the correct tuning of parameters - as compared to using a 'normal' forward-simulation method. But, there is also one notable advantage: by using the existing OCP structure to perform our model verification, we end up with a way to test not only the predictions of the model on its own, but also its implementation within the OCP formulation.

3.1 Induction verification under rotationally-steady conditions

The work of Haas and Meyers (2017) includes a Large Eddy Simulation (LES) prediction of induction factors over the plane of the kites' spans, for a steady-inflow scenario, of three kites flying an axisymmetric, constant-radius, constant-force trajectory. This case is an ideal verification case because it is simplified enough that a comparison should be able to meaningfully check the predictions of a wake model under steady conditions, without being caught up too much by the details of the trajectory.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



(a) from Haas and Meyers (2017)

(b) from `awebox`, with RLL model ($\mathfrak{I} = 30, \mathfrak{P} = 2$)

Figure 6. The downstream-direction projection of the induced velocity, non-dimensionalized by the reference wind speed (a.k.a. the axial induction factor), over the plane of the kites' spans, under the steady-flow situation described by (Haas and Meyers, 2017). The contours describe the relative levels of induction factor ($a = -0.05, 0$, and 0.2), with the swept annulus shown in gray. Left (a) gives the levels published by Haas and Meyers (2017), included with permission. Right (b) gives the levels predicted by the `awebox`, when solving the simulation OCP (23) with the RLL model, for $\mathfrak{I} = 30$ and $\mathfrak{P} = 2$.

The LES in Haas and Meyers (2017), specifically integrates its partial differential equation with a pseudo-spectral method in the horizontal directions $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$, an energy-conserving finite-difference method in the vertical direction $\hat{\mathbf{z}}$, and a fourth-order Runge-Kutta method with a time-step chosen to give a Courant-Friedrichs-Lewy number of 0.4 in the temporal direction t. The LES grid cells have dimensions $5.7\,\mathrm{m} \times 5.7\,\mathrm{m} \times 2.85\,\mathrm{m}$, which Haas and Meyers (2017) claim gives sufficient resolution to fully capture the shedding of the wing-tip vortices.

In the specific situation described by Haas and Meyers (2017), the free-stream wind velocity is $\boldsymbol{v}_\infty = v_{\mathrm{Ref}}\hat{\mathbf{x}} = 10\,\mathrm{m\,s^{-1}}\hat{\mathbf{x}}$, and the kites are assumed to be flying high enough that the atmosphere, including wind velocity profile, is uniform. The air density is assumed to be the standard, dry, sea-level value of $\rho_{\mathrm{air}} = 1.225\,\mathrm{kg\,m^{-3}}$. When solving the RLL Simulation OCP in the `awebox`, we'll additionally ignore the asymmetries of the modeled universe - gravity, the presence of the ground, as well as those asymmetries already neglected such as atmospheric baroclinity - and the tether drag.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

Each kite in this simulation is modeled as a Lockheed C-5a Galaxy, with a wingspan $b = 68\,\mathrm{m}$ (Haas and Meyers, 2017), and a wing area of 580 m^2 (Wikipedia contributors, 2024). This gives a mean aerodynamic chord of 8.5294 m.

The three kites are evenly spaced in a steady, circular trajectory. The kites are assumed to be oriented with their span-wise axis $\hat{e}_{2,\mathrm{k},\mathfrak{k}}$ in the circle's radial coordinate, and their up axis $\hat{e}_{3,\mathrm{k},\mathfrak{k}}$ pointing along the wind axis. In Haas and Meyers (2017), the rotation of the trajectory is about the $-\hat{\mathbf{x}}$ axis; in `awebox` parlance, the trajectory is 'counter-clockwise' when looking along a positive $\hat{\mathbf{x}}$. The first kite is located at an angle of 40° from the $\hat{\mathbf{z}}$ axis in the direction of the $-\hat{\mathbf{y}}$ axis.

In this verification case, we will use fictitious forces to allow the kites to track a circular trajectory with winding number $\mathfrak{L} = 1$. If there is reel-out motion, it is parallel to the freestream wind velocity, so that it can be subsumed into the freestream wind velocity term in the following calculations. This circular trajectory has a fixed radius (at the kite mid-span position) R of 155.77 m. At the kites' exterior wingtips, the radial distance to the axis of rotation is $R_{\mathrm{ext}} = R + \frac{1}{2}b = 189.77\,\mathrm{m}$; at the interior wingtips, it is $R_{\mathrm{int}} = R - \frac{1}{2}b = 121.77\,\mathrm{m}$. The tracked trajectory has a uniform rotational velocity, with one full rotation requiring a period of $t_{\mathrm{final}} = 13.98\,\mathrm{s}$. This corresponds to a kite speed ratio $\Lambda = \frac{\Omega R}{v_0} = \frac{\boldsymbol{v}_{\mathrm{app}}\cdot(-\hat{\mathbf{t}})}{\boldsymbol{v}_{\mathrm{app}}\cdot\hat{\mathbf{x}}}$ of 7 and an angular speed $\Omega = \Lambda v_\infty / R = 0.45\,\mathrm{rad\,s^{-1}}$.

The associated thesis by Haas (2022) gives the lift-per-unit-span L' on each kite as:

$$L'(r) = \rho_{\mathrm{air}} \frac{4\pi}{3} \frac{{v_0}^2}{\Omega} \overline{a}\,(1-\overline{a})\,\|\boldsymbol{v}_{\mathrm{eff}}(r)\|_2\,, \text{ where } \|\boldsymbol{v}_{\mathrm{eff}}(r)\|_2 = \sqrt{\left((1-\overline{a})\,v_0\right)^2 + \left((1+\overline{a}')\,\omega r\right)^2}. \tag{25}$$

From Haas and Meyers (2017), we have the azimuthally-averaged axial- and tangential-direction induction factors (respectively, $\overline{a}$ and $\overline{a}'$); from Burton et al. (2011), we have the definition of the flow angle φ_{flow}:

$$\overline{a} = \frac{1}{3}, \quad \overline{a}'(r) = \frac{\overline{a}(1-\overline{a})}{\Lambda^2 \left(\frac{r}{R}\right)^2}, \quad \varphi_{\mathrm{flow}}(r) = \tan^{-1}\left(\frac{v_0\,(1-\overline{a})}{\Omega r\,(1+\overline{a}')}\right). \tag{26}$$

By integrating the lift-per-unit-span over the kite span, and projecting according to the flow angle, we can find the lift force to apply to each $\mathfrak{s}^{\mathrm{th}}$ kite:

$$\boldsymbol{L}_{\mathrm{eff},\mathfrak{s}} = \int_{R_{\mathrm{int}}}^{R_{\mathrm{ext}}} L'(r)\left(\sin\varphi_{\mathrm{flow}}(r)\hat{\mathbf{t}} + \cos\varphi_{\mathrm{flow}}(r)\hat{\mathbf{x}}\right)\mathrm{d}r = \left(1.15\cdot 10^5\,\mathrm{N}\right)\hat{\mathbf{t}} + \left(1.25\cdot 10^6\,\mathrm{N}\right)\hat{\mathbf{x}}, \tag{27}$$

when $\hat{\mathbf{x}}$ points axially downwind, and $\hat{\mathbf{t}}$ is pointing in the tangential direction.

When we simulate the above trajectory by solving the RLL Simulation OCP (23) with, for example, $\mathfrak{I} = 30$ control intervals and $\mathfrak{P} = 2$ periods tracked, we find the vorticity distribution shown in Figure 5, and the axial induction factor $(\boldsymbol{v}_{\mathrm{ind},\mathfrak{k}} \cdot (-\hat{\mathbf{x}})/v_{\mathrm{Ref}})$ distribution over the plane of the kites' spans in Figure 6b. A side-by-side comparison between the axial induction distribution presented in Haas and Meyers (2017) (Figure 6a) and that predicted by solving the above-described RLL Simulation OCP (23) in the `awebox` (Figure 6b) shows that the RLL model captures many of the expected features of the induced velocity distribution: the axial-direction induction concentration behind the kites' current locations; the region of downwash, forming a section-wise-expanding annulus; the upwash immediately before the kites, as well as spilling off of the wingtips.

The main difference between the LES prediction of Haas and Meyers (2017) and the result produced with the RLL model, is the 'length' of the region of significant downwash, in the tangential sense. For example, the LES result of Haas and Meyers

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

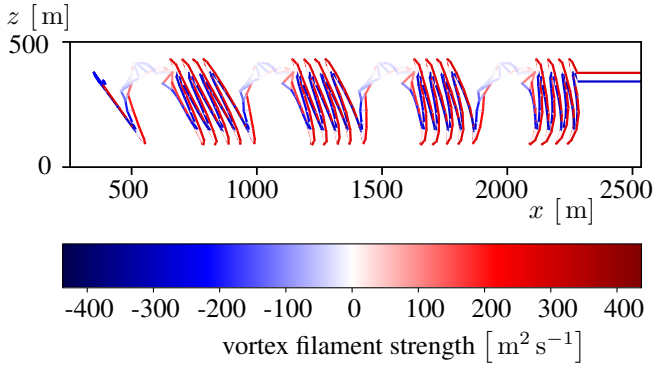


Figure 7. Vorticity distribution predicted by the RLL Simulation OCP in the Haas (2019) verification test, with $\mathfrak{I} = 80$ control intervals and $\mathfrak{P} = 4$ periods tracked. Notice the extremely distinct sections of wake, shed during the reel-in (pale) vs reel-out (dark) phases.

(2017) predicts the region of induction factor 0.2 from the right-most kite (whose lifting-line is drawn in red in Figure 6b.) extending all the way to the line of $y = 0R$, where the RLL model has that region end at about $y = R/4$. The reason for this difference is likely to be caused by the differing vorticity distribution between the two simulations. First, Haas and Meyers (2017)'s elliptical lift distribution, as well as the Gaussian convolution filter that distributes forces from the actuator line into the flow, would result in more vorticity being present immediately behind the (shaded-gray) swept area, than in the RLL model. And since the 3D Biot-Savart relationship includes its inverse-squared-distance relationship, this differing vorticity distribution would impact the induction factors observed on the swept annulus while leaving the areas (white) outside the swept area less affected. Second, we ought not to forget that the RLL model explicitly assumes rigid convection. In a situation where the azimuthally-averaged induction factor is $\overline{a} = \frac{1}{3}$, the induction within the flow is clearly not negligible, upsetting the assumption built into the RLL model's rigid convection. The flow convection speed predicted by (Haas and Meyers, 2017)'s LES simulation is undoubtedly lower than the free-stream velocity, meaning that the pitch of the vortex helices must also be tighter and more irregular than is modeled using the rigidly-convected, lifting-line vortex model. This irregularity can be seen in the slight waver in the boundaries of the contour produced by (Haas and Meyers, 2017)'s LES, which do not show up when the wake is represented with the very regular rigidly-convected, lifting-line vortex model.

Despite this local difference within the swept area, the qualitative similarity between the two results suggests that the proposed RLL model and implementation can describe steady wake effects with a reasonable degree of physicality. Of course, we still need to gain confidence in the model's behavior under conditions more typical of AWE OCP solutions.

3.2 Induction verification under dynamic conditions

In order to verify the RLL model's behavior under dynamic conditions, we'll turn to Haas et al. (2019), in which a single-kite trajectory (including aerodynamic force information) created in an earlier version (2019) of the `awebox` is transferred into the same LES solver used in the rotationally-steady verification. The study contained results computed at different grid resolutions, including one dimensioned as $2700\,\mathrm{m} \times 800\,\mathrm{m} \times 800\,\mathrm{m}$ with a resolution of $6.25\,\mathrm{m} \times 6.25\,\mathrm{m} \times 3.125\,\mathrm{m}$; and another

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

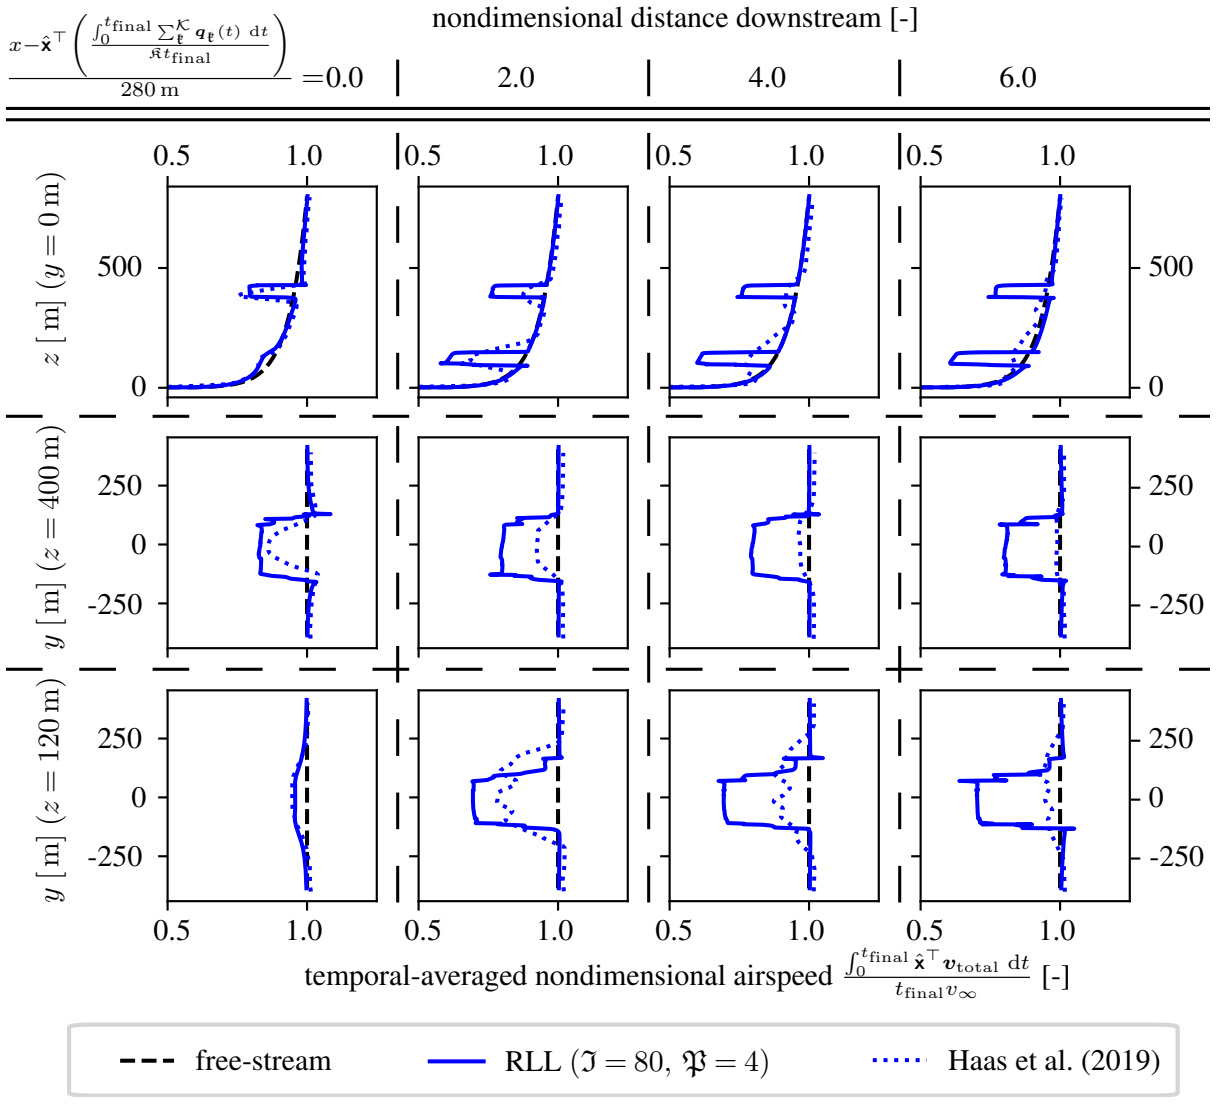


Figure 8. Temporal-averaged downstream-direction velocity distribution, at several downstream locations, under the dynamic-flow conditions of Haas et al. (2019). The top row contains vertical profiles, the bottom two rows contain horizontal profiles. The black-dashed curves are the freestream logarithmic profile; the blue-dotted curves are digitized from Haas (2022, Fig. 3.14); the blue-solid curves are produced by solving the RLL Simulation OCP (23) in the `awebox`.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

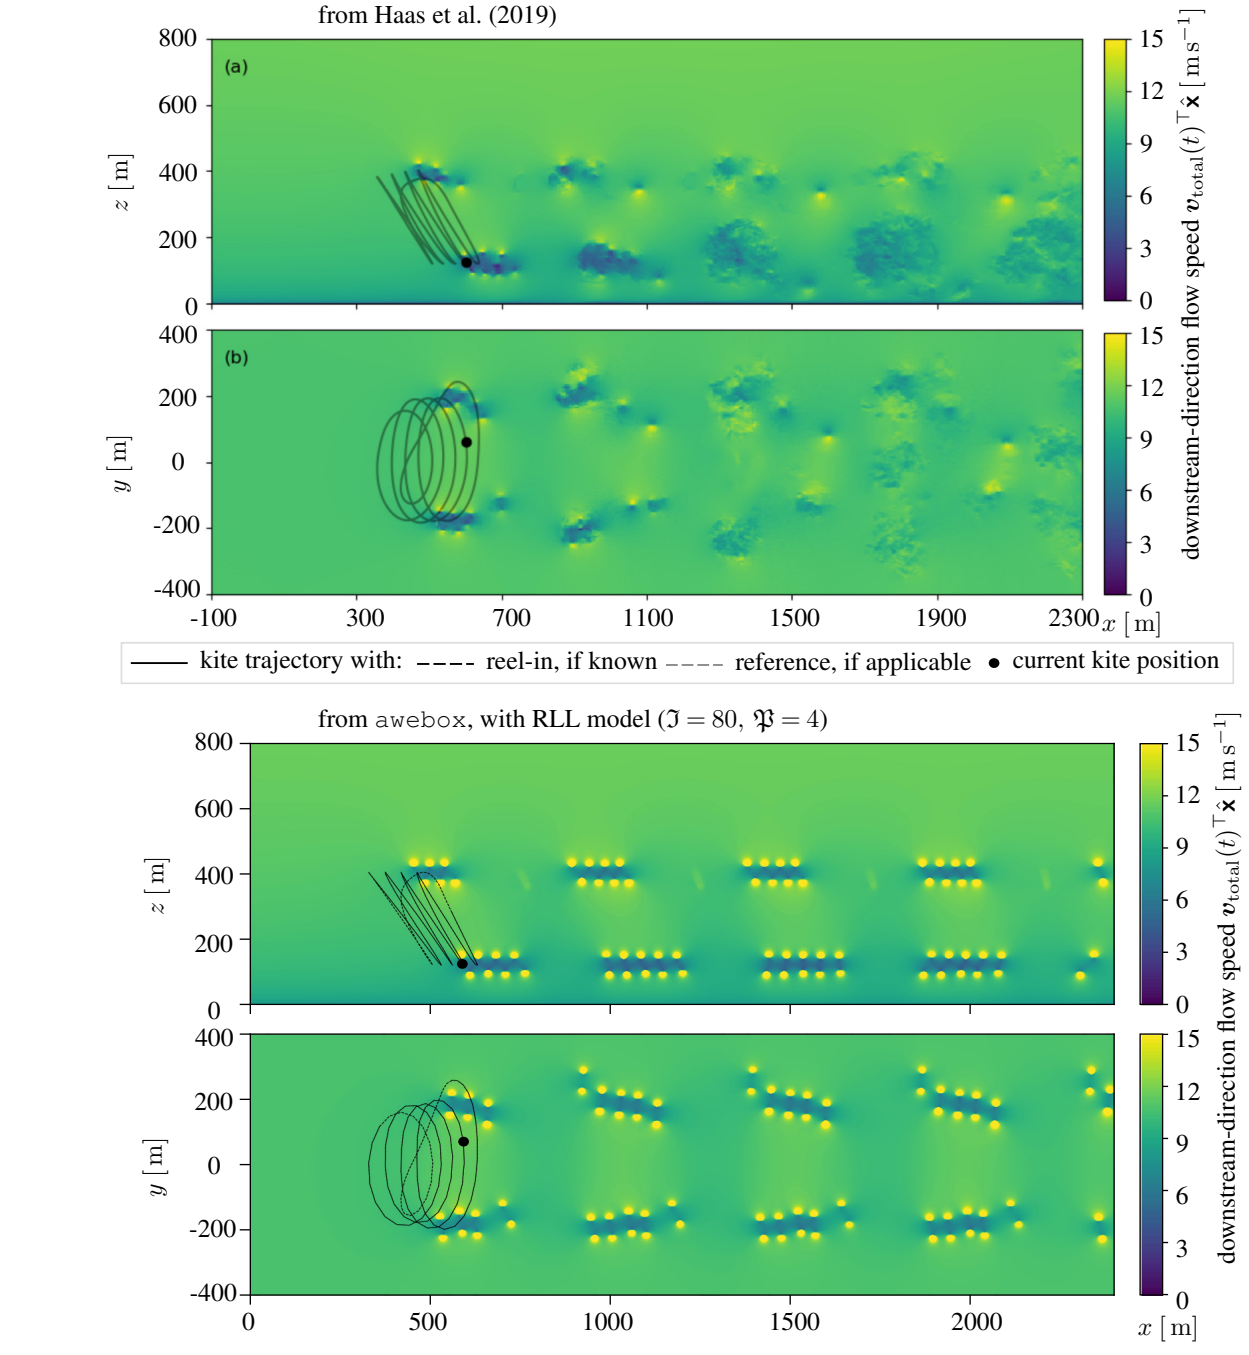


Figure 9. The instantaneous $\hat{\mathbf{x}}$-direction flow speed, under the dynamic-flow conditions of Haas et al. (2019). The top set of plots is included (with coordinates adapted per Haas et al. (2019, Fig.1)) from Haas et al. (2019, Fig.5, included and adapted with permission); the bottom set is produced by solving the RLL Simulation OCP (23). The sub-figures are aligned so that the $\hat{\mathbf{x}}$ coordinates are on the same axis, with the ground-stations mutually positioned at $\boldsymbol{q}_0 = \mathbf{0}$, and without cropping the LES-generated reference. For an overlapping $\hat{\mathbf{y}}, \hat{\mathbf{z}}$-plane projection of the two trajectories, see Figure C1.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

dimensioned as $2720\,\mathrm{m} \times 800\,\mathrm{m} \times 800\,\mathrm{m}$, with a resolution of $4\,\mathrm{m} \times 4\,\mathrm{m} \times 2\,\mathrm{m}$. In both cases, a total period of $100\,\mathrm{s}$ (roughly 2.8 periods) was simulated, with a timestep of roughly $0.1\,\mathrm{s}$. And, while Haas et al. (2019) include results generated with a turbulence inflow and without, we will limit situational differences by focusing here exclusively on the non-turbulent inflow cases.

As mentioned in Table C1, the kite is modeled as a 3-degree-of-freedom point-mass with a controllable lift-coefficient derivative $\dot{C}_\mathrm{L}$ and roll rate $\dot{\varphi}_\mathrm{roll}$, in the style of Houska and Diehl (2007). Then, the kite's drag force is determined with an elliptical-wing drag coefficient that depends exclusively on the zero-lift drag coefficient $C_{\mathrm{D}_0} = 0.0054$ (Haas, 2022, pg. 48), the lift coefficient, and the aspect ratio $\mathit{AR} = 26$ (Haas et al., 2019):

$$C_\mathrm{D} = C_{\mathrm{D}_0} + \frac{C_\mathrm{L}^2}{\pi \mathit{AR}}. \tag{28}$$

Also relevant to the aerodynamic force and the shed-vorticity is the determination of the spanwise and upwards reference axes, respectively $\hat{\boldsymbol{e}}_{2,\mathrm{k},\mathfrak{k}}$ and $\hat{\boldsymbol{e}}_{3,\mathrm{k},\mathfrak{k}}$, of the kite. These unit vectors are determined by rolling the span by an angle φ_roll out of a neutral direction perpendicular to both the tether and the inflowing effective velocity and into a more rotationally-active direction along the tether, then completing a right-handed coordinate system with the chord-wise unit vector $\hat{\boldsymbol{e}}_{1,\mathrm{k},\mathfrak{k}}$ oriented along the apparent inflow:

$$\hat{\boldsymbol{a}} = \frac{\left(\boldsymbol{q}_\mathfrak{k} - \boldsymbol{q}_{\mathcal{P}(\mathfrak{k})}\right) \times \boldsymbol{v}_{\mathrm{app},\mathfrak{k}}}{\left\|\left(\boldsymbol{q}_\mathfrak{k} - \boldsymbol{q}_{\mathcal{P}(\mathfrak{k})}\right) \times \boldsymbol{v}_{\mathrm{app},\mathfrak{k}}\right\|_2}, \quad \hat{\boldsymbol{b}} = \frac{\boldsymbol{v}_{\mathrm{app},\mathfrak{k}} \times \hat{\boldsymbol{e}}_v}{\left\|\boldsymbol{v}_{\mathrm{app},\mathfrak{k}} \times \hat{\boldsymbol{e}}_v\right\|_2}, \tag{29}$$

$$\hat{\boldsymbol{e}}_{1,\mathrm{k},\mathfrak{k}} = \frac{\boldsymbol{v}_{\mathrm{app},\mathfrak{k}}}{\left\|\boldsymbol{v}_{\mathrm{app},\mathfrak{k}}\right\|_2}, \quad \hat{\boldsymbol{e}}_{2,\mathrm{k},\mathfrak{k}} = \hat{\boldsymbol{a}}\cos\varphi_\mathrm{roll} + \hat{\boldsymbol{b}}\sin\varphi_\mathrm{roll}, \quad \hat{\boldsymbol{e}}_{3,\mathrm{k},\mathfrak{k}} = \hat{\boldsymbol{b}}\cos\varphi_\mathrm{roll} - \hat{\boldsymbol{a}}\sin\varphi_\mathrm{roll}.$$

Here, the aerodynamic forces are defined with the lift- and drag- components projected and non-dimensionalized with the apparent velocity:

$$\boldsymbol{F}_\mathrm{aero} = \frac{1}{2}\rho_{\mathrm{air}\infty}(\boldsymbol{q}_\mathfrak{k})\left\|\boldsymbol{v}_{\mathrm{app},\mathfrak{k}}\right\|_2^2 s\left(C_\mathrm{L}\hat{\boldsymbol{e}}_{3,\mathrm{k},\mathfrak{k}} + C_\mathrm{D}\hat{\boldsymbol{e}}_{1,\mathrm{k},\mathfrak{k}}\right). \tag{30}$$

Notice that the twin facts that - in the 3-degree-of-freedom model - the lift and drag coefficients are directly determined by the tracked state C_L and control $\dot{C}_\mathrm{L}$ in (28) and that the direction of the forces is determined by the apparent velocity without any influence from the induced velocity-field in (30), mean that aerodynamic forces in the verification test are prescribed from the reference trajectory.

To help direct this reference trajectory into the local minimum reported in Haas et al. (2019), we will also introduce three additional constraints into the problem: we will pin (with equality constraints) the optimization period and the tether diameter to the solution values reported in Haas et al. (2019) and Haas (2022), and we will cap the maximum node altitude slightly above the addition of the "280 m... difference between [the] highest and lowest altitudes flown by the AWES" to the $120\,\mathrm{m}$ lower altitude bound(Haas et al., 2019). The applied inequality constraints are listed in Table C2, with the newly added constraints underlined.

Beyond this, the most relevant input to the baseline AWE OCP, when trying to match Haas et al. (2019)'s trajectory, turns out to be the tether drag model. The paper reports that the tether drag force is modeled using Argatov and Silvennoinen (2013)'s method of computing the drag-force in relation to the aerodynamic force on a point-mass kite (the model named 'kite-only'

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

in the `awebox`). However, the baseline AWE OCP finds trajectories that are much-closer matches to those shown in the results published in Haas et al. (2019) when the then-default `awebox` tether drag model ('equivalence-buggy') is used instead. Since the RLL wake model does not include the tether-drag among the aerodynamic forces that result in shed vorticity, we will proceed with this verification test using the 'equivalence-buggy' tether drag model. For the sake of reproducibility, the implementation of this tether drag model is discussed in greater detail in Appendix Section C.

With these inputs, we can solve the baseline AWE OCP (3) with the parameters and initial guess information given in, respectively, Tables C1 and C3, and find a reference trajectory that corresponds to the trajectory simulated in Haas et al. (2019). We will then track that reference with the RLL Simulation OCP (23) using $\mathfrak{I} = 80$ control intervals, $\mathfrak{D} = 4$ collocation nodes per control interval, and $\mathfrak{P} = 4$ periods tracked[23], simulate the wake modeled by the RLL model, and make two comparisons against the results published in Haas et al. (2019) and Haas (2022).

In the first planned comparison (see Figure 8), we compare the temporal-average $\hat{\mathbf{x}}$-component velocity deficits at different locations (top row) sweeping vertically along $\hat{\mathbf{z}}$, and (middle and bottom row) sweeping horizontally along $\hat{\mathbf{y}}$, through the virtual-wind-tunnel.

When this non-dimensionalized downstream-direction flow speed is considered at, or near, the trajectory - where the viscosity has not had much time to dissipate any shed vorticity - the RLL model (blue, solid) predicts (wake-center) values with magnitudes similar to the LES-determined result in Haas et al. (2019) (blue, dotted). In this region, the RLL model also predicts that the regions where the flow speed deviates from the freestream's logarithmic values - that is, the wake - will be similarly located with a similar width, compared to that computed with the LES model. Beyond this region, there are some notable differences.

First, though the match in velocity deficit at the center of the 'upper' and 'lower' wakes when first shed (left column) is very good, the rigidly-convected and inviscid RLL predicts very little change to the induced velocity as the vorticity is convected. (Further downstream than the tracked four-period's-worth of convected wake, the induced velocity will very quickly drop to zero, as the influence of the modeled vorticity cuts off.) Since, by contrast, the LES simulation's shrinking-amplitude and widening-width Gaussians represent the momentum distribution that occurs due to both mixing and viscous decay, the RLL model begins to over-predict and over-concentrate the velocity deficit as the evaluation point moves further and further downstream (columns further to the right).

Second, the velocity deficits predicted by the RLL are (generally) much more uniform (box- or hat-shaped) within the wake, than predicted by the LES 'ground truth'. It's uncertain how much of this difference is due to the assumptions included in the LES (the Gaussian convolution filter used by the actuator line method to distribute the momentum-change due to the aerodynamic force into the flow) and how much is due to the assumption included in the RLL model of a uniform spanwise circulation distribution. We'll observe, however, that the vorticity distribution (see Figure 7) predicted by the RLL Simulation OCP includes segments (shed during the reel-out phase) when the pitch of the shed vortices appears fairly tight. The tighter the

[23] We chose to use $\mathfrak{P} = 4$ because of the convection distance shown in (Haas et al., 2019, Fig.5). This corresponds to: number of vortex elements modeled per instant $\mathfrak{E} = 963$ and the total number of vortex elements (at all discrete times) $\mathfrak{V} \approx 3.8 \cdot 10^5$.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

pitch of a constant-strength vortex helix, the more a vortex-structure will resemble a vortex cylinder, whose interior would be expected (Branlard, 2017, p.433) to have a uniform induced velocity.

The third difference is the exception to the box-like uniform induced velocity: at the edges of the wake, there are occasional small 'spikes' of velocity. These spikes are always found at the edges of the wake whether one sweeps horizontally or vertically, suggesting that the spikes occur because - at some point during the temporal averaging - the induced velocity must have been evaluated within the (regularized) core radius of one of the trailing vortices, resulting in a comparatively-large induced velocity being included. Such a spike will never appear in the LES result since the Navier-Stokes equations don't contain the Biot-Savart expression's (regularized) singularity.

In the end, when one considers the difference in the wake-resolution between a vortex filament model with $\mathfrak{E} = 963$ elements at any one instance in time and a LES with 14-million[24] control volumes, the match between the predictions made by the two models is quite good.

An instantaneous comparison between Haas et al. (2019, Fig.5)'s LES-simulated 'ground truth' and the RLL solution (see Figure 9) again needs to take into account the assumptions made while developing the RLL model. First, the RLL model assumes rigid convection in an inviscid flow: this explains why the 'bursts' of dark blue (low speed) and yellow (high speed), which represent the slicing-through of the convected vorticity, do not dissipate into clouds of (intermediate-speed) green as one moves downstream from the kite's trajectory. Second, the RLL model assumes a uniform circulation distribution, which means that the induced velocity will not have a smooth (uniform) distribution of downwash between the 'edges' of the convecting vorticity 'ribbon', but will rather take much stronger values (brighter yellow) specifically at the edges of that ribbon, along the tip vortices.

Those differences being acknowledged, we do see the expected trend of slow flow speeds (blue) occurring within 'already-harvested' flow, while the flow immediately around this decelerated air is accelerated (yellow) in a Venturi-like effect. We also see that the spacing and intensity of the 'bursts' of velocity difference reflect the phases of the trajectory, with the long horizontal 'empty' space corresponding to slicing through the small-magnitude vorticity shed during reel-in, and more extreme colors being found when slicing through the large-magnitude vorticity shed during reel-out. One can very distinctly count the five windings of the trajectory, being repeated exactly four times. All of this is to say that the instantaneous flow speed reflects what someone would expect from applying the RLL model's assumptions to the 'true' instantaneous flow speed calculated in Haas et al. (2019).

Having hopefully gained some confidence that the RLL model describes the flow as well as the assumptions will allow, our next step is to consider the numerical behavior of the model within its desired framework: optimal control.

4 Tradeoffs in convergence and expense

We've now formulated a RLL AWE OCP (22) and confirmed that the model works reasonably considering the underlying assumptions. But, anyone who sets out to use the model in an AWE OCP, will also need to know whether increasing the

[24] Haas et al. (2019)'s 'Grid 2' has $432 \times 128 \times 256 = 14,155,776$ control volumes.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

resolution of the resulting RLL AWE OCP leads to the solution converging at a reasonable computational expense. So, let's leave aside the exact details of the solution, and solve an example problem for a two Ampyx-AP2 system and increasing numbers of control intervals $\mathfrak{I}$ and periods' worth of wake tracked $\mathfrak{P}$. The below solutions are all for single-winding trajectories, with the system properties given in Table B1-B2, the inequality constraints given in Table B3, and the initialization information given in Table B4.

4.1 Convergence

For all of the solved combinations of $\mathfrak{I}$ and $\mathfrak{P}$ (see Figure 10), the RLL AWE OCP finds the same single-winding-number local minimum[25]: a trajectory with two overlapping loops, the loop centered further in the $+\hat{\mathbf{y}}$ and $+\hat{\mathbf{z}}$ directions being 'oval-shaped' and the one centered further in the $-\hat{\mathbf{y}}$ and $-\hat{\mathbf{z}}$ directions being 'egg-shaped'. This trajectory is shown in the inset figure in Figure 10c, with the vorticity distributions at the Pareto-extreme wake-resolutions shown in Figure 4. Increasing the number of control intervals $\mathfrak{I}$ - that is, allowing each kite to progress from 15 sequential control inputs per loop to 30 - allows the kite to make fine course-corrections, and therefore increases the power available. (The special case of $\mathfrak{I} = 10$ control intervals will be discussed later.)

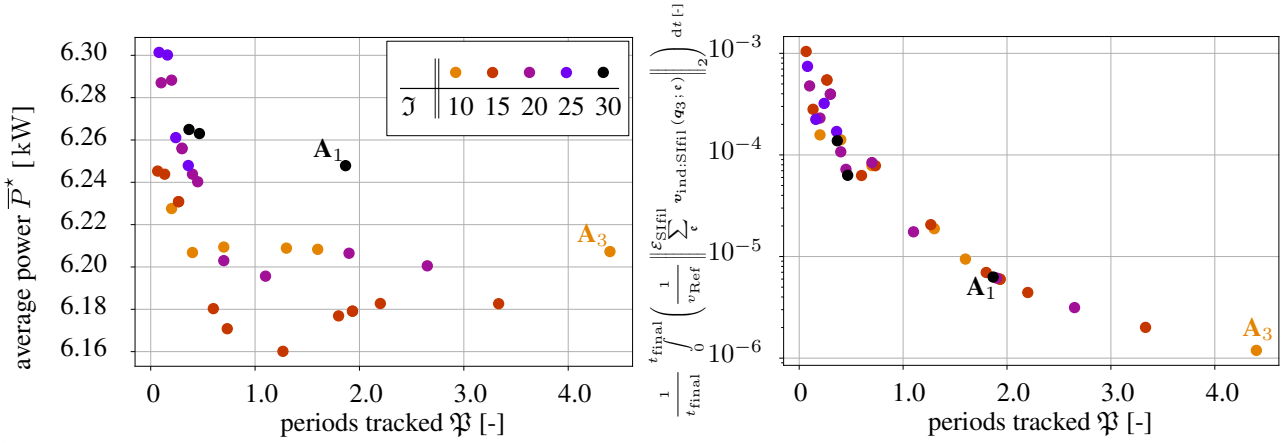
In general, tracking more of the wake's history (increasing $\mathfrak{P}$) decreases the optimal power (see Figure 10c); this general trend reflects the fact that - with an underlying periodic trajectory and an assumption of externally-inviscid flow - a longer resolved wake means that more vorticity will be modeled, resulting in greater induced velocities at the kites and less power available to be extracted from the flow. Consistent with this trend is the observation that the largest drop in power happens while the lengthening-tracked-wake causes the kites to overfly-for-the-first-time more of each other's shed- and trailing-vorticity. Since the first kite ($\mathfrak{k} = 2$) overflies[26] the second kite's wake at roughly $.11t^\star_{\text{final}}$ and $0.39t^\star_{\text{final}}$ and the second kite ($\mathfrak{k} = 3$) overflies at roughly $.07t^\star_{\text{final}}$ and $0.36t^\star_{\text{final}}$, the power found with a longer tracked wake ($\mathfrak{P} > 0.39$) is fairly stable. (The case of $\mathfrak{I} = 15$ control intervals will be discussed in a later paragraph.) Since the increasing-length-wake acts as a source of aerodynamic inefficiency, the optimizer uses the assumptions of the model - the assumption of rigid convection and, separately, the discretization of the wake into a finite number of vortex filaments - to help reduce that inefficiency.

The impact of this first modeling assumption - rigid convection - means that the optimizer makes a small increase to the system's operating altitude. Specifically: when there is no wake tracked, the trajectory saturates the tether length upper bound (700m) at the end of the reel-out period, keeping the elevation angles relatively small (13.5deg) to avoid cosine losses. But, as more of the wake is tracked, the optimizer adjusts its relative prioritization of the losses due to the wake versus to tether drag. For $\mathfrak{I} = 15$ control intervals (where the impact on the power in Figure 10a is easiest to see), the altitude increases by 4 meters (to 164.5m) between $\mathfrak{P} = 0.07$ and 2.2, to increase the log-law freestream wind speed. (That increasing the flight altitude leads to increased aerodynamic forces on the kites, is visible in the increased circulation - and correspondingly-stronger induced velocity from the stronger recently-shed vorticity - experienced by the kite flying the higher-average-altitude 'oval-

[25] even using the not-guaranteed-to-be-repeatable linear-solver `MA86`

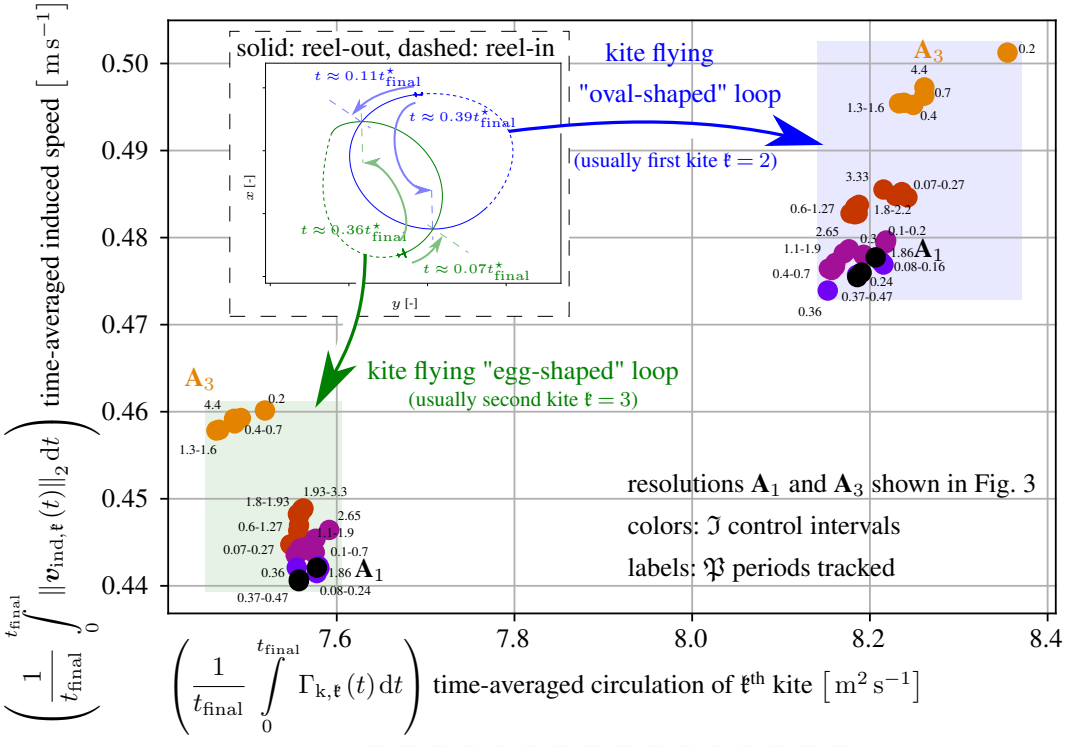
[26] The overflying times are calculated based on the projection of the kite locations in the plane perpendicular to the convection direction, the $\hat{\mathbf{y}}$-$\hat{\mathbf{z}}$ plane, as seen in the inset of Figure 10c.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



(a) Optimal system power converges to stable values as length of tracked wake increases; increasing the control-authority by increasing the number of control intervals (with piecewise-constant controls) increases the optimal power.

(b) The time-averaged, non-dimensionalized contribution of the placeholder 'far-wake' semi-infinite filaments to the induced velocity at the second kite ($\mathfrak{k} = 3$) decays with inverse-convected-distance.



(c) As the number of control intervals $\mathfrak{J}$ (colors) increases, each kite's time-averaged circulation and the time-averaged norm of the induced velocity at each kite's mid-span converge towards gradually overlapping centers. Overflight times $[0.07t_{\mathrm{final}} - 0.39t_{\mathrm{final}}]$ suggest that 'dot jumps' are related to the number of tracked overflights. Note that the induced speed is an order-of-magnitude smaller than the reference (at 10m altitude) wind speed $v_{\mathrm{Ref}} = 5\,\mathrm{m\,s^{-1}}$.

Figure 10. Convergence of the two-kite-system's optimal power and experienced induced velocities with control intervals $\mathfrak{J}$ and the periods-tracked of the wake history $\mathfrak{P}$.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

shaped' trajectory in Figure 10c, compared to the kite flying the lower-average-altitude 'egg-shaped' trajectory.) However, an increased freestream velocity on the overall-system will also impact the tether-drag, causing the optimizer to reduce the tether lengths - the maximum main tether extension decreases by 33 meters (to 666m) without changing the reel-in-distance, while the secondary tethers, needed for crosswind-flight, shorten by 0.3 meters (to 63.9m) - with the result that the system's elevation angle increases by a degree (to 14.5deg). This 'priority-adjustment' happens most intensely during the first period-of-wake-tracked, but it continues at a slower pace as further-rounds of wake-overflying take place: for $\mathfrak{I} = 15$ control intervals, the elevation angle increases by 1.42-degrees-per-period-tracked between $\mathfrak{P} \in [0.27, 0.73]$, whereas there is almost no change after 1.5 periods have been tracked (0.12-degrees-per-period-tracked between $\mathfrak{P} \in [1.8, 2.2]$).

The curious thing is that the optimal power recovers back a small amount of this dip, and the amount of the recovery is resolution dependent: notice that Figure 10a shows much less recovery in the $\mathfrak{I} = 20$ control intervals solutions than in the $\mathfrak{I} = 15$ solutions. The resolution-dependence implies that the optimizer is taking advantage of the specific-details of how the shed vorticity is modeled to 'break physics' in some way. Of these specific-implementation-details, we might first consider that the optimizer is nudging the semi-infinite 'far-wake' filaments - standing-in for the wake shed during the infinitely-periodic system history - into the most-beneficial position with respect to the kites' locations; but, this explanation appears unlikely considering the fact that, by $\mathfrak{P} > 1$, (2D) Biot-Savart's inverse-distance relationship reduces the velocities induced on the kites by the far wake to a negligible fraction of the incoming flow (see Figure 10b). The second specific implementation detail to consider would be the impact of discretizing what should be a continuous vorticity-field into finite vortex filaments, especially the relatively-low resolution 'rungs' of shed vorticity in the 'vorticity-ladder' that the single-spanwise-discretization lifting-line model trails behind each kite. At a low resolution of the shedding-grid - which, recall, is coupled to the control grid - the optimizer might use minute-adjustments to the control-grid to slide the discrete shed vortices out of the way of the oncoming kites. We would expect this manipulation to have the largest impact for $\mathfrak{P} \in [0.07, 0.39]$ with a decreasing marginal return in power as each additional overflown vortex is further downstream. We would also expect that the optimizer's opportunity to conduct such manipulation would decrease as the resolution of the shedding-grid would tighten with the resolution of the control-grid. And these two expectations match what we see in the data: for the case of $\mathfrak{I} = 15$ control intervals, for a length of resolved wake that includes at-least all four of the most-recently-shed period's overflights, the optimal optimization period is a very-stable $t^\star_{\text{final}} = 11.832\,\text{s} \pm 0.0017\,\text{s}$ while the power is still decreasing ($\mathfrak{P} \in [0.6, 1.27]$), and shifts to another very-stable $t^\star_{\text{final}} = 11.861\,\text{s} \pm 0.0021\,\text{s}$ for $\mathfrak{P} \in [1.8, 2.2]$. The point here is not that the printed decimal points are significant figures, but that the investigated data-points - chosen for the independent reason of computational memory availability when batching tasks - include very-stable clusters with differences that are small with-respect-to the full trajectory but large with-respect-to the clusters' local standard-deviations. Coupled with the aforementioned observation that the $\mathfrak{I} = 20$ solutions show much less 'recovery' of power after the first period tracked than do the lower-resolution $\mathfrak{I} = 15$ solutions; the fact that investigated trials for $\mathfrak{I} = 20$ control intervals did not find a similar clustering in the optimization period[27]; as well as the fact that De Schutter et al. (2026) - with a vorticity-patch rather than vorticity-filament near-wake - makes no mention of any similar power-recovery-

[27] For $\mathfrak{I} = 20$ control intervals, the investigated trials show a roughly-linear decrease in optimal optimization period from 11.86s to the minimum value at ($\mathfrak{P} = 0.7, t^\star_{\text{final}} = 11.83\,\text{s}$), followed by a gradual increase to 11.91 s.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

effect despite only using 8 control intervals per half-period, suggests that the optimizer is manipulating the coupled control- and shedding-grids to adjust the locations of the 'rungs' of shed vorticity wherever the wake resolution is rough enough to make 'cheating physics' attractive.

However, if the argument is that the power-dip-then-recovery is related to the resolution of the shedding-grid, then why is this behavior not seen in the solutions with the lowest wake resolution, those with only $\mathfrak{I} = 10$ shedding instances per winding? Extrapolating the argument suggests that the reason that the $\mathfrak{I} = 10$ solutions show no dip-then-recovery is exactly-because the 'rungs' of the shed vorticity are too far apart. After the first overflights-of-the-most-recent-periods'-worth-of-wake have been completed ($\mathfrak{P} > 0.39$), the optimizer is always able to 'dodge' geometrically-inconvenient Biot-Savart cross-products, with the result that the solution is not even pushed to a faster-freestream operating altitude. For the data-points $\mathfrak{P} \geq 0.4$, the average altitude only increases a fairly-constant 1m above its $\mathfrak{P} = 0.2$ value of 161m; the elevation angle only increases by 0.1deg, and the maximum tether length upper-bound remains saturated.

All of the above is to say that: for wake lengths that track one-full-optimization-period plus a reel-out phase ($\mathfrak{P} \geq 1 + t^\star_{\text{switch}}/t^\star_{\text{final}}$ with the optimal switching time $t^\star_{\text{switch}}$ being the time that the system switches from the reel-out phase to the reel-in phase), the tracked wake typically includes enough overflight locations that the solution will have reached its final optimal-average-power value, regardless of the number of control intervals. For control intervals $\mathfrak{I} \geq 20$ the resolution of the wake is fine enough that the optimal power no longer demonstrates much optimizer trickery, with 'the aerodynamic experience' of the kites - as judged by time-averaged circulation and the average induced speeds experienced - no longer changing much with further increases in resolution. The observed convergence trends appear reasonable, considering the modeling assumptions: specifically, rigid convection, a coupled shedding- and control-grid, and a low-frequency shedding of shed-vortices whose strengths are based on the time-averaged-circulation.

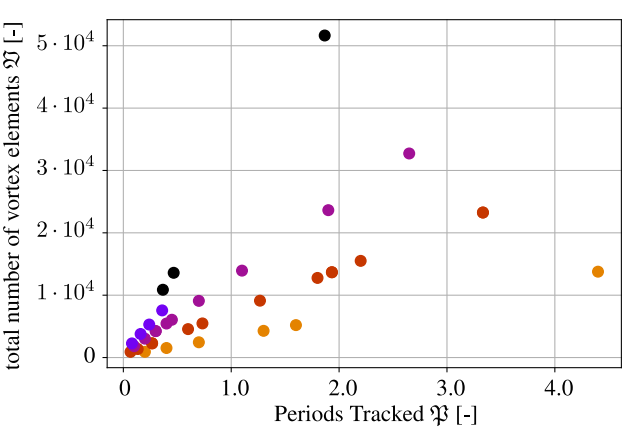
4.2 Expense

When the problem is solved along the full homotopy path, the inclusion of the rigidly-convected, lifting-line vortex model into the AWE OCP unfortunately comes with a large cost, both in terms of computational memory needed and in-terms of processor time. These costs are related to the number of vortex elements that have to be 'remembered' (that is, included in the model, then instantiated within the OCP's discretization) and 'processed' (by requiring the solution of additional dynamics and equality constraints). This means that the value we need to consider when discussing the problem expense is the total number of vortex elements (at all discrete times) $\mathfrak{V}$:

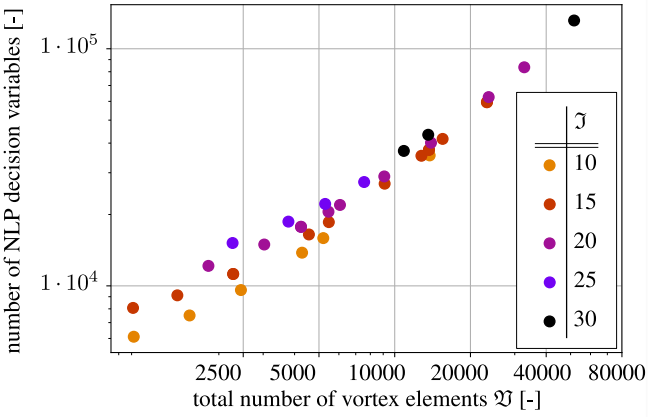
$$\mathfrak{V} = \mathfrak{E}\left((\mathfrak{I}+1)+\mathfrak{I}\mathfrak{D}\right) = (3\mathfrak{K}\mathfrak{A})\left((\mathfrak{I}+1)+\mathfrak{I}\mathfrak{D}\right) = 3\mathfrak{K}\text{ceil}\left(\mathfrak{I}\mathfrak{P}+1\right)\left((\mathfrak{I}+1)+\mathfrak{I}\mathfrak{D}\right), \tag{31}$$

where $\text{ceil}\,()$ is the ceiling function, since the number of wake nodes $\mathfrak{A}$ must be a positive integer. In this expression, we see that the value $\mathfrak{V}$ increases with the number of vortex elements modeled at any time-instance $\mathfrak{E} = \lfloor \mathcal{E} \rceil$ - a value that itself increases linearly with the number of wake nodes $\mathfrak{A}$ and therefore, according to (8), with the number of periods tracked $\mathfrak{P}$ - as well as with the number of control intervals (roughly) squared $\mathfrak{I}^2$ and the collocation order $\mathfrak{D}$. This relationship is shown in the top-left plot of Figure 11.

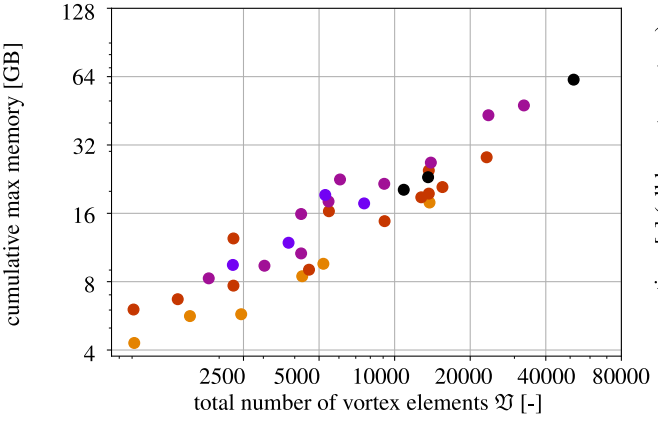
https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



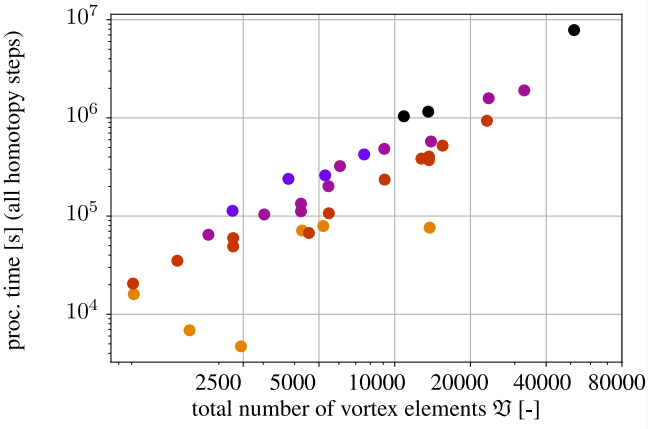
(a) Increasing the wake resolution increases the total number of vortex elements (at all discrete times) $\mathfrak{V}$.



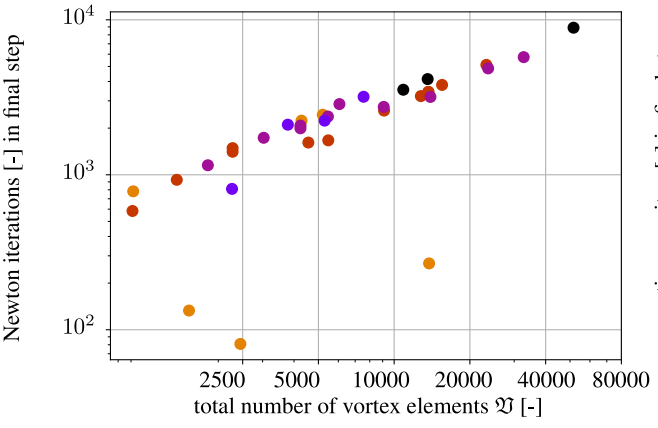
(b) The additional algebraic variables from increasing the total number of vortex elements $\mathfrak{V}$ dominate the RLL AWE OCP's size.



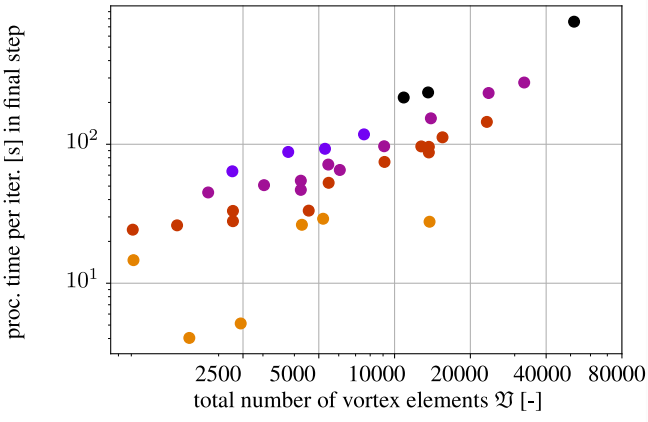
(c) The memory increases roughly linearly ($R^2 = 0.91$) with the total number of vortex elements $\mathfrak{V}$.



(d) The processor time totalled over all homotopy steps increases more than linearly ($\propto \mathfrak{V}^{2.3}$ with $R^2 = 0.95$, note log-log scale).



(e) The processor time is the product of the number of iterations (roughly $\propto \mathfrak{V}$, $R^2 = 0.83$) and the processor time per iteration.



(f) The linear solver uses most (95.6 - 99.8 percent) of this time, so that processor time per iteration grows $\propto \mathfrak{V}^{1.3}$, $R^2 = 0.85$).

Figure 11. Breaking down the computational cost of the RLL AWE OCP, for a two-kite test problem.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

The size of the RLL AWE OCP is heavily dominated by $\mathfrak{V}$, with the number of vortex-related variables vastly outnumbering the baseline variables. When the number of OCP variables is plotted for increasing numbers of control intervals $\mathfrak{I}$ and periods tracked $\mathfrak{P}$, the result is a close-to-linear relationship with the total number of vortex elements (at all discrete times). This can be seen in the top-right plot of Figure 11.

Since the variables added as a result of the RLL model are all algebraic variables - they do not add any actual degrees of freedom to the problem - each newly added variable requires one additional constraint. And, since these variables and constraints have to be stored and handled, the required computational memory grows (roughly) linearly with the total number of vortex elements (at all discrete times). Should a programmer want to add further inequality constraints to the vortex-related variables, such as enforcing a minimum altitude to the wake nodes, this memory should be expected to grow even further.

With the collected constraints of the RLL AWE OCP given in (22), the problem size and the required memory end up matching our expectation that a vortex method's cost should be related to the number of evaluations of the Biot-Savart expression: proportional to the number of vortex elements for a rigid-wake[28] model with a fixed number of observation points. This trend changes when the processor-time needed to solve the final homotopy step of the RLL AWE OCP is considered: this form of 'computational cost' is proportional to (roughly) $\mathfrak{V}^{2.4}$, which is markedly more than the linear proportionality predicted by Biot-Savart. (See the middle row of Figure 11.) The point to notice, here, is that the processor time is not predominantly driven by the evaluation of Biot-Savart when setting up each Exact-Newton's linear-system, but is rather used to solve that linear system. This linear-algebra cost is the product of the cost per linear-solve - which grows with (roughly) $\mathfrak{V}^{1.3}$ as the KKT matrix is close-to sparse - and the number of linear solves necessary - which grows with (roughly) $\mathfrak{V}^{1}$ as a result of imperfect scaling.[29] The processor time per linear-solve and the number of iterations are shown, respectively, in the bottom row of Figure 11.

The above discussion of processor time suggests that, because the rigidly-convected, lifting-line vortex model is embedded within an OCP structure, the methods of estimating computation time that have been theoretically established (see, eg., Branlard (2017)) for the forward-simulation of vortex models, are likely to be misleading. At this point, the authors are unaware of any published heuristic for computational cost in the solution of OCPs with included vortex models, within the literature (DeVries and Paley, 2013; Zheng and Lyu, 2023; van den Broek et al., 2022, 2024; De Schutter et al., 2026). This has a clear downside, in that it can be difficult to predict a-priori how long any particular problem will need for solution, which has consequences in the batching of problems for parallel solution. But, this also hints at a hopeful upside: since the cost[30] of evaluating Biot-Savart is so vastly dwarfed by the cost of the linear system-solves, it is likely that future research will be able to drastically reduce the cost of such problems through improved tuning or more-tailored algorithms.

[28] More commonly stated is the free-wake version of this relationship, wherein the cost is proportional to the number of vortex elements squared(Branlard, 2017, Ch. 41) since, in the free-wake simulation, the Biot-Savart expression needs to be evaluated from every vortex element that can induce velocity, on every vortex element that needs to be convected.

[29] The above-described trends in memory and processor time are determined according to least-square curve-fits of types $w_2 = \phi_1 + \phi_2 \mathfrak{V}$ and $w_2 = \phi_1 + \phi_2 \mathfrak{V}^{\phi_3}$ subject to $\phi_1 \geq 0$, $\phi_2 \geq 0$ and $0.1 \leq \phi_3 \leq 20$ on the data shown in Figure 11. The fit with the best 'goodness of fit' R^2 is the one shown.

[30] Practical experience suggests that using different scaling strategies, such as `IPOPT`'s `'linear_scaling_on_demand=no'` option which rescales each Newton iteration's linear system instead of only rescaling 'on demand', does frequently reduce both the number of iterations and processor time needed to solve the RLL AWE OCP, but introduces much more variability in both metrics and so was not used in the included convergence-and-expense study.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

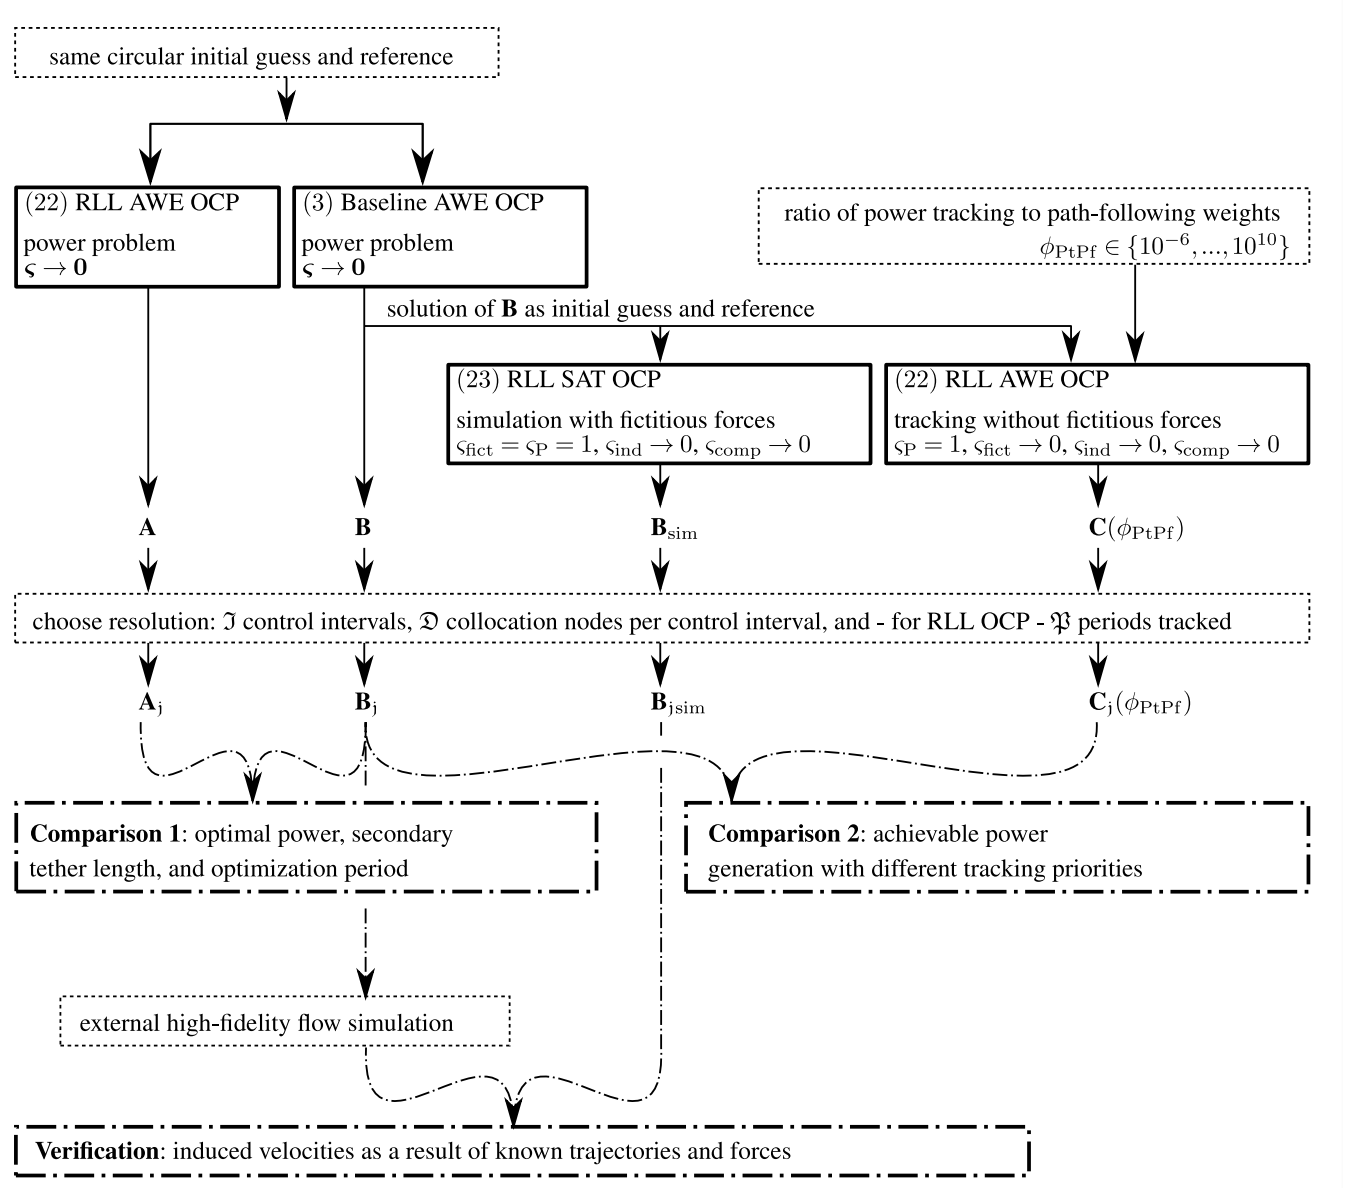


Figure 12. A visual description of the comparisons made with the RLL AWE OCP.

5 Comparison against the baseline optimal control problem

As we've just seen, the computational cost of including the RLL model into the AWE OCP is substantial. The obvious question is: is this cost 'worth it'? This one question has different interpretations based on the asker's practical concerns. The first possibility is: would early design decisions be made differently based on the results of the RLL AWE OCP vs. an AWE OCP built on the assumption that the induced velocity-field will remain negligible whatever trajectory is flown (a.k.a., the

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

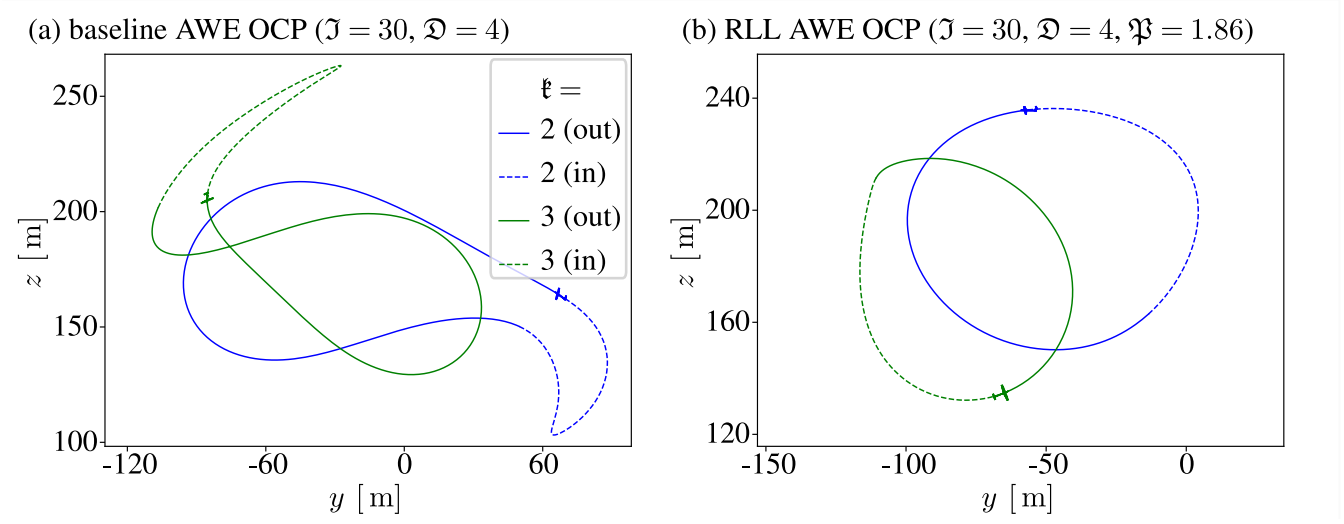


Figure 13. The optimal trajectory found when induction is assumed to be negligible (a, baseline AWE OCP) involves many tight turns including a dramatic reel-in (dashed) swoop to increase the angle between the secondary tethers. By contrast, the optimal trajectory found when induction is included with the RLL model (b) combines the desire to avoid having the kites fly over each other's shed vorticity with a small-distance-between-the-kites during reel-out (solid) and a large-distance during reel-in (dashed). Balanced against the airspeed bounds and the tether drag, this results in a much shorter winding period.

baseline AWE OCP)? And second, would it be difficult to fly a trajectory found by solving this baseline (no-induction) AWE OCP in a 'universe' where momentum is conserved?

5.1 Direct comparison of the optimal solutions of the baseline (no-wake) AWE OCP and the RLL AWE OCP

In having described the aerodynamic forces and moments on the $\mathfrak{s}^{\text{th}}$ kite as a function of the effective velocity -

$$\boldsymbol{F}_{\text{aero},\mathfrak{s}} = \boldsymbol{F}_{\text{aero},\mathfrak{s}}\left(\boldsymbol{v}_{\text{eff},\mathfrak{s}}\right) = \boldsymbol{F}_{\text{aero},\mathfrak{s}}\left(\boldsymbol{v}_{\infty}\left(\boldsymbol{q}_{\mathfrak{s}}\right) - \dot{\boldsymbol{q}}_{\mathfrak{s}} + \boldsymbol{v}_{\text{ind}}\left(\boldsymbol{q}_{\mathfrak{s}}\right)\right)$$

- we have set up the interpretation that the baseline AWE OCP (3) is the specific case of a wake-aware AWE OCP wherein the induced velocities experienced by the kite are assumed to be negligible: $\boldsymbol{v}_{\text{ind}}\left(\boldsymbol{q}_{\mathfrak{s}}\right) \approx \boldsymbol{0}$. There's some justification for that assumption: the former industry-leader Makani Power reportedly observed that AWE systems, as a matter of principle, "sweep out much more wind area relative to their power, [such that] wake losses will be negligible" (Echeverri et al., 2020). But, this is still an assumption. It would be important to know how much of an impact the inclusion of this assumption is likely to make on an (offline) performance prediction, especially on the sort of highly-simplified prediction that would be made at the beginning of the design process when foundational decisions about the system - and the system's eventual niche within an already-developed electricity market - are still being considered.

As a result, we will make a direct comparison - between the baseline (no-wake) AWE OCP and the RLL AWE OCP - of certain solution-outputs that would contribute to the levelized cost of electricity (the maximum average power), the system

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

			units	baseline	RLL	difference
soltn.	average power	$\overline{P}^\star$	kW	9.32	6.25	-33%
	sec. tether length	$\ell^\star_{\text{scdr}}$	m	92.8	64.3	-31%
	optimization period	$t^\star_{\text{final}}$	s	23.1	11.9	-48%
problem size	number of vortex elements modeled per instant	$\mathfrak{E}$	-	–	342	–
	total number of vortex elements (at all discrete times)	$\mathfrak{V}$	-	–	51642	–
	model variables	$\dim(\boldsymbol{w}_{\text{model}})$	-	126	936	$6.42 \cdot 10^0$
	NLP variables	$\dim(\boldsymbol{W})$	-	$1.01 \cdot 10^4$	$1.32 \cdot 10^5$	$1.21 \cdot 10^1$
	number of discrete equality constraints	$\dim(\boldsymbol{G})$	–	$9.50 \cdot 10^3$	$1.31 \cdot 10^5$	$1.28 \cdot 10^1$
	number of discrete inequality constraints	$\dim(\boldsymbol{H})$	–	$2.16 \cdot 10^3$	$2.16 \cdot 10^3$	0%
cost	processor time in set-up	-	CPU-s	$1.25 \cdot 10^2$	$8.08 \cdot 10^3$	$6.34 \cdot 10^1$
	processor time in optimization	-	CPU-s	$1.24 \cdot 10^3$	$7.84 \cdot 10^6$	$6.32 \cdot 10^3$
	cumulative max. memory	-	GB	6.27	62.1	$8.90 \cdot 10^0$

Table 1. The impact of the rigidly-convected, lifting-line vortex model on the two-kite AWE OCP, comparing the outputs of the AWE OCP with the RLL wake model $\mathbf{A}_1$ and without any wake model $\mathbf{B}_1$ (for $\mathfrak{I} = 30$ control intervals, $\mathfrak{D} = 4$ collocation nodes per interval, and, in the RLL case, $\mathfrak{P} = 1.86$ periods tracked). This table shows optimal outputs relevant to the levelized-cost-of-electricity, the system design, and the system operation; resulting problem size; and solution cost in computational time and memory.

design (the optimal secondary tether length), and the system operation (the optimal optimization period). These comparisons are printed in Table 1, where the first thing to notice is the one-third decrease in average power harvested with the RLL wake model. Then, the optimal secondary tether and optimization period found by the RLL AWE OCP are both shorter than those found with the baseline problem. This result is unusual, considering that most wake-aware AWE OCP solutions (Zanon et al., 2014; Leuthold et al., 2017; De Schutter et al., 2026) tend to find exactly the opposite trend: that including a wake model should increase the flight radius by extending both the secondary tether length and flight time. However, referring back to Figure 13a, we see that the baseline optimal trajectory found does not have a uniformly-large flight radius. Instead, the tight turns of the 'interlocked-goldfish' pattern provide many points where the kites would overfly their mutually-shed vorticity. With a model of physics where the specific proximity-to-vorticity becomes relevant because induction is modeled with Biot-Savart, these overflight locations themselves become very relevant.

Indeed, the optimal trajectory found by solving the RLL AWE OCP via homotopy (see Figure 13b) ends up finding a Venn-diagram-like shape with two partially-overlapping loops, instead of the complicated maneuvers of the baseline problem (see Figure 13a). Notice that this solution does align with other expectations: the Venn-trajectory is still oriented so that the reel-out phase corresponds with the kites flying close together - better-aligning the secondary-tethers so that their tensions will transmit to the main tether - while the kites fly at their largest distances from one-another during the reel-in phase - allowing the secondary tethers' tension forces to partially cancel each other as they reach the main tether. But, the natural consequence of the

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

optimizer settling upon a trajectory that can be summarized as 'simple-circles-with-minimal-overflying', is that it becomes very important to reduce the tether drag. This, plus a saturated upper bound on the airspeed, leads to shorter tethers and flight-times.

Further, the fact that the convection is assumed to be rigid means that there's no incentive to make large $\left|\dot{C}_{\mathrm{L}}\right|$ maneuvers in order to drop locally-advantageous vorticity into the flow. Instead, the kites fly their simple Venn-diagram-shaped trajectory, with their shed vorticity convecting downstream in an almost-regular helix. The resulting vorticity distribution (see Figure 4) does much less 'bunching-up' when shed during the reel-out phase, than would-be-shed by the baseline trajectory (see the insets in Figure 14). Even with the shorter optimal optimization period, the combination of this period (see Table 1) and the freestream wind-speed ($8.2 \pm 0.17\,\mathrm{m\,s^{-1}}$ on the first kite, $\mathfrak{k} = 2$) means that the wake convects far enough downstream ($97.7\,\mathrm{m} = 17.8b$) that (3D) Biot-Savart's inverse-squared-radius relationship probably also exerts little preference for the sort of longer-and-more-complicated trajectory that would further stretch the wake shed behind one winding.

Now, there is a reasonable objection that the large differences described in this section's direct-comparison-of-solutions could easily be explained as artifacts of a solution-procedure that schedules the induction homotopy step ($\varsigma_{\mathrm{ind}} \to 0$) of the RLL AWE OCP before the power step ($\varsigma_{\mathrm{P}} \to 0$): since it is specifically in the power-term of the objective that overflights of the kites' mutually-shed vorticity effectively are disincentivized, solving the power-step after the induction model has been enforced might be expected to discourage the trajectory from deviating from the (uniform-circular) shape of the reference. However, a very similar optimal solution is found for the RLL AWE OCP at $\mathbf{A}_1$ when the step order is reversed (that is: first $\varsigma_{\mathrm{P}} \to 0$ then $\varsigma_{\mathrm{ind}} \to 0$). In that case, the optimal trajectory's yz projection has the same Venn-shape as shown in Figure 13b; the optimal average power, secondary tether length, and optimization period are very similar ($\overline{P}^{\star} = 6.27\,\mathrm{kW}$, $\ell^{\star}_{\mathrm{scdr}} = 64.2\,\mathrm{m}$, $t^{\star}_{\mathrm{final}} = 11.9\,\mathrm{s}$) to those reported in Table 1; and even the computation cost is almost the same (differing only because this 'induction-step-after-power' variant RLL AWE OCP was solved outside the parallel task-batching that was used to generate the convergence plots) with $7.35 \cdot 10^3$ CPU-s of processor time needed for set-up, $7.72 \cdot 10^6$ CPU-s for optimization, and a cumulative-maximum memory use of 62.0 GB.

From the above discussion, we determine that the optimal power (proxy for the levelized-cost-of-electricity), secondary tether length (for the system design as a whole), and optimization period (for the system operation) presented in Table 1 are sufficiently different between the RLL AWE OCP and the baseline AWE OCP where the induction is assumed to be negligible, that iterative design processes might be driven in very different directions from these two starting points. From that perspective, it could be quite valuable to know what both of these potential 'starting points' are when approaching preliminary design review.

Unfortunately, making the sort of direct-comparison-of-solutions represented by the first-block of Table 1 will always be difficult in the context of non-convex problems because large differences make it less certain that the outputs-and-trajectories being compared originate from similar local minima in their Lagrangian-landscapes. For this reason, we will switch focus to comparing similar local minima.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

			units	baseline	$\mathbf{C}_2(10^{-6})$	difference
soltn.	average power	$\overline{P}^\star$	kW	9.23	7.65	-17%
	sec. tether length	$\ell^\star_{\mathrm{scdr}}$	m	93.9	93.9	0%
	optimization period	$t^\star_{\mathrm{final}}$	s	23.5	23.5	0%
problem size	number of vortex elements modeled per instant	$\mathfrak{E}$	-	–	186	-
	total number of vortex elements (at all discrete times)	$\mathfrak{V}$	-	–	18786	-
	model variables	$\dim(\boldsymbol{w}_{\mathrm{model}})$	-	126	572	$3.54 \cdot 10^0$
	NLP variables	$\dim(\boldsymbol{W})$	-	$6.74 \cdot 10^3$	$5.13 \cdot 10^4$	$6.62 \cdot 10^0$
	number of discrete equality constraints	$\dim(\boldsymbol{G})$	–	$6.35 \cdot 10^3$	$5.10 \cdot 10^4$	$7.03 \cdot 10^0$
	number of discrete inequality constraints	$\dim(\boldsymbol{H})$	–	$1.44 \cdot 10^3$	$1.44 \cdot 10^3$	0%
cost	processor time in set-up	-	CPU-s	68.1	676	$8.93 \cdot 10^0$
	processor time in optimization	-	CPU-s	220	$1.90 \cdot 10^5$	$8.60 \cdot 10^2$
	cumulative max. memory	-	GB	3.78	14.1	$2.72 \cdot 10^0$

Table 2. The impact of the rigidly-convected, lifting-line vortex model on the two-kite AWE OCP, comparing the outputs of the baseline (no-wake) AWE OCP $\mathbf{B}_2$ with those of the RLL SAT OCP $\mathbf{C}_2(10^{-6})$ where $\phi_{\mathrm{PtPf}} = 10^{-6}$. These problems are solved with $\mathfrak{I} = 20$ control intervals, $\mathfrak{D} = 4$ collocation nodes per interval, and, in the RLL case, $\mathfrak{P} = 1.5$ periods tracked. This table shows optimal outputs relevant to the levelized-cost-of-electricity, the system design, and the system operation; resulting problem size; and solution cost in computational time and memory.

5.2 The difficulty of tracking a baseline (no-wake) AWE OCP solution in the RLL AWE OCP

One straightforward alternative for judging the impact of replacing the assumption of negligible induction with the RLL wake model, while remaining in a comparable location within the Lagrangian-landscape, is to take the RLL AWE OCP; set the baseline optimal solution as the initial guess and reference; then solve along the homotopy path until the interior point-, induction-, and fictitious- homotopy parameters have been driven to zero. That is to say: we can try tracking the baseline problem's optimal trajectory, including the RLL model among the constraints. Conceptually, this would be similar to computing an optimal trajectory with the baseline ($\boldsymbol{v}_{\mathrm{ind}} \approx 0$) AWE OCP, then trying to fly that trajectory with a real kite on a day with 'perfect conditions'. Presumably, if it is easy to follow the given trajectory while producing the 'anticipated' amount of power, then the approximation of $\boldsymbol{v}_{\mathrm{ind}} \approx 0$ must be very close to being true. Conversely, if it turns out to be very difficult to track the trajectory and produce the promised power even without any of the complexities of real-world experimental work, then: the assumption was maybe too strong.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

Among the decision variables are variables that describe the translational-and-angular positions-and-velocities of the nodes ($\boldsymbol{q}_{\mathfrak{b}}$ and $\dot{\boldsymbol{q}}_{\mathfrak{b}}$ $\forall \mathfrak{b} \in \mathcal{B}$, as well as $\underline{\underline{\boldsymbol{R}}}_{\mathrm{k},\mathfrak{k}}$ and $\boldsymbol{\Omega}_{\mathfrak{k}}$ $\forall \mathfrak{k} \in \mathcal{K}$), as well as variables that describe the instantaneous mechanical power[31] as the product of the tether tension and reel-out speed ($\tau_{\mathfrak{b}}$ $\forall \mathfrak{b} \in \mathcal{B}$, as well as the ℓ_{t} and $\dot{\ell}_{\mathrm{t}}$). Let's call the first set the 'path-following' variables, and the second set the 'power-tracking' variables.

Now, if we try to track a reference trajectory, the optimizer will 'try the hardest' to track the variables whose weighting exerts the most pressure. So, in order to avoid the potential objection that, whichever value comes out of this tracking experiment will purely be a result of the weights chosen, let's re-weight the RLL AWE OCP. Specifically, let's take a ratio ϕ_{PtPf} and use it to adjust the weights of the power-tracking variables in comparison to the path-following variables. To do this, we will multiply the weights $\underline{\underline{\boldsymbol{K}}}_{\boldsymbol{q}_{\mathfrak{b}}}$, $\underline{\underline{\boldsymbol{K}}}_{\dot{\boldsymbol{q}}_{\mathfrak{b}}}$, $\underline{\underline{\boldsymbol{K}}}_{\underline{\underline{\boldsymbol{R}}}_{\mathrm{k},\mathfrak{k}}}$, and $\underline{\underline{\boldsymbol{K}}}_{\boldsymbol{\Omega}_{\mathfrak{k}}}$ by $1/\sqrt{\phi_{\mathrm{PtPf}}}$; and multiply the weights $\underline{\underline{\boldsymbol{K}}}_{\tau_{\mathfrak{b}}}$, $K_{\ell_{\mathrm{t}}}$, and $K_{\dot{\ell}_{\mathrm{t}}}$ by $\sqrt{\phi_{\mathrm{PtPf}}}$. Then, we will solve the RLL AWE OCP with this ratio-of-power-tracking-to-path-following-weights ϕ_{PtPf} covering a range of orders-of-magnitude from very small (here, 10^{-6}) to very large (here, 10^{10}).

The results of this tracking are given in Figure 14, for $\mathfrak{I} = 20$ control intervals and $\mathfrak{P} = 1.5$ periods tracked. We see that intermediate values of ϕ_{PtPf} - where the optimizer tries to track all of the variables with similar weight and consequently is unable to track any of them particularly well - are the furthest from the power 'Pareto front'. For values of ϕ_{PtPf} smaller than this intermediate, the kites' positions and velocities track their reference (baseline) trajectory visually very well, but the resulting aerodynamic forces are unable to produce the required tether tension. For values of ϕ_{PtPf} that are larger than the intermediate ratio, the tether tension-per-unit-lengths track very well, but the reduced regularization on the secondary tethers' orientations (that is, the kite positions, which cannot be uniquely determined from the tether lengths) means that the vector-sum of the forces doesn't recover the anticipated baseline power.

So, while the baseline (no-wake) AWE OCP anticipates[32] 9.23kW from this trajectory, the inclusion of the RLL model into the SAT OCP means that we wouldn't actually harvest more than, variously, 7.65kW at $\phi_{\mathrm{PtPf}} = 10^{-6}$ or 7.49kW at $\phi_{\mathrm{PtPf}} = 10^{8}$. Finding solutions that harvest (roughly) 80% of the anticipated baseline problem is a substantial improvement over the value found in Section 5.1 when the optimizer followed the full-homotopy-path through the decision space (including the system parameters and optimization period) to the 67%-of-anticipated-baseline-power local minimum. The fact that tracking with the RLL model can find 'better' solutions than full solution with the RLL AWE OCP, suggests that the solution is heavily influenced by the specific homotopy-initialization used and that the RLL AWE OCP would benefit from future research investigating adjusted- or alternative- strategies in order to restore `IPOPT`'s typical insensitivity to initialization. In any case: it is up to future AWE OCP users to decide whether being unable to harvest 17%-19% of anticipated power in exchange for still-substantial computational expense and an early-lock-in on system-design-parameter space, is an acceptable tradeoff. For that comparison, the computational cost of solving the $\phi_{\mathrm{PtPf}} = 10^{-6}$ problem is listed in Table 2.

One further note about the central plot shown in Figure 14: when the ratio-of-power-tracking-to-path-following-weights ϕ_{PtPf} becomes very small and the regularization applied to the power-tracking variables disappears, those variables are still

[31] Recall from page 6 and (5) that all of the terms that enter the instantaneous mechanical power are included within the (continuous-form) OCP's decision variables as algebraic variables.

[32] This value is smaller than the power found in Section 5.1 despite being at 'the same' local minimum of the solution landscape, because the trajectory here was only discretized with $\mathfrak{I} = 20$ control intervals, compared to the 30 used in Section 5.1.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

uniquely determined by the model's holonomic constraints. By contrast, when the ratio ϕ_{PtPf} becomes very large and the regularization applied to the path-following variables disappears, the KKT matrix will approach singularity. Due to its slack-formulation, automatic rescaling, and regularization strategies, the interior point solver with which these problems are being solved (Wächter and Biegler (2006)'s `IPOPT`) is very robust to ill-posed problems; however, the cost of these strategies is that we will lose the certainty that the observed outcomes are actually a result of ϕ_{PtPf} and not of other numerical effects. Therefore, we should not expect that 'simply continuing to increase $\phi_{\mathrm{PtPf}} \to \infty$' will lead the optimizer to find trajectories that continuously improve towards the anticipated baseline power. This is why we have cut-off the domain of the ratio ϕ_{PtPf} at 10^{10}, because this is where the power-tracking worsens in comparison to the next-smallest ratio $\phi_{\mathrm{PtPf}} = 10^{8}$.

In the two comparisons made in this section, modeling the conservation of momentum within the flow using a rigidly-convected, lifting-line vortex model has a fairly large impact on the power that a given AWE system is predicted to be able to harvest in its given freestream wind. We specifically find a very different local minimum with the RLL AWE OCP than when solving the AWE OCP with an assumption of negligible induction effects $\boldsymbol{v}_{\mathrm{ind}}(\boldsymbol{q}) \approx 0$. And, upon trying to 'fly' that baseline (no-wake) optimal trajectory in a universe that does conserve flow-momentum (here, approximated by the RLL model), we find that we're only able to harvest (roughly) 80% of the anticipated power. These observations suggest that a wake model - such as the RLL model - should be included, in some way, during the initial design of an AWE system despite the additional computational cost.

6 Conclusions

The RLL model makes reasonable predictions of the induced velocity field, considering the modeling assumptions, for both rotationally-steady and dynamic situations. When included among the sub-models of an AWE OCP, the resulting power predictions converge as the length of the tracked wake is increased: the main effects are included as soon as enough wake has been tracked that the kites have overflown their most-recent periods' worth of mutually-shed vorticity, though two-periods worth-of-overflown-vorticity should probably be tracked to limit the optimizer's capacity for trickery. For the case that one doesn't know a-priori where the overflight locations will be on the trajectory, the approximation $\mathfrak{P} \approx 1 + (t^{\star}_{\mathrm{switch}}/t^{\star}_{\mathrm{final}})$, or even - conservatively - $\mathfrak{P} \approx 2$ could be used. Limiting the opportunity for optimizer trickery also suggests that, if an AWE OCP's shedding grid is coupled to its control grid, one should avoid using fewer than 20 control intervals per winding/looping of the trajectory.

Between the homotopy approach, the simplification of rigid convection, and the reduction in the number of vortex elements by shedding vorticity based on the time-averaged circulation and a uniform spanwise circulation distribution, an AWE OCP with an included rigidly-convected, lifting-line vortex model can indeed be successfully - if expensively - solved. That expense is driven by the problem's size which is itself dominated by the additional lifted-variables and equality constraints of the RLL model: the required computational memory scales with the total number of vortex elements (at all discrete times), and the solution processor time scales with that number (roughly) squared.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

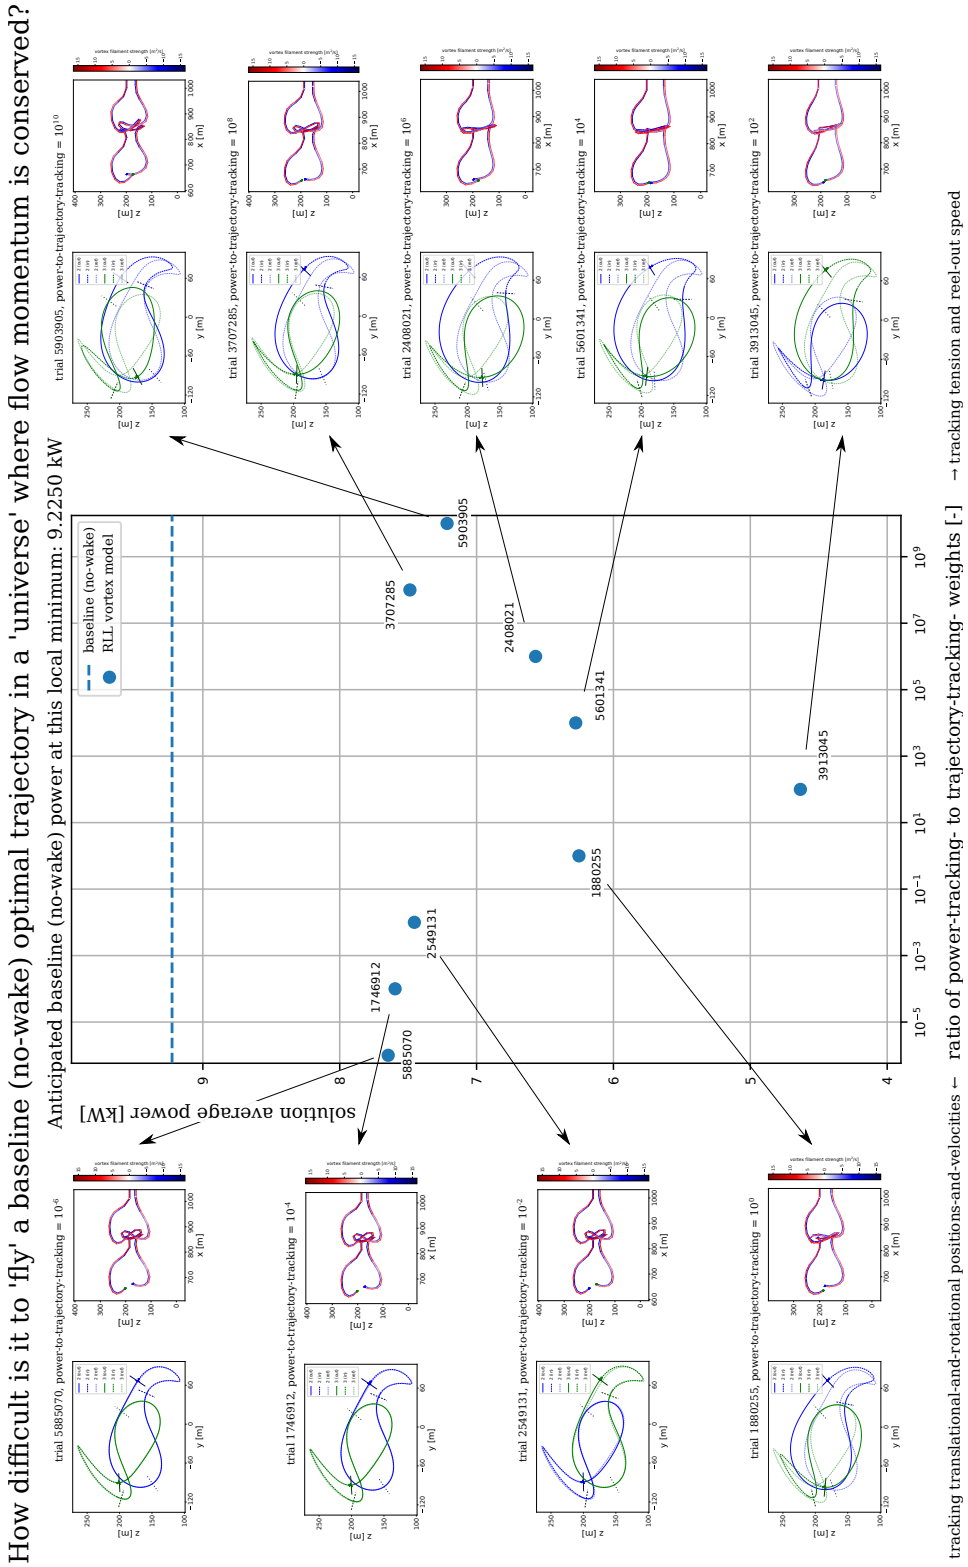


Figure 14. The available power when tracking a baseline (no-wake) AWE OCP optimal trajectory in the RLL AWE OCP with varying relative weights of the power-tracking to path-following variables $\phi_{\texttt{PtPf}}$. This $\mathfrak{I} = 20$ and $\mathfrak{P} = 1.5$ solution results from the 'full-resolution' execution of the approach discussed in Leuthold et al. (2026).

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

Determining whether this cost is warranted, depends on the priorities of the considered task. In this paper's investigation, tracking the optimal trajectory found from an AWE OCP with an assumption of negligible induced velocity-fields using the RLL AWE OCP resulted in only being able to harvest (roughly) 80% of the anticipated power - a deficit that may be large enough to cause financial difficulties for a designer or customer of airborne wind energy systems, but may also be 'close enough' when considering the other sub-component-model-simplifications made. When the RLL is solved as a power-maximization power, the solution outcome is heavily influenced by the homotopy strategy used for initialization: the assumption of rigid wake convection tends to preserve features of the chosen reference trajectory in the fully-converged RLL AWE OCP solution, even more than does a homotopy-initialization on its own.

Naturally, this raises many further questions. Would modeling free-convection help the optimizer to escape local minima by encouraging the kites to 'maneuver', possibly with the goal of cheating physics, or would it simply increase the penalty on overflights? How specifically would the solutions-and-cost of the wake-aware AWE OCP change if the problem is adjusted so that the main opportunities the optimizer has to 'cheat physics' at low resolution are closed: for example, by using quadrilateral vorticity-patches (as used in De Schutter et al. (2026)'s near-wake representation) rather than vortex filaments? And - returning specifically to the context of a preliminary AWE system design with very little knowledge of a kite's specific geometry and load-distribution - how much of the wake effect that dominates the computation cost of the RLL AWE OCP actually requires an expensive vorticity-based model, vs. how much could be included using a cheaper, actuator-based representation? These questions are left for future work.

Appendix A: Vortex elements modeled according to the rigidly-convected, lifting-line vortex model

The vortex elements modeled according to the rigidly-convected, lifting-line vortex model are described in Table A1.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

Str.	Element $\mathfrak{e}$	Shed From $\mathfrak{s}$	$\in \mathcal{E}_{\mathrm{Ffil}}$ (2^{nd} kite in $\mathcal{K}$)	$\in \mathcal{E}_{\mathrm{SIfil}}$	Wake Node $\mathfrak{a}$	$\boldsymbol{q}_{\mathrm{start},\mathfrak{e}}$	$\boldsymbol{q}_{\mathrm{end},\mathfrak{e}}$	Strength $\Gamma_{\mathfrak{e}}$	Type
bound wake	0	1^{st} kite in $\mathcal{K}$	✓	–	0	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,0}$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,0}$	$\overline{\Gamma}_{\mathfrak{s},0}$	finite filament
	1	2^{nd} kite in $\mathcal{K}$	✓	–	0	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,0}$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,0}$	$\overline{\Gamma}_{\mathfrak{s},0}$	
	2	3^{rd} kite in $\mathcal{K}$	✓	–	0	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,0}$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,0}$	$\overline{\Gamma}_{\mathfrak{s},0}$	
	...		✓	–	...			...	
	$\mathfrak{K}-1$	last kite in $\mathcal{K}$	✓	–	0	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,0}$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,0}$	$\overline{\Gamma}_{\mathfrak{s},0}$	
near wake	$\mathfrak{K}$	1^{st} kite in $\mathcal{K}$	✓	–	$0\to 1$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,0}$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,1}$	$-\overline{\Gamma}_{\mathfrak{s},0}$	finite filament
	$\mathfrak{K}+1$		✓	–	$0\to 1$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,0}$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,1}$	$+\overline{\Gamma}_{\mathfrak{s},0}$	
	$\mathfrak{K}+2$		✓	–	1	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,1}$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,1}$	$\overline{\Gamma}_{\mathfrak{s},1}-\overline{\Gamma}_{\mathfrak{s},0}$	
	$\mathfrak{K}+3$		✓	–	$1\to 2$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,1}$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,2}$	$-\overline{\Gamma}_{\mathfrak{s},1}$	
	$\mathfrak{K}+4$		✓	–	$1\to 2$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,1}$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,2}$	$+\overline{\Gamma}_{\mathfrak{s},1}$	
	$\mathfrak{K}+5$		✓	–	2	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,2}$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,2}$	$\overline{\Gamma}_{\mathfrak{s},2}-\overline{\Gamma}_{\mathfrak{s},1}$	
	...	...	✓	–	...			...	
	$\mathfrak{K}+3\mathfrak{a}-1$		✓	–	$\mathfrak{a}$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,\mathfrak{a}}$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,\mathfrak{a}}$	$\overline{\Gamma}_{\mathfrak{s},\mathfrak{a}}-\overline{\Gamma}_{\mathfrak{s},(\mathfrak{a}-1)}$	
	$\mathfrak{K}+3\mathfrak{a}$		✓	–	$\mathfrak{a}\to\mathfrak{a}+1$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,\mathfrak{a}}$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,\mathfrak{a}+1}$	$-\overline{\Gamma}_{\mathfrak{s},\mathfrak{a}}$	
	$\mathfrak{K}+3\mathfrak{a}+1$		✓	–	$\mathfrak{a}\to\mathfrak{a}+1$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,\mathfrak{a}}$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,\mathfrak{a}+1}$	$+\overline{\Gamma}_{\mathfrak{s},\mathfrak{a}}$	
	...		✓	–	...			...	
	$\mathfrak{K}+3(\mathfrak{A}-1)-1$	1^{st} kite in $\mathcal{K}$	✓	–	$\mathfrak{A}-1$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,\mathfrak{A}-1}$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,\mathfrak{A}-1}$	$\overline{\Gamma}_{\mathfrak{s},(\mathfrak{A}-1)}-\overline{\Gamma}_{\mathfrak{s},(\mathfrak{A}-2)}$	
	$\mathfrak{K}+3(\mathfrak{A}-1)$	2^{nd} kite in $\mathcal{K}$	✓	–	$0\to 1$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,0}$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,1}$	$-\overline{\Gamma}_{\mathfrak{s},0}$	
	$\mathfrak{K}+3(\mathfrak{A}-1)+1$		✓	–	$0\to 1$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,0}$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,1}$	$+\overline{\Gamma}_{\mathfrak{s},0}$	
	$\mathfrak{K}+3(\mathfrak{A}-1)+2$	...	✓	–	1	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,1}$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,1}$	$\overline{\Gamma}_{\mathfrak{s},1}-\overline{\Gamma}_{\mathfrak{s},0}$	
	...		✓	–	...			...	
	$\mathfrak{E}-2\mathfrak{K}-1$	last kite in $\mathcal{K}$	✓	–	$\mathfrak{A}-1$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,\mathfrak{A}-1}$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,\mathfrak{A}-1}$	$\overline{\Gamma}_{\mathfrak{s},(\mathfrak{A}-1)}-\overline{\Gamma}_{\mathfrak{s},(\mathfrak{A}-2)}$	
far wake	$\mathfrak{E}-2\mathfrak{K}$	1^{st} kite in $\mathcal{K}$	–	✓	$(\mathfrak{A}-1)\to\infty$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,\mathfrak{A}-1}$	–	$-\overline{\Gamma}_{\mathfrak{s},(\mathfrak{A}-1)}$	semi-infinite filament
	$\mathfrak{E}-2\mathfrak{K}+1$	1^{st} kite in $\mathcal{K}$	–	✓	$(\mathfrak{A}-1)\to\infty$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,\mathfrak{A}-1}$	–	$+\overline{\Gamma}_{\mathfrak{s},(\mathfrak{A}-1)}$	
	$\mathfrak{E}-2\mathfrak{K}+2$	2^{nd} kite in $\mathcal{K}$	–	✓	$(\mathfrak{A}-1)\to\infty$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,\mathfrak{A}-1}$	–	$-\overline{\Gamma}_{\mathfrak{s},(\mathfrak{A}-1)}$	
	$\mathfrak{E}-2\mathfrak{K}+3$	2^{nd} kite in $\mathcal{K}$	–	✓	$(\mathfrak{A}-1)\to\infty$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,\mathfrak{A}-1}$	–	$+\overline{\Gamma}_{\mathfrak{s},(\mathfrak{A}-1)}$	
	...		–	✓	...			...	
	$\mathfrak{E}-2$	last kite in $\mathcal{K}$	–	✓	$(\mathfrak{A}-1)\to\infty$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},-,\mathfrak{A}-1}$	–	$-\overline{\Gamma}_{\mathfrak{s},(\mathfrak{A}-1)}$	
	$\mathfrak{E}-1$	last kite in $\mathcal{K}$	–	✓	$(\mathfrak{A}-1)\to\infty$	$\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},+,\mathfrak{A}-1}$	–	$+\overline{\Gamma}_{\mathfrak{s},(\mathfrak{A}-1)}$	

Table A1. Indexing the $\mathfrak{E}=\lfloor\mathcal{E}\rceil=3\mathfrak{K}+3\mathfrak{K}(\mathfrak{A}-1)=3\mathfrak{K}\mathfrak{A}$ vortex elements in the algebraically-represented lifting-line model, using the `Python` convention of beginning indexing at 0. The wake node positions are subscripted so that the subscript of $\boldsymbol{q}_{\mathrm{wake},\mathfrak{s},\mathfrak{t},\mathfrak{a}}$ refers, respectively, to the index of the shed kite $\mathfrak{s}$ within the set of kites $\mathcal{K}$, the wingtip $\mathfrak{t}$ referring either to the positive or negative side of the kite, and the wake node index $\mathfrak{a}$ that runs from 0 to $\mathfrak{A}-1$. Remember that the semi-infinite vortex filaments in the far wake do not need 'ending coordinates', because their longitudinal direction is chosen instead: parallel to the freestream velocity.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

type	property	symbol	value	units	source
wind	wind profile	–	log	–	
	reference wind speed	v_{Ref}	5	$\mathrm{m\,s^{-1}}$	
	reference altitude	z_{Ref}	10	m	
	roughness length	z_0	0.1	m	
	atmospheric density profile	–	international standard	–	`awebox` default
	reference density	$\rho_{\mathrm{air\,Ref}}$	1.225	$\mathrm{kg\,m^{-3}}$	`awebox` default
kite	number of kites	$\mathfrak{K}$	2	–	
	kite mass	m_{k}	36.8	kg	(Licitra et al., 2017)
	wing area	s	3	$\mathrm{m^2}$	(Licitra et al., 2017)
	wing span	b	5.5	m	(Licitra et al., 2017)
	mean aero. chord	c_{MA}	0.5455	m	area/span (s/b)
	aspect ratio	$\mathcal{AR}$	10.08	–	span/mean aero. chord (b/c_{MA})
	moment of inertia	$\underline{\underline{J}}_{\mathrm{k}}$	$\begin{pmatrix} 25 & 0 & 0.47 \\ 0 & 32 & 0 \\ 0.47 & 0. & 56 \end{pmatrix}$	$\mathrm{kg\,m^2}$	(Licitra et al., 2017)
	degrees of freedom	–	6	–	(Licitra et al., 2017)
tether	control variable	–	$\ddot{\ell}_{\mathrm{t}}$	–	(Licitra et al., 2019)
	density	ρ_{t}	1464	$\mathrm{kg\,m^{-3}}$	(Licitra et al., 2019)
	drag coeff.	C_{Dt}	1.2	–	(Licitra et al., 2019)
	aerodynamic elements	$\mathfrak{Z}$	5	–	2025-`awebox` default
	drag model	–	'multi'	–	2025-`awebox` default
	max. reel-out factor	Φ_{routub}	∞	–	2025-`awebox` default

Table B1. System properties used when evaluating the convergence and expense of the RLL AWE OCP, as when comparing the baseline AWE OCP to the RLL AWE OCP in Sections 4 and 5. All other values are 2025-version `awebox` default values.

Appendix B: Replication Information for the two-kite Optimal Control Problem with the rigidly-convected, lifting-line vortex model

The AWE OCP used to generate the convergence and expense plots of Section 4, as well as in the comparison between the baseline 'no-wake' AWE OCP and the RLL AWE OCP of Section 5, uses the system properties given in Table B1-B2, the inequality constraints given in Table B3, and the initialization information given in Table B4.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

Table B2. Stability derivatives for the Ampyx AP2, from Malz et al. (2019), used when evaluating the convergence and expense of the RLL AWE OCP, as when comparing the baseline AWE OCP to the RLL AWE OCP in Sections 4 and 5

	C_X	C_Y	C_Z	C_l	C_m	C_n
0	-0.0293	–	-0.5526	–	-0.0307	–
α	0.4784	–	-5.0676	–	-0.6027	–
α^2	2.5549	–	5.7736	–	–	–
β	–	-0.1855	–	-0.063	–	0.0577
$\beta\ \alpha$	–	-0.0299	–	-0.0003	–	-0.0849
$\beta\ \alpha^2$	–	0.0936	–	0.0312	–	–
δ_ailn	–	-0.0514	–	-0.2489	–	0.01903
$\delta_\mathrm{ailn}\ \alpha$	–	-0.0024	–	-0.0087	–	-0.1147
$\delta_\mathrm{ailn}\ \alpha^2$	–	0.0579	–	0.2383	–	–
δ_elev	-0.0106	–	-0.315	–	-1.0427	–
$\delta_\mathrm{elev}\ \alpha$	0.1115	–	-0.0013	–	-0.0061	–
$\delta_\mathrm{elev}\ \alpha^2$	–	–	0.2923	–	0.9974	–
δ_rudr	–	0.1032	–	0.00436	–	-0.0404
$\delta_\mathrm{rudr}\ \alpha$	–	0.0268	–	-0.0013	–	-0.0117
$\delta_\mathrm{rudr}\ \alpha^2$	–	-0.1036	–	–	–	0.04089
$\dot{\varphi}_\mathrm{roll}$	–	-0.1022	–	-0.5632	–	-0.0565
$\dot{\varphi}_\mathrm{roll}\ \alpha$	–	-0.014	–	-0.0247	–	-0.9137
$\dot{\varphi}_\mathrm{roll}\ \alpha^2$	–	0.0496	–	0.2813	–	–
$\dot{\varphi}_\mathrm{pitch}$	-0.6029	–	-7.556	–	-11.3	–
$\dot{\varphi}_\mathrm{pitch}\ \alpha$	4.4124	–	0.1251	–	-0.0026	–
$\dot{\varphi}_\mathrm{pitch}\ \alpha^2$	–	–	6.1486	–	5.2885	–
$\dot{\varphi}_\mathrm{yaw}$	–	0.1694	–	0.1811	–	-0.0553
$\dot{\varphi}_\mathrm{yaw}\ \alpha$	–	0.1368	–	0.6448	–	0.029
$\dot{\varphi}_\mathrm{yaw}\ \alpha^2$	–	–	–	–	–	0.0257

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

type	inequality	lower bound	value	upper bound	value	units	other	value
named constraints	tether stress	–	–	$\sigma_{\mathrm{t,yield}}$	$3.6 \cdot 10^9$	$\mathrm{N\,m^{-2}}$	$\Phi_{\mathrm{safety}_\sigma}$	1.5
	airspeed	$v_{\mathrm{eff},\mathfrak{k}_{\mathrm{lb}}}$	10	$v_{\mathrm{eff},\mathfrak{k}_{\mathrm{ub}}}$	32	$\mathrm{m\,s^{-1}}$	–	–
	aero. validity	$\alpha_{\mathfrak{k}_{\mathrm{lb}}}$	-6	$\alpha_{\mathfrak{k}_{\mathrm{ub}}}$	9	deg.	–	–
		$\beta_{\mathfrak{k}_{\mathrm{lb}}}$	-20	$\beta_{\mathfrak{k}_{\mathrm{ub}}}$	20	deg.	–	–
	anticollision	–	–	–	–	–	$\Phi_{\mathrm{safety}_{\mathrm{acoll}}}$	2
	yaw-rotation	$\varphi_{\mathrm{yaw}_{\mathfrak{k}_{\mathrm{lb}}}}$	-40	$\varphi_{\mathrm{yaw}_{\mathfrak{k}_{\mathrm{ub}}}}$	40	deg.	–	–
simple bounds	tether length	$\ell_{\mathrm{t_{lb}}}$	10	$\ell_{\mathrm{t_{ub}}}$	700	m	–	–
	tether speed	$\dot{\ell}_{\mathrm{t_{lb}}}$	-15	$\dot{\ell}_{\mathrm{t_{ub}}}$	20	$\mathrm{m\,s^{-1}}$	–	–
	node position	$\boldsymbol{q}_{\mathfrak{b}_{\mathrm{lb}}}$	$[-\infty, -\infty, 100]^\top$	–	–	m	–	–
	angular velocity	$\boldsymbol{\Omega}_{\mathfrak{k}_{\mathrm{lb}}}$	$[-50, -50, -50]^\top$	$\boldsymbol{\Omega}_{\mathfrak{k}_{\mathrm{ub}}}$	$[50, 50, 50]^\top$	$\mathrm{deg\,s^{-1}}$	–	–
	control surf. defln.	$\delta_{\mathrm{cs}\mathfrak{k}_{\mathrm{lb}}}$	$[-20, -30, -30]^\top$	$\delta_{\mathrm{cs}\mathfrak{k}_{\mathrm{ub}}}$	$[20, 30, 30]^\top$	deg.	–	–
	tension multiplier	τ_{lb}	0	–	–	$\mathrm{N\,m^{-1}}$	–	–
	tether acceleration	$\ddot{\ell}_{\mathrm{t_{lb}}}$	-2.4	$\ddot{\ell}_{\mathrm{t_{ub}}}$	2.4	$\mathrm{m\,s^{-2}}$	–	–
	control surf. def. rate	$\dot{\delta}_{\mathrm{cs}\mathfrak{k}_{\mathrm{lb}}}$	$[-2, -2, -2]^\top$	$\dot{\delta}_{\mathrm{cs}\mathfrak{k}_{\mathrm{ub}}}$	$[2, 2, 2]^\top$	$\mathrm{rad\,s^{-1}}$	–	–
	main tether diameter	$\varnothing_{\mathrm{t_{lb}}}$	10^{-4}	$\varnothing_{\mathrm{t_{ub}}}$	0.1	m	–	–
	secondary tether diameter	$\varnothing_{\mathrm{scdr_{lb}}}$	10^{-4}	$\varnothing_{\mathrm{t_{ub}}}$	0.1	m	–	–
	secondary tether length	$\ell_{\mathrm{scdr_{lb}}}$	0.01	$\ell_{\mathrm{scdr_{ub}}}$	10^3	m	–	–
	optimization period	$t_{\mathrm{final_{lb}}}$	5	$t_{\mathrm{final_{ub}}}$	35	s	–	–

Table B3. Parameters for path inequality constraints $\boldsymbol{h}_{\mathrm{path}}$ and parameter inequality constraints $\boldsymbol{h}_{\boldsymbol{\theta}}$, used in the RLL AWE OCPs solved in this paper. Here, the kite indices $\mathfrak{k} \in \mathcal{K}$ where $\mathcal{K}$ is the set of all nodes within the system's tree-architecture that contain kites, such that the set of kite-nodes $\mathcal{K}$ is a subset of the set of architectural nodes $\mathcal{B}$. The node index $\mathfrak{b}$ iterates over the set of nodes $\mathcal{B} = \{1, 2, ..., \mathfrak{B}\}$.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

property	symbol	value	units	source
main tether diameter	$\varnothing_{\mathrm{t}}$	$5 \cdot 10^{-3}$	m	`awebox` default
secondary tether diameter	$\varnothing_{\mathrm{scdr}}$	$4 \cdot 10^{-3}$	m	`awebox` default
main tether length	ℓ_{t}	200	m	(Licitra et al., 2019)
secondary tether length	ℓ_{scdr}	50	m	`awebox` default
traj. winding number	$\mathfrak{L}$	1	–	–
traj. inclination angle	φ_{incl}	45	deg.	(Licitra et al., 2019)
traj. cone angle	φ_{cone}	15	deg.	(Licitra et al., 2019)
kite groundspeed	$\|\dot{\boldsymbol{q}}_{\mathbf{t}}\|_2$	15	$\mathrm{m\,s^{-1}}$	(Licitra et al., 2019)

Table B4. Initial guess values used to create a reference trajectory for the AWE OCPs of Sections 4 and 5, partially determined using the solution given in Licitra et al. (2019).

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

type	property	symbol	value	units	source
wind	wind profile	–	logarithmic	–	(Haas et al., 2019)
	reference wind speed	v_{Ref}	10	$\mathrm{m\,s^{-1}}$	(Haas, 2022, pg. 51)
	reference altitude	z_{Ref}	100	m	(Haas, 2022, pg. 51)
	roughness length	z_0	$2 \cdot 10^{-4}$	m	(Haas et al., 2019)
	atmospheric density profile	–	international standard	–	`awebox` default
	reference density	$\rho_{\mathrm{air\,Ref}}$	1.225	$\mathrm{kg\,m^{-3}}$	`awebox` default
kite	number of kites	$\mathfrak{K}$	1	–	(Haas et al., 2019)
	kite mass	m_{k}	$8 \cdot 10^3$	kg	(Haas et al., 2019)
	wing area	s	$1.385 \cdot 10^2$	$\mathrm{m^2}$	(Haas et al., 2019)
	wing span	b	60	m	(Haas et al., 2019)
	mean aero. chord	c_{MA}	2.3083	m	area/span (s/b)
	aspect ratio	$\mathcal{AR}$	26	–	(Haas et al., 2019)
	zero-lift drag coeff.	C_{D_0}	0.0054	–	(Haas, 2022, pg. 48)
	degrees of freedom	–	3	–	(Haas et al., 2019)
tether	control variable	–	$\ddot{\ell}_{\mathrm{t}}$	–	(Haas et al., 2019)
	density	ρ_{t}	970	$\mathrm{kg\,m^{-3}}$	(Haas et al., 2019)
	drag coeff.	C_{Dt}	0.0501	–	(Haas et al., 2019)
	aerodynamic elements	$\mathfrak{Z}$	10	–	2019-`awebox` default
	drag model	–	'equivalence-buggy'	–	2019-`awebox` default
	max. reel-out factor	Φ_{routub}	∞	–	2025-`awebox` default

Table C1. System properties used to create a reference trajectory matching the one published in Haas et al. (2019) using the baseline AWE OCP (3), and then to track that reference trajectory in the RLL Simulation OCP, during the dynamic verification test of Section 3.2. All other values are 2025-version `awebox` default values.

Appendix C: Replication Information for the Dynamic-Conditions Verification Test

The AWE OCP used in the dynamic-conditions verification test of Section 3.2 uses the system properties given in Table C1, with the inequality constraints given in Table C2, and the initial guess information given in Table C3.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

inequality	lower	value	upper	value	units	source
tether stress	–	–	$\sigma_{\mathrm{t},\mathfrak{b}_{\mathrm{ub}}}$	$1.03 \cdot 10^9$	$\mathrm{N\,m^{-2}}$	(Haas et al., 2019)
acceleration	–	–	$\|\ddot{\boldsymbol{q}}_{\mathfrak{k}}\|_{2_{\mathrm{ub}}}$	$(12)(9.81)$	$\mathrm{m\,s^{-2}}$	(Haas, 2022, pg. 131)
tether length	$\ell_{\mathrm{t}_{\mathrm{lb}}}$	10	$\ell_{\mathrm{t}_{\mathrm{ub}}}$	$1 \cdot 10^3$	m	(Haas, 2022, pg. 49)
tether speed	$\dot{\ell}_{\mathrm{t}_{\mathrm{lb}}}$	$-\infty$	$\dot{\ell}_{\mathrm{t}_{\mathrm{ub}}}$	∞	$\mathrm{m\,s^{-1}}$	(Haas, 2022, pg. 49)
tether acceleration	$\ddot{\ell}_{\mathrm{t}_{\mathrm{lb}}}$	-10	$\ddot{\ell}_{\mathrm{t}_{\mathrm{ub}}}$	10	$\mathrm{m\,s^{-2}}$	(Haas, 2022, pg. 130)
node position	$\boldsymbol{q}_{\mathfrak{b}_{\mathrm{lb}}}$	$[-\infty, -\infty, 120]^\top$	$\boldsymbol{q}_{\mathfrak{b}_{\mathrm{ub}}}$	$[\infty, \infty, \underline{405}]^\top$	m	(Haas et al., 2019) and tuning
lift coefficient	$C_{\mathrm{L}_{\mathrm{lb}}}$	-0.5	$C_{\mathrm{L}_{\mathrm{ub}}}$	1.675	–	(Haas, 2022, pg. 47)
roll angle	$\varphi_{\mathrm{roll}_{\mathrm{lb}}}$	-80	$\varphi_{\mathrm{roll}_{\mathrm{ub}}}$	80	deg.	(Haas, 2022, pg. 47)
tension multiplier	τ_{lb}	0	–	–	$\mathrm{N\,m^{-1}}$	`awebox` default
tether jerk	$\dddot{\ell}_{\mathrm{t}_{\mathrm{lb}}}$	-100	$\dddot{\ell}_{\mathrm{t}_{\mathrm{ub}}}$	100	$\mathrm{m\,s^{-3}}$	(Haas, 2022, pg. 131)
lift coeff. derivative	$\dot{C}_{\mathrm{L}_{\mathrm{lb}}}$	-5	$\dot{C}_{\mathrm{L}_{\mathrm{ub}}}$	5	$\mathrm{s^{-1}}$	(Haas, 2022, pg. 131)
roll rate	$\dot{\varphi}_{\mathrm{roll}_{\mathrm{lb}}}$	-285	$\dot{\varphi}_{\mathrm{roll}_{\mathrm{ub}}}$	285	$\mathrm{deg\,s^{-1}}$	(Haas, 2022, pg. 131)
optimization period	$t_{\mathrm{final}_{\mathrm{lb}}}$	$\underline{43.8178}$	$t_{\mathrm{final}_{\mathrm{ub}}}$	$\underline{43.8178}$	s	(Haas et al., 2019, Fig. 2, digitized)
tether diameter	$\varnothing_{\mathrm{t}_{\mathrm{lb}}}$	$\underline{9.55 \cdot 10^{-2}}$	$\varnothing_{\mathrm{t}_{\mathrm{lb}}}$	$\underline{9.55 \cdot 10^{-2}}$	m	(Haas, 2022, pg. 50)

Table C2. Inequalities used to create a reference trajectory matching the one published in Haas et al. (2019) using the baseline AWE OCP (3), during the dynamic verification test of Section 3.2. Underlined values are new bounds introduced by the authors to help compensate for any changes in `awebox` AWE OCP formulation and scaling since Haas et al. (2019) was published.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

property	symbol	value	units	source
optimization period	t_{final}	43.8178	s	(Haas et al., 2019, Fig. 2, digitized)
tether diameter	$\varnothing_{\mathrm{t}}$	$9.55 \cdot 10^{-2}$	m	(Haas, 2022, pg. 50)
tether length	ℓ_{t}	500	m	2019-`awebox` default
traj. winding number	$\mathfrak{L}$	5	–	2019-`awebox` default
traj. mean altitude	$\overline{z}$	260	m	(Haas et al., 2019)
traj. radius	R	140	m	$\overline{z} - \hat{\mathbf{z}}^{\top} \boldsymbol{q}_{\mathfrak{t}_{\mathrm{lb}}}$
traj. inclination angle	φ_{incl}	31.3323	deg.	$\arcsin(\overline{z}/\ell_{\mathrm{t}})$
traj. cone angle	φ_{cone}	16.2602	deg.	$\arcsin(R/\ell_{\mathrm{t}})$
kite groundspeed	$\|\dot{\boldsymbol{q}}_{\mathfrak{t}}\|_2$	100.3754	$\mathrm{m\,s^{-1}}$	$2\pi R\mathfrak{L}/t_{\mathrm{final}}$

Table C3. Initial guess values used to create a reference trajectory matching the one published in Haas et al. (2019) using the baseline AWE OCP (3), during the dynamic verification test of Section 3.2.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

The 'equivalence' tether drag model referenced in Table C1, was initially developed to distribute between the two endpoints of a straight, rigid tether, a set of equivalent forces that would exactly match a numerical integral of the tether drag (in all three directions) on the segment, as well as the numerical integral of the tether moment about the tether center in the two directions that include the length-wise direction of the tether, while prescribing zero tether moment twisting the tether about its own length. Unfortunately, the initial implementation of this model in `awebox` had a faulty numerical integration step, which is described below. This model has been left un-corrected in the code, to allow verification tests to be constructed against papers - such as Haas et al. (2019) - produced before the bug was noticed.

In this model, the drag force applied to the node containing a single-kite system's single kite, is determined as:

$$\boldsymbol{D}_{\mathrm{t},\mathfrak{k}} = \overbrace{\underline{\underline{\boldsymbol{R}}}_{\mathrm{t},\mathfrak{k}}}^{\text{return to Earth-fixed coordinates}} \underbrace{\left(\begin{array}{ccc|ccc} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \end{array}\right)}_{\substack{\text{select only the kite-node's forces,} \\ \text{but you could swap the blocks} \\ \text{to select only the kite node's} \\ \text{parent's 'equivalent' forces}}} \overbrace{\left(\begin{array}{cccccc} \frac{1}{2} & 0 & 0 & 0 & 1/\ell_{\mathrm{t}} & 0 \\ 0 & \frac{1}{2} & 0 & 1/\ell_{\mathrm{t}} & 0 & 0 \\ 0 & 0 & \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ \frac{1}{2} & 0 & 0 & 0 & -1/\ell_{\mathrm{t}} & 0 \\ 0 & \frac{1}{2} & 0 & -1/\ell_{\mathrm{t}} & 0 & 0 \\ 0 & 0 & \frac{1}{2} & 0 & 0 & -\frac{1}{2} \end{array}\right)}^{\text{distribute the equivalent forces}} \left(\begin{array}{c} \overbrace{\underline{\underline{\boldsymbol{R}}}_{\mathrm{t},\mathfrak{k}}^{-1}}^{\substack{\text{convert to body-} \\ \text{fixed coordinates}}} \overbrace{\sum_{\mathfrak{z}=0}^{\mathfrak{Z}-1} \boldsymbol{D}_{\mathrm{app},\mathfrak{k},\mathfrak{z}}}^{\substack{\text{integrate the} \\ \text{earth-fixed} \\ \text{forces to} \\ \text{distribute}}} \\ \underbrace{\left(\begin{array}{ccc} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{array}\right)}_{\substack{\text{ignore the} \\ \text{'twisting' moment}}} \underbrace{\underline{\underline{\boldsymbol{R}}}_{\mathrm{t},\mathfrak{k}}^{-1}}_{\substack{\text{convert to body-} \\ \text{fixed coordinates}}} \underbrace{\sum_{\mathfrak{z}=0}^{\mathfrak{Z}-1} \boldsymbol{T}_{D,\mathrm{app},\mathfrak{k},\mathfrak{z}}}_{\substack{\text{integrate the} \\ \text{earth-fixed} \\ \text{moments to} \\ \text{distribute}}} \end{array}\right) \quad \text{, (C1)}$$

$$\boldsymbol{D}_{\mathrm{app},\mathfrak{b},\mathfrak{z}} = C_{\mathrm{D,t}} \frac{1}{2} \rho_{\mathrm{air}} \left(\boldsymbol{q}_{\mathfrak{b},\mathfrak{z}}\right) \left\|\boldsymbol{v}_{\mathrm{app},\mathfrak{b},\mathfrak{z}}\right\|_2 \varnothing_{\mathrm{t}} \ell_{\perp,\mathfrak{b},\mathfrak{z}} \boldsymbol{v}_{\mathrm{app},\mathfrak{b},\mathfrak{z}}, \qquad \boldsymbol{T}_{D,\mathrm{app},\mathfrak{b},\mathfrak{z}} = \left(\left(\frac{1}{2} - \frac{\mathfrak{z}+\frac{1}{2}}{\mathfrak{Z}}\right)\left(\boldsymbol{q}_{\mathfrak{b}} - \boldsymbol{q}_{\mathcal{P}(\mathfrak{b})}\right)\right) \times \boldsymbol{D}_{\mathrm{app},\mathfrak{b},\mathfrak{z}}, \quad \text{(C2)}$$

relying on the rotation matrix $\underline{\underline{\boldsymbol{R}}}_{\mathrm{t},\mathfrak{k}} = [\hat{e}_{1,\mathrm{t},\mathfrak{k}}, \hat{e}_{2,\mathrm{t},\mathfrak{k}}, \hat{e}_{3,\mathrm{t},\mathfrak{k}}]$ to transform the forces and moments into- and out-of a tether-segment-fixed coordinate frame. This rotation matrix is the horizontal concatenation of tether unit-vectors $\hat{e}_{3,\mathrm{t},\mathfrak{k}} = \frac{\boldsymbol{q}_{\mathfrak{k}} - \boldsymbol{q}_{\mathcal{P}(\mathfrak{k})}}{\left\|\boldsymbol{q}_{\mathfrak{k}} - \boldsymbol{q}_{\mathcal{P}(\mathfrak{k})}\right\|_2}$ pointing up along the tether, $\hat{e}_{1,\mathrm{t},\mathfrak{k}} = \frac{\hat{\mathbf{y}} \times \hat{e}_{3,\mathrm{t},\mathfrak{k}}}{\left\|\hat{\mathbf{y}} \times \hat{e}_{3,\mathrm{t},\mathfrak{k}}\right\|_2}$ pointing perpendicular to the tether and the earth-fixed axis $\hat{\mathbf{y}}$, and $\hat{e}_{2,\mathrm{t},\mathfrak{k}} = \frac{\hat{e}_{3,\mathrm{t},\mathfrak{k}} \times \hat{e}_{1,\mathrm{t},\mathfrak{k}}}{\left\|\hat{e}_{3,\mathrm{t},\mathfrak{k}} \times \hat{e}_{1,\mathrm{t},\mathfrak{k}}\right\|_2}$ completing the right-handed set.

To this point, everything was implemented according to the above description. Unfortunately, the terms entering the tether's $\mathfrak{z}^{\mathrm{th}}$ element's drag-force contribution $\boldsymbol{D}_{\mathrm{app},\mathfrak{b},\mathfrak{z}}$, which should have been evaluated at the mid-points of the $\mathfrak{z}^{\mathrm{th}}$ element and should therefore have been functions of $\mathfrak{z}$ and $\mathfrak{Z}$, were instead always evaluated at the mid-point of the entire tether (halfways between the parent node $\mathcal{P}(\mathfrak{b})$ and the tether upper node $\mathfrak{b}$). This mis-implementation reads as follows:

$$\ell_{\perp,\mathfrak{b},\mathfrak{z}} = \sqrt{\left\|\left(\boldsymbol{q}_{\mathfrak{b}} - \boldsymbol{q}_{\mathcal{P}(\mathfrak{b})}\right)\right\|_2^2 - \left(\left(\boldsymbol{q}_{\mathfrak{b}} - \boldsymbol{q}_{\mathcal{P}(\mathfrak{b})}\right)^{\top} \frac{\boldsymbol{v}_{\mathrm{app},\mathfrak{b},\mathfrak{z}}}{\left\|\boldsymbol{v}_{\mathrm{app},\mathfrak{b},\mathfrak{z}}\right\|_2}\right)^2}, \quad \boldsymbol{v}_{\mathrm{app},\mathfrak{b},\mathfrak{z}} = \frac{1}{2}\left(\dot{\boldsymbol{q}}_{\mathcal{P}(\mathfrak{b})} + \dot{\boldsymbol{q}}_{\mathfrak{b}}\right) - \boldsymbol{v}_{\infty}\left(\boldsymbol{q}_{\mathfrak{b},\mathfrak{z}}\right), \quad \text{(C3)}$$

$$\boldsymbol{q}_{\mathfrak{b},\mathfrak{z}} = \frac{1}{2}\left(\boldsymbol{q}_{\mathcal{P}(\mathfrak{b})} + \boldsymbol{q}_{\mathfrak{b}}\right). \quad \text{(C4)}$$

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

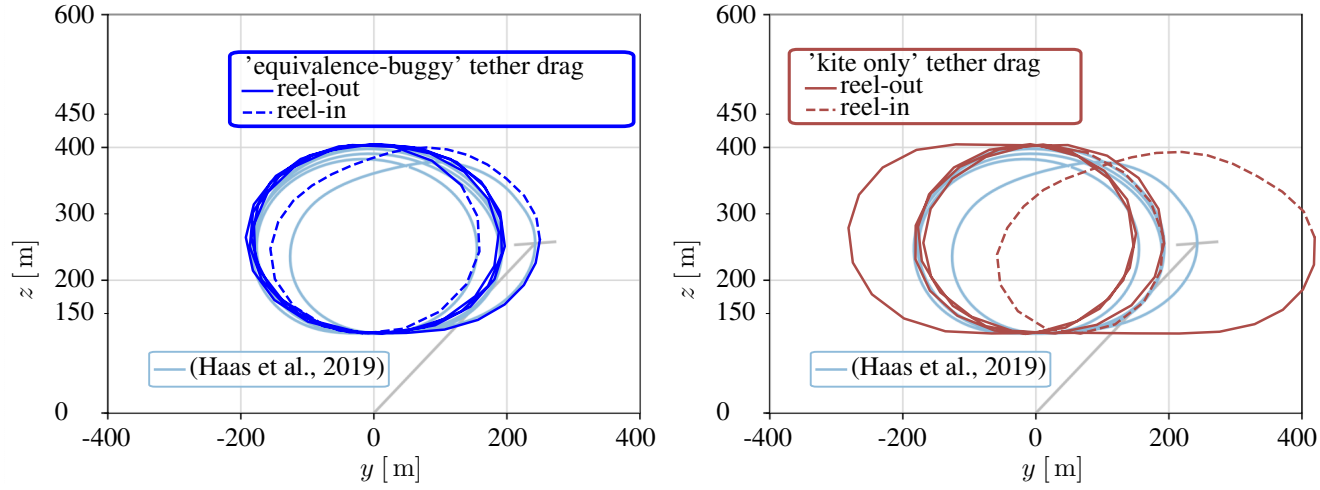


Figure C1. The trajectory produced with the 'equivalence' tether drag model (blue) is a much better match for the trajectory reported in Haas et al. (2019) (turquoise), than that produced with Argatov and Silvennoinen (2013)'s 'kite only' model (red). Both plots are shown over adaptations of Haas (2022, Fig. 3.9b), included with permission.

The end effect is that the implemented model counts the same (poor estimate) of the tether drag 3-times over, giving results that are very different from those produced with other tether drag models.

To produce the trajectory seen in Haas et al. (2019) with Argatov and Silvennoinen (2013)'s 'kite only' model (or any other more reasonable tether drag model) would require substantially more un-noted tuning and manipulation than the few modifications described in Section 3.2 and underlined in Table C2. You can see a comparison between the trajectory reported in (Haas et al., 2019), and the solutions of the baseline AWE OCP with the properties, inequalities, and initial guess given in Tables C1-C3 with both the 'kite only' and the 'equivalence' models in Figure C1.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

Code availability. The results generated here, were produced with `awebox` version 1.0.1 (Sept. 17, 2026, tag: "v1.0.1-wes-2026") available at https://github.com/awebox/awebox. The scripts used are available in the folder `awebox/examples/Leuthold_2026_RLL_paper_scripts`.

Author contributions. Conceptualization: MD (lead), RL (equal). Data curation: RL. Funding acquisition: MD. Investigation: RL. Methodology: RL (lead), MD (supporting), SG (supporting). Project administration: MD (lead), RL (equal). Resources: MD (lead), RL (equal). Software: RL (sole for all wake-related `awebox` OCP aspects, equal for all baseline/non-wake `awebox` OCP aspects), JdS (equal for all baseline/non-wake `awebox` OCP aspects). Supervision: SG (lead for baseline OCP formulation and general OCP tractability tasks, plus the tracking comparison), CC (lead for wake-modeling tasks), MD (lead for combined tasks). Verification: RL. Visualization: RL. Writing – original draft: RL. Writing – review & editing: RL (lead), MD (equal), SG (supporting), CC (supporting).

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. This research was supported by DFG via projects 504452366 (SPP 2364) and 525018088 (MAWERO), by BMWK via 03EI4057A and 03EN3054B, and by the EU via projects AWESCO 642682 and ELO-X 953348. ChatGPT was used to identify spelling- and grammar mistakes and to plot already-generated data, but not for brainstorming, scientific coding, writing, or editing. We also thank Mike McWilliam of the Technical University of Denmark (DTU) for help contextualizing the problem difficulty; Matthieu Duponcheel of the Catholic University of Louvain for a discussion of which reference velocity was used in the generation of the applied Ampyx AP2 stability derivatives; Jakob Hartzer of Albert-Ludwigs-University Freiburg for discussions about optimizer trickery and integration errors as well as sketch-assembly suggestions; Erin Cold of the German Aerospace Center (DLR) for a philosophical discussion leading to the inclusion of the dynamic verification test; Thomas Haas of Ghent University for help reconstructing the verification problems as well as image-reproduction permissions; Xinliang Dai of the Karlsruhe Institute of Technology and François Pacaud of Mines Paris for discussions of memory usage and solver repeatability; as well as Christoph Sieg, Hans Leuthold, and Katharina Steiner for many varied technical discussions as well as encouragement contributing to this work.

Nomenclature

Directions

$\hat{e}_{X,\mathrm{u},\mathfrak{k}}$	in a kite's control reference frame, pointing chordwise, from tail to nose
$\hat{e}_{Y,\mathrm{u},\mathfrak{k}}$	in a kite's control reference frame, pointing spanwise
$\hat{e}_{Z,\mathrm{u},\mathfrak{k}}$	in a kite's control reference frame, pointing upwards, along the axis of symmetry in the half-plane of positive lift
$\hat{\ell}$	semi-infinite filament longitudinal direction

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

$\hat{\mathbf{t}}$	tangential direction
$\hat{\mathbf{x}}$	in the Earth-fixed direction, pointing downstream in the dominant wind direction
$\hat{\mathbf{y}}$	in the Earth-fixed direction, pointing to the edge of the wind-window
$\hat{\mathbf{z}}$	in the Earth-fixed direction, pointing up against gravity

Operations and Mathematical Notation

$\dddot{\square}$	third derivative with respect to time
$\ddot{\square}$	second derivative with respect to time
$\dot{\square}$	first derivative with respect to time
$\hat{\square}$	a unit-vector
$\mathrm{mod}\,(\square, \mathfrak{J})$	the modulo operator, $\square \% \mathfrak{J}$
$\lvert\square\rvert$	the absolute value, element-wise if applied to a vector
$\overline{\square}$	temporal-average
$\lfloor\square\rceil$	the cardinality of a set, a.k.a. the number of elements in the set
$\mathrm{ceil}\,(\square)$	ceiling, the smallest integer value $\geq \square$
$\dim(\square)$	dimension, for example $\dim \mathbb{R}^3 = 3$
$\mathrm{floor}\,(\square)$	floor, the largest integer value $\leq \square$
$\underline{\underline{\square}}$	a matrix or tensor quantity
$\square^\star$	the optimal value
$\mathcal{P}(\square)$	the parent of architectural node $\square$, for example $\mathcal{P}(1) = 0$
$\square'$	derivative with respect to the spanwise distance, as in lift-per-unit-span

Subscripts

∞	the freestream flow
X	component along the kite control axis $\hat{e}_{X,\mathrm{u},\mathfrak{k}}$
Y	component along the kite control axis $\hat{e}_{Y,\mathrm{u},\mathfrak{k}}$

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

	Z	component along the kite control axis $\hat{e}_{Z,\mathrm{u},\mathfrak{k}}$
995	acoll	anti-collision, preventing collision between the kites
	aero	aerodynamic forces and moments acting on a kite
	ailn	aileron, the control surface
	air	the material properties of the air
	app	the apparent velocity (relative inflow not including induced velocity)
1000	bsln	the baseline (no wake) model or OCP
	comp	the inequality constraints' complementarity conditions, or the interior point method of handling them
	cone	the trajectory cone angle (between the 'line of symmetry' of a kite's full trajectory in time and the tether), where meaningful
	cs	kite control surfaces
1005	c	vortex core
	dyn	the model dynamics
	D	drag force
	eff	the effective velocity (relative inflow including induced velocity)
	elev	elevator, the control surface
1010	end	the 'end' of a vortex filament
	ext	an exterior position, especially radially on an annulus
	fict	fictitious forces and moments acting on a kite
	final	a final time, especially for the optimization period
	flow	the flow angle (between the axial-direction inflow and the tangential-direction inflow)
1015	Ffil	finite vortex filament
	hmtp	the homotopy solution process
	incl	the inclination angle (between the main tether and the ground/wind)

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



	ind	the induced velocity, via Biot-Savart
	int	an interior position, especially radially on an annulus
1020	k	the kite
	lb	lower bound
	LL	the lifting line, or lifting line theory
	LMC	Lagrangian momentum correction for open systems
	L	lift force
1025	l	roll moment
	MA	mean aerodynamic, as in 'mean aerodynamic chord'
	model	a local model, as determined by a least-squares fit
	m	pitch moment
	n	yaw moment
1030	path	an OCP's path constraints
	pitch	the pitch angle
	PtPf	the ratio between the power tracking and path-following weights
	P	the mechanical system power
	Ref	the reference solution used in OCP regularization or tracking
1035	reg	OCP regularization
	RLL	rigidly-convected, lifting-line vortex model
	roll	the roll angle
	rout	the main tether's reel-out motion
	rudr	rudder, the control surface
1040	safety	a safety constraint or safety factor
	scdr	a secondary tether (the tether holding a kite, if not the main tether)

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



	shed	a shed vortex, or the instant of shedding
	sim	simulation
	SIfil	semi-infinite vortex filament
	sprptn	the superposition of vortex elements
	start	the 'start' of a vortex filament
	switch	switching from the reel-out to reel-in phase, as in switching time
	S	side force
	tip	a kite's wingtips
	track	tracking the reference trajectory
	t	the main tether, or the material properties of the tether
	ub	upper bound
	wake	the wake
	yaw	the yaw angle
	yield	yield stress, as of the tether (purely material-based, does not include any safety factor)
	$\perp$	in a perpendicular direction, as for tether drag

Physical Symbols

$\boldsymbol{\Omega}$	angular velocity (with magnitude $\Omega = \|\boldsymbol{\Omega}\|_2$)	$[\,\mathrm{rad\,s^{-1}}]$
$\boldsymbol{D}$	drag force (with magnitude $D = \|\boldsymbol{D}\|_2$)	[N]
$\boldsymbol{F}$	force (with magnitude $F = \|\boldsymbol{F}\|_2$)	[N]
$\boldsymbol{L}$	lift force (with magnitude $L = \|\boldsymbol{L}\|_2$)	[N]
$\boldsymbol{S}$	side force (with magnitude $S = \|\boldsymbol{S}\|_2$)	[N]
$\boldsymbol{v}$	a flow velocity (with magnitude $v = \|\boldsymbol{v}\|_2$)	$[\,\mathrm{m\,s^{-1}}]$
α	angle of attack	[rad]
$\mathcal{AR}$	kite aspect ratio	[-]

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

	β	sideslip angle	[rad]
	δ	deflection, as in control surface deflection	[rad]
	$\varnothing$	diameter	[m]
	ℓ	length	[m]
1070	Γ	circulation, or vortex filament strength	[$m^2\,s^{-1}$]
	τ	tether tension per unit length, the Lagrange-multiplier of the holonomic tether-length constraint(s)	[$N\,m^{-1}$]
	Λ	kite speed ratio (kite flight speed to freestream flow speed)	[-]
	$\mathcal{L}$	Lagrangian, in a mechanical or optimization sense	[as needed]
	R^2	coefficient of determination, as the goodness of fit of a regression model	[-]
1075	$\underline{\underline{\boldsymbol{J}}}$	moment of inertia tensor	[$kg\,m^2$]
	$\underline{\underline{\boldsymbol{R}}}$	rotation matrix	[-]
	Φ	a constant or factor, as in 'safety factor'	[-]
	ϕ	any local parameterization	[as needed]
	ρ	density	[$kg\,m^{-3}$]
1080	σ	stress	[$N\,m^{-2}$]
	φ	angle	[rad]
	$\boldsymbol{q}$	position in Earth-Fixed Cartesian coordinates, as generalized coordinates of Lagrangian dynamics	[m]
	$\boldsymbol{T}$	torque or moment	[N m]
	a	axial-direction induction factor	[-]
1085	b	kite wingspan	[m]
	C	a (3D) aerodynamic coefficient, as in the 'lift coefficient'	[-]
	c	kite chord length	[m]
	m	mass	[kg]
	P	the mechanical system power	[W]

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

R	radius, specifically for circular trajectories	[m]
s	area, specifically the planform area for a kite	[m^2]
t	time	[s]
z	altitude	[m]

Optimization Symbols

Continuous form		Discrete form		Description
scalar value	all values	single instance	all instances	
d	$\boldsymbol{d}$	–	–	model dynamics, as components of equality constraints
ψ	–	Ψ	–	objective function
g	$\boldsymbol{g}$	G	$\boldsymbol{G}$	equality constraints
h	$\boldsymbol{h}$	H	$\boldsymbol{H}$	inequality constraints
k	$\underline{\underline{\boldsymbol{k}}}$	K	$\underline{\underline{\boldsymbol{K}}}$	objective weights
r	$\boldsymbol{r}$	–	–	residual functions, as components of equality constraints
u	$\boldsymbol{u}$	U	$\boldsymbol{U}$	controls
w	$\boldsymbol{w}$	W	$\boldsymbol{W}$	decision variables
x	$\boldsymbol{x}$	X	$\boldsymbol{X}$	states
y	$\boldsymbol{y}$	Y	$\boldsymbol{Y}$	algebraic variables
θ	$\boldsymbol{\theta}$	$\boldsymbol{\theta}$	–	system parameters, as decision variables
ξ	$\boldsymbol{\xi}$	Ξ	$\boldsymbol{\Xi}$	boundary conditions
ς	$\boldsymbol{\varsigma}$	$\boldsymbol{\varsigma}$	–	homotopy parameters

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

Indices, Sets, and Counting

Index	Total count	Set of indices	Pronounced as	Counts the
$\mathfrak{a}$	$\mathfrak{A}$	$\mathcal{A}$	"a", "A"	wake nodes, per kite per wingtip
$\mathfrak{b}$	$\mathfrak{B}$	$\mathcal{B}$	"b", "B"	architectural nodes of the system
$\mathfrak{d}$	$\mathfrak{D}$	–	"d", "D"	collocation indices
$\mathfrak{e}$	$\mathfrak{E}$	$\mathcal{E}$	"e", "E"	number of vortex elements modeled per instant
$\mathfrak{i}$	$\mathfrak{I}$	–	"i", "I"	control intervals
$\mathfrak{j}$	$\mathfrak{J}$	–	"j", "J"	any local indexing or counting purpose
$\mathfrak{k}$	$\mathfrak{K}$	$\mathcal{K}$	"k", "K"	kites, especially as observers of induced velocity
–	$\mathfrak{L}$	–	"L"	trajectory loopings or windings in one period
–	$\mathfrak{P}$	–	"P"	periods' worth of wake tracked
$\mathfrak{s}$	$\mathfrak{K}$	$\mathcal{K}$	"s", "K"	kites, as shedders of vorticity
$\mathfrak{t}$	$\mathfrak{T}$	$\mathcal{T}$	"t", "T"	wingtips, per kite
–	$\mathfrak{V}$	–	"V"	vortex elements (at all time instances)
$\mathfrak{z}$	$\mathfrak{Z}$	–	"z", "Z"	tether elements on a segment between two nodes

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

References

Andersson, J. A. E., Gillis, J., Horn, G., Rawlings, J. B., and Diehl, M.: CasADi - A software framework for nonlinear optimization and optimal control, Mathematical Programming Computation, 11, 1–36, https://doi.org/10.1007/s12532-018-0139-4, 2019.

Argatov, I. and Silvennoinen, R.: Efficiency of Traction Power Conversion Based on Crosswind Motion, in: Airborne Wind Energy, edited by Ahrens, U., Diehl, M., and Schmehl, R., chap. 4, pp. 65–80, Springer Berlin Heidelberg, Berlin, Heidelberg, ISBN 978-3-642-39965-7, https://doi.org/10.1007/978-3-642-39965-7_4, 2013.

Bayes, T.: An essay towards solving a problem in the doctrine of chances, Philosophical Transactions of the Royal Society, 53, https://doi.org/10.1098/rstl.1763.0053, 1763.

Beale, J. T. and Majda, A.: Vortex Methods I: Convergence in Three Dimensions, Mathematics of Computation, 39, https://doi.org/10.2307/2007617, 1982.

Betts, J. T.: Practical Methods for Optimal Control and Estimation Using Nonlinear Programming, Second Edition, Society for Industrial and Applied Mathematics, second edn., https://doi.org/10.1137/1.9780898718577, 2010.

Biegler, L. T.: Nonlinear Programming, Society for Industrial and Applied Mathematics, ISBN 978-0-89871-702-0, https://doi.org/10.1137/1.9780898719383, 2010.

Branlard, E.: Wind Turbine Aerodynamics and Vorticity-Based Methods, vol. 7 of *Research Topics in Wind Energy*, Springer International Publishing, ISBN 978-3-319-55163-0, https://doi.org/10.1007/978-3-319-55164-7, 2017.

Burton, T., Sharpe, D., Jenkins, N., and Bossanyi, E.: Wind energy handbook, Wiley, ISBN 9780470699751, 2011.

Crismer, J.-B., Haas, T., Duponcheel, M., and Winckelmans, G.: Airborne wind energy systems flying optimal trajectories in turbulent wind using flight path tracking, Journal of Physics: Conference Series, 2767, 072 021, https://doi.org/10.1088/1742-6596/2767/7/072021, 2024.

Crismer, J.-B., Haas, T., Duponcheel, M., and Winckelmans, G.: Large-eddy simulation of airborne wind energy systems flying in turbulent wind using model predictive control, Wind Energy Science, 11, 2669–2694, https://doi.org/10.5194/wes-11-2669-2026, 2026.

De Schutter, J., Leuthold, R., Bronnenmeyer, T., Malz, E., Gros, S., and Diehl, M.: AWEbox: An Optimal Control Framework for Single- and Multi-Aircraft Airborne Wind Energy Systems, Energies, 16, https://doi.org/10.3390/EN16041900, 2023.

De Schutter, J., Mühleck, A., Leuthold, R., and Diehl, M.: Optimal control of crosswind kite systems with an engineering wake model based on vortex loops and dipoles, Wind Energy Science, 11, 1751–1770, https://doi.org/10.5194/wes-11-1751-2026, 2026.

DeVries, L. and Paley, D. A.: Wake estimation and optimal control for autonomous aircraft in formation flight, AIAA Guidance, Navigation, and Control (GNC) Conference, https://doi.org/10.2514/6.2013-4705, 2013.

Echeverri, P., Fricke, T., Homsy, G., and Tucker, N.: The Energy Kite: Selected Results From the Design, Development and Testing of Makani's Airborne Wind Turbines, Part 1/3, Makani Technologies, https://storage.googleapis.com/x-prod.appspot.com/files/Makani_TheEnergyKiteReport_Part1.pdf, 2020.

Gaunaa, M., Forsting, A. M., and Trevisi, F.: An engineering model for the induction of crosswind kite power systems, Journal of Physics: Conference Series, 1618, 032 010, https://doi.org/10.1088/1742-6596/1618/3/032010, 2020.

Goit, J. P. and Meyers, J.: Optimal control of energy extraction in wind-farm boundary layers, Journal of Fluid Mechanics, 768, 5–50, https://doi.org/10.1017/JFM.2015.70, 2015.

Goit, J. P., Munters, W., and Meyers, J.: Optimal Coordinated Control of Power Extraction in LES of a Wind Farm with Entrance Effects, Energies 2016, Vol. 9, Page 29, 9, 29, https://doi.org/10.3390/EN9010029, 2016.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

Gros, S. and Diehl, M.: Modeling of airborne wind energy systems in natural coordinates, in: Airborne Wind Energy, edited by Ahrens, U., Diehl, M., and Schmehl, R., chap. 10, pp. 181–203, Springer, Berlin, Heidelberg, https://doi.org/10.1007/978-3-642-39965-7_10, 2013.

Haas, T.: Simulation of Airborne Wind Energy Systems in the Atmospheric Boundary Layer, Ph.D. thesis, KU Leuven, https://lirias.kuleuven.be/retrieve/658988, 2022.

Haas, T. and Meyers, J.: Comparison study between wind turbine and power kite wakes, Journal of Physics: Conference Series, 854, 12 019, https://doi.org/10.1088/1742-6596/854/1/012019, 2017.

Haas, T., Schutter, J. D., Diehl, M., and Meyers, J.: Wake characteristics of pumping mode airborne wind energy systems, Journal of Physics: Conference Series, 1256, 012 016, https://doi.org/10.1088/1742-6596/1256/1/012016, 2019.

Haas, T., De Schutter, J., Diehl, M., and Meyers, J.: Large-eddy simulation of airborne wind energy farms, Wind Energy Science, 7, 1093–1135, https://doi.org/10.5194/wes-7-1093-2022, 2022.

Hansen, J., Kharecha, P., Morgan, D., and Vest, J.: Yes, 2026 is on Track to be the Hottest Year, techreport 2026.06.12, Climate Science, Awareness and Solutions Program. Columbia University., 475 Riverside Drive (Room 520-C) New York, NY 10115 USA, https://www.columbia.edu/~jeh1/mailings/2026/2026GlobalTemperature.2026.06.12.pdf, 2026.

Harzer, J., Schutter, J. D., and Diehl, M.: Integration Error Regularization in Direct Optimal Control using Embedded Runge Kutta Methods, in: 2025 European Control Conference (ECC), 2025.

Hogg, J. and Scott, J.: New Parallel Sparse Direct Solvers for Multicore Architectures, Algorithms, 6, 702–725, https://doi.org/10.3390/a6040702, 2013.

Horn, G., Gros, S., and Diehl, M.: Numerical Trajectory Optimization for Airborne Wind Energy Systems Described by High Fidelity Aircraft Models, in: Airborne Wind Energy, edited by Ahrens, U., Diehl, M., and Schmehl, R., pp. 205–218, Springer Berlin Heidelberg, Berlin, Heidelberg, ISBN 978-3-642-39965-7, https://doi.org/10.1007/978-3-642-39965-7_11, 2013.

Houska, B. and Diehl, M.: Optimal control for power generating kites, in: Proc. European Control Conference, pp. 1–14, Institute of Electrical and Electronics Engineers Inc., ISBN 9783952417386, https://doi.org/10.23919/ecc.2007.7068861, 2007.

Janssens, N. and Meyers, J.: Towards real-time optimal control of wind farms using large-eddy simulations, Wind Energy Science, 9, 65–95, https://doi.org/10.5194/wes-9-65-2024, 2024.

Karypis, G.: METIS: A Software Package for Partitioning Unstructured Graphs, Partitioning Meshes, and Computing Fill-Reducing Orderings of Sparse Matrices, University of Minnesota, https://papers.karypis.org/glaros/files/sw/metis/manual.pdf, 2013.

Karypis, G. and Kumar, V.: METIS: A Software Package for Partitioning Unstructured Graphs, Partitioning Meshes, and Computing Fill-Reducing Orderings of Sparse Matrices, techreport TR 97-061, University of Minnesota, Department of Computer Science/ Army HPC Research Center, Minneapolis, MN 55455 USA, https://hdl.handle.net/11299/215346, 1997.

Katz, J. and Plotkin, A.: Low-speed aerodynamics, Cambridge University Press, 2 edn., 2010.

Kawajir, Y., Margot, F., Laird, C., Vigerske, S., and Wächter, A.: Quasi-Newton Approximation of Second Derivatives, COIN-OR, https://coin-or.github.io/Ipopt/SPECIALS.html#QUASI_NEWTON.

Lee, S.-G., Yang, H.-H., Addo-Akoto, R., and Han, J.-H.: Transition Flight Trajectory Optimization for a Flapping-Wing Micro Air Vehicle with Unsteady Vortex-Lattice Method, Aerospace, 9, https://doi.org/10.3390/aerospace9110660, 2022.

Leishman, J. G.: Principles of Helicopter Aerodynamics, Cambridge University Press, 2 edn., ISBN 0521858607, 2006.

Leuthold, R., Gros, S., and Diehl, M.: Induction in Optimal Control of Multiple-Kite Airborne Wind Energy Systems, in: Proceedings of the 20th World Congress The International Federation of Automatic Control, 2017.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



Leuthold, R., Schutter, J. D., Malz, E., Licitra, G., Gros, S., and Diehl, M.: Operational Regions of a Multi-Kite AWE System, in: European Control Conference, 2018.

Leuthold, R., Crawford, C., Gros, S., and Diehl, M.: Engineering Wake Induction Model For Axisymmetric Multi-Kite Systems, Journal of Physics: Conference Series, 1256, 012 009, https://doi.org/10.1088/1742-6596/1256/1/012009, 2019.

Leuthold, R., Schutter, J. D., Crawford, C., Gros, S., and Diehl, M.: Trajectory Tracking in AWE Optimal Control with a Rigidly-Convected, Lifting-Line Vortex Model, in: Book of Abstracts of the Airborne Wind Energy Conference 2026, edited by Fontes, F. A. C. C., Thoms, S., and Schmehl, R., Universidade do Porto and Delft University of Technology, Porto, Portugal, ISBN 978-94-6186-846-6, https://doi.org/10.4233/uuid:20d74952-bf62-4b10-9df0-bfa1910da6d2, 2026.

Licitra, G.: Identification and Optimization of an Airborne Wind Energy System, phdthesis, University of Freiburg, https://d-nb.info/1164374370/34, 2018.

Licitra, G., Williams, P., Gillis, J., Ghandchi, S., Sieberling, S., Ruiterkamp, R., Diehl, M., V, A. P. B., Hague, T., Leuven, K. U., and Breisgau, F. I.: Aerodynamic Parameter Identification for an Airborne Wind Energy Pumping System, in: Proceedings of the 20th World Congress The International Federation of Automatic Control, vol. 50, pp. 11 951–11 958, Elsevier B.V., ISSN 24058963, https://doi.org/10.1016/j.ifacol.2017.08.1038, 2017.

Licitra, G., Koenemann, J., Bürger, A., Williams, P., Ruiterkamp, 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.

Malz, E., Koenemann, J., Sieberling, S., and Gros, S.: A reference model for airborne wind energy systems for optimization and control, Renewable Energy, 140, 1004–1011, https://doi.org/10.1016/j.renene.2019.03.111, 2019.

Malz, E. C.: Airborne Wind Energy – to fly or not to fly? A study on the power production of airborne wind energy systems and their integration in the electricity generation system, phdthesis, Chalmers University of Technology, 2020.

Mula, S. M., Stephenson, J. H., Tinney, C. E., and Sirohi, J.: Dynamical characteristics of the tip vortex from a four-bladed rotor in hover, Exp Fluids, 54, https://doi.org/10.1007/s00348-013-1600-9, 2013.

Munters, W.: Large-eddy simulation and optimal coordinated control of wind-farm boundary layers, Ph.D. thesis, KU Leuven, https://lirias.kuleuven.be/retrieve/481140, 2017.

Munters, W. and Meyers, J.: An optimal control framework for dynamic induction control of wind farms and their interaction with the atmospheric boundary layer, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 375, https://doi.org/10.1098/RSTA.2016.0100, 2017.

Munters, W. and Meyers, J.: Towards practical dynamic induction control of wind farms: Analysis of optimally controlled wind-farm boundary layers and sinusoidal induction control of first-row turbines, Wind Energy Science, 3, 409–425, https://doi.org/10.5194/WES-3-409-2018, 2018.

Nocedal, J. and Wright, S. J.: Numerical Optimization, Springer, 2 edn., ISBN 2006923897, 2000.

Paraschivoiu, I.: Double-Multiple Streamtube Model for Studying Vertical-Axis Wind Turbines, J. Propulsion, 4, 370–377, https://doi.org/10.2514/3.23076, 1988.

Payne, P. R. and McCutchen, C.: Self-erecting windmill, https://patents.google.com/patent/US3987987A/en, 1976.

Rancourt, D.: Method for the flight path optimization of the electric-powered reconfigurable rotor (EPR2) VTOL concept, phdthesis, Georgia Institute of Technology, https://repository.gatech.edu/entities/publication/20e01f81-2c74-412c-ae61-2ae1087514a9, 2016.

Rancourt, D. and Mavris, D. N.: Efficient Aerodynamic Method for Interacting Lifting Surfaces over Long Distances, Journal of Aircraft, 55, 2466–2475, https://doi.org/10.2514/1.C034905, 2018.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.

Raymer, D. P.: Aircraft Design: A Conceptual Approach, AIAA Education Series, American Institute of Aeronautics and Astronautics, 1992.

Saffman, P. G.: Vortex Dynamics, Cambridge Monographs on Mechanics and Applied Mathematics, Cambridge University Press, ISBN 0-521-477395, 1995.

Sant, T., Kuik, G. V., and Bussel, G. J. V.: Estimating the angle of attack from blade pressure measurements on the NREL phase VI rotor using a free wake vortex model: Axial conditions, Wind Energy, 9, 549–577, https://doi.org/10.1002/WE.201, 2006.

SCD Computational Mathematics Theme: Coin-HSL: A collection of HSL packages for use with IPOPT, https://licences.stfc.ac.uk/product/coin-hsl, 2024.

Sterle, A., Hans, C. A., and Raisch, J.: Model predictive control of wakes for wind farm power tracking, Journal of Physics: Conference Series, 2767, 032 005, https://doi.org/10.1088/1742-6596/2767/3/032005, 2024.

Tagg, A. C., Anderson, R., Joseph, C., and Ning, A.: Trajectory Optimization of eVTOL and Conventional Aircraft: A Comparative Analysis of Vortex Particle Method and Vortex Lattice + Blade Element Momentum Theory, in: AIAA AVIATION FORUM AND ASCEND 2024, https://doi.org/10.2514/6.2024-3587, 2024.

Trevisi, F., Castro-Fernández, I., Pasquinelli, G., Riboldi, C. E. D., and Croce, A.: Flight trajectory optimization of Fly-Gen airborne wind energy systems through a harmonic balance method, Wind Energy Science, 7, 2039–2058, https://doi.org/10.5194/wes-7-2039-2022, 2022.

Trevisi, F., Riboldi, C. E., and Croce, A.: Vortex model of the aerodynamic wake of airborne wind energy systems, Wind Energy Science, 8, 999–1016, https://doi.org/10.5194/WES-8-999-2023, 2023.

Trevisi, F., Sabug, L. J., and Fagiano, L.: A Gaussian wake model for Airborne Wind Energy Systems, IOP Science home Accessibility Help Journal of Physics: Conference Series, 3016, 2025.

Vali, M., Petrović, V., Boersma, S., van Wingerden, J.-W., Pao, L. Y., and Kühn, M.: Model Predictive Active Power Control of Waked Wind Farms, in: 2018 Annual American Control Conference (ACC), pp. 707–714, https://doi.org/10.23919/ACC.2018.8431391, 2018.

van den Broek, M. J., Tavernier, D. D., Sanderse, B., and van Wingerden, J.-W.: Adjoint Optimisation for Wind Farm Flow Control with a Free-Vortex Wake Model, Renewable Energy, 201, 752–765, https://arxiv.org/abs/2208.11516, 2022.

van den Broek, M. J., van den Berg, D., Sanderse, B., and van Wingerden, J. W.: Optimal Control for Wind Turbine Wake Mixing on Floating Platforms, IFAC-PapersOnLine, 56, 7656–7661, https://doi.org/10.1016/J.IFACOL.2023.10.1165, 2023.

van den Broek, M. J., Becker, M., Sanderse, B., and Wingerden, J.-W. V.: Dynamic wind farm flow control using free-vortex wake models, Wind Energ. Sci, 9, 721–740, https://doi.org/10.5194/wes-9-721-2024, 2024.

van der Wall, B. G. and Richard, H.: Analysis methodology for 3C-PIV data of rotary wing vortices, Experiments in Fluids, 40, 798–812, https://doi.org/10.1007/s00348-006-0117-x, 2006.

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, Annual Reviews in Control, 52, 330–357, https://doi.org/10.1016/J.ARCONTROL.2021.03.002, 2021.

Wikipedia contributors: Lockheed C-5 Galaxy, Wikipedia, The Free Encyclopedia, https://en.wikipedia.org/w/index.php?title=Lockheed_C-5_Galaxy&oldid=1259240605; accessed 27-November-2024, 2024.

Wächter, A. and Biegler, L. T.: On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming, Mathematical Programming, 106, 25–57, https://doi.org/10.1007/s10107-004-0559-y, 2006.

Zanon, M., Gros, S., Horn, G., and Diehl, M.: Airborne Wind Energy Based on Dual Airfoils, in: Airborne Wind Energy Conference 2013, http://homes.esat.kuleuven.be/~mzanon/, 2013.

https://doi.org/10.5194/wes-2026-211
Preprint. Discussion started: 12 November 2026
© Author(s) 2026. CC BY 4.0 License.



Zanon, M., Gros, S., Meyers, J., and Diehl, M.: Airborne Wind Energy: Airfoil-Airmass Interaction, in: Proceedings of the 19th World Congress The International Federation of Automatic Control, vol. 19, pp. 5814–5819, IFAC Secretariat, ISBN 9783902823625, ISSN 14746670, https://doi.org/10.3182/20140824-6-za-1003.00258, 2014.

Zheng, R. and Lyu, Y.: Nonlinear tight formation control of multiple UAVs based on model predictive control, Defence Technology, 25, 69–75, https://doi.org/10.1016/j.dt.2023.03.011, 2023.