Author's response to Anonymous Reviewer #2

Jochem De Schutter, Antonia Mühleck, Rachel Leuthold, Moritz Diehl

April 11, 2026

Dear Reviewer,

Thank you for the time and effort you invested in providing detailed feedback and thoughtful comments. We believe these have significantly strengthened the manuscript. In the following, we address each comment point by point and indicate the corresponding revisions made in the manuscript. When referring to Figure and Table numbers, we use the numbering in the updated manuscript.

Kind regards,
Jochem De Schutter on behalf of the authors

Major Comment 1: Validation of the simplified wake model within the optimal control framework
The continuous vortex-loop wake model is convincingly validated against DUST. However, the optimal control problem employs a simplified hybrid wake model with reduced window-of influence assumptions and finite wake history. While such simplifications are understandable for computational tractability, it remains unclear to what extent the reduced model preserves the quantitative accuracy demonstrated for the full formulation. Although Figure 12 compares the discretised wake against an analytical vortex-loop solution, it is not evident how the reduced model compares to the DUST-validated formulation. Since the dual-kite optimisation results rely on the simplified wake model, it would strengthen the manuscript to clarify this validation chain, for example through a direct comparison with DUST for the reference solution or by quantifying the modelling error introduced by the OCP simplifications in the reference solution .

Response:
Thank you for this comment. In the previous version of the manuscript, the validation chain consisted of three steps:

1. Validate the analytic vortex-loop model with DUST simulations.
2. Validate a reference hybrid model transcription with a large window of influence to the analytic vortex-loop model.

3. Discretize more coarsely, choose a smaller window of influence and compare the change in optimal cost and solution compared to when using the reference transcription.

We believe your comment relates to Step 2, which we agree is not sufficiently elaborated on. This step was previously somewhat "hidden" in Fig. 12 and we agree that it would also benefit from a direct comparison of the reference transcription to the DUST simulation. We have added a comparison of the resultant aerodynamic forces prominently and with accompanying Fig. 9 in Section 6.1, showing very good agreement.

Major Comment 2: Representativeness of the case study and interpretation of wake effects

The numerical case study considers a configuration with fixed tether length and without reeling dynamics and with a very optimistic mass scaling (see https://doi.org/10.5194/wes-7-1847 2022). This appears to represent either a fly-gen system without explicitly accounting for the additional mass and drag of onboard turbines, or a ground-gen system without accounting for its reel-out operation. In both cases, the resulting flight speed and loading would be reduced in real operation, which directly affects period duration and wake strength. The analysed dual-kite configuration therefore seems close to a worst-case scenario in terms of wake sensitivity. While this does not invalidate the results, it affects how broadly they can be generalised. A more explicit discussion of the operational interpretation and limitations of the chosen case study would help readers assess the scope and transferability of the conclusions to other systems.

Response:

The case study was chosen as a dual-wing "ship-towing" or "airborne actuator" case intentionally, to indeed create a "worst-case" scenario for induction. The reason for this is was to put our model to the test. The model is able to capture the induced velocity field in great detail when the kites are flying very close to their own wake (very high effective gliding ratio). For this, the vortex-loop model is necessary. For the significantly lower effective gliding ratios that will likely be encountered in realistic airborne wind energy operation, the distance to the wake will increase. In terms of validity, this does not pose a problem: the vortex-loop model does not become less accurate at larger distances. Rather the opposite is true: the vortex-dipole model becomes more accurate. This is also clear from the newly added comparison in Fig. 9: the hybrid model, with only one layer of vortex-loops and further only vortex dipoles, gives almost a perfect fit compared to the full vortex-loop model in this extreme scenario.

In terms of computational expense, a lower effective gliding ratio mainly makes the OCP cheaper: the wake is convected faster, creating a larger distance, up to a point wherer one can use only the cheaper vortex-dipole model. Because of the large wake convection distances, one can also cut off the wake history sooner. This we point we had already made in the conclusion.

Finally, we have also added a comment on the mass-scaling, which in this case is indeed very optimistic considering the high loading. However, for the dual-kite case, where the kites are balancing each other, the influence of the mass on the optimal flight trajectory is typically much lower compared to the single-kite case. Hence, we deem the model acceptable for the validation purposes of this paper.

Major Comment 3: Half-periodicity assumption for the dual-kite problem
The manuscript states that, because the two airfoils follow nearly identical trajectories over a full orbit, the dynamics can be integrated over half an orbit by enforcing half-periodicity (line 411), following Zanon et al. (2013). This reduction appears insufficiently justified for the present setting. For a helical trajectory, the "up-loop" and "down-loop" segments typically experience different apparent wind conditions, speeds, and load distributions due to gravity, so the orbit need not exhibit a strict half-period symmetry. This makes the choice of the two half-orbit boundary points (where the state matching is imposed) critical, and it is not clear from the text where along the orbit the half horizon starts/ends, nor under which assumptions half-periodicity is expected to hold. Related optimal-control studies (eg. https://doi.org/10.3390/en16041900) explore dual-aircraft configurations without taking only half-period and the resulting trajectories are not generally symmetric. This suggests that enforcing half-periodicity may exclude physically relevant asymmetric solutions unless the symmetry assumptions are carefully justified. I recommend that the authors (i) clearly define the half-orbit start/end points along the trajectory, (ii) explain the conditions under which the half-periodicity constraint is valid for their chosen path class and (iii) ideally demonstrate that solving the full-orbit problem does not yield materially different (asymmetric) optimal solutions for at least one representative case.

Response:
Thank you for this comment. We agree that this step has received insufficient attention in the manuscript. The main issue here is the meaning of the term "half-periodicity". This does not mean that only one half of the loop is simulated and the other half is chosen symmetric to it. Rather, we exploit the fact that we have a two-kite system in combination with a single-loop trajectory (typical case for all non-reeling systems). In this case, while one wing flies the upstroke, the other simultaneously flies the downstroke. After half a period, the roles are reversed and the first wing flies the downstroke and the second one the upstroke. Thus, by implementing this "role reversal" as a constraint instead of imposing full periodicity, we avoid simulating the same up- and downstroke twice. We have merely imposed that both wings fly the same downstroke and upstroke paths, which can be asymmetrical (this can be seen, e.g. in the force profiles). This is of course crucial, since our central claim is that we capture the non-axisymmetric nature of kite system flight trajectories. The start- and endpoints can be chosen freely by the solver, and do not influence the solution, as the flight loop is simulated completely for all possible choices. We have

elaborated on this matter in the text and have explicitly stated the new "role reversal" constraints.
In the conclusion we had already raised the point that allowing differing up- and downstrokes ("multi-loop") for the two wings might lead to a different, more performant solution, as the wings would in this case be able to adjust their flight path to "evade" the trailing wake of the other one. While this is a very interesting research path, such an investigation falls outside of the scope of this paper.

Minor comments

Comment 1.
Line 14-15: Specify what metric are you using for the accuracy and comparison with reference solution. It should be clear by only reading the abstract.

Response:
The metric is the combined error of optimal cost and design. We have added this to the abstract.

Comment 2.
Line 80: "To limit the scope of this paper, we consider a kite system with fixed tether length, thereby excluding pumping-style AWE systems, which rely on tether reeling." Also mention that no onboard power generators are included, such that they purely act as actuators

Response:
We have added this information to the text.

Comment 3.
Line 81: "However, the developed methods in this paper can be directly applied to these type of systems as well." The present method assumes a fully developed wake, whilst in pumping systems wake development becomes increasingly important, particularly if the amount of flown loops is small. Please justify further how these methods could be applied to pumping systems or remove this phrase.

Response:
The proposed unsteady model can be used for time-domain simulations and does not in itself assume a fully developed wake. Because of its unsteady nature, it is in fact very much suited for modeling pumping-style wake patterns. However, when used in a periodic OCP, we do assume a fully developed periodic wake. For a pumping system, this just means that the pumping wake pattern (including build-up during reel-out as well as the reel-in phase) is repeated indefinitely when going down in the wake's history. This does not pose a limitation to the proposed method's applicability. We have elaborated on this in the text.

Comment 4.
Figure 1. Add a label to the methods (a,b). There are some symbols that are only introduced later in the text so the reference should indicate a) or b) as well.

Response:
We have added labels to the figures and when referencing to them in the text.

Comment 5.
Line 228: "Recall that for each wing we ignore the self-induced velocity for the first half rotation interval with length T." It says recall while it is the first time a half-rotation is mentioned.

Response:
Thank you for this comment, we have corrected the text.

Comment 6.
Line 250: additional clarification of the DUST simulation setup would improve transparency. It is not entirely clear whether the wing kinematics (positions and velocities) are prescribed and the pitch angle is chosen to match a target C_L without accounting for induction effects, or whether the induced velocities computed by DUST are already incorporated in a fully coupled manner such that the local angle of attack and resulting aerodynamic loads are determined self-consistently.

Response:
The first option is the case: the wing kinematics are prescribed, with the addition that the chosen pitch angle takes into account the induction effects predicted by the proposed wake model. We have clarified in the text.

Comment 7.
Figure 4 is somewhat difficult to interpret in its current form. The readability would improve if the flown trajectories of the kites were explicitly included in the visualisation. In addition, it would be helpful to clearly distinguish in the figure or caption which results correspond to the DUST simulation and which to the proposed vortex wake model.

Response:
We have added the kite trajectory to the plot, added more information to the legend and the caption so that all visual information is clearly defined.

Comment 8.
Lines 275-280: It would be helpful for the reader to understand how to qualitatively asses over or underestimation of convection speed in Figure 4 (eg. number of tip vortices seen in DUST vs vortex method)

Response:

We have added an explanation in this paragraph. Indeed we rely on looking at the location of the tip vortices as they cross the xz-plane.

Comment 9.
Figures 6,7: it would be helpful to mark the positions where the wing crosses this plane

Response:
We have added the full flight trajectories as well as the wing position and orientation at the time of the snapshot.

Comment 10
Figure 9. The time intervals N are difficult to see.

Response:
We have updated the plot so that the intervals (now correctly labeled N_{w}) should be clearly visible.

Comment 11.
Subsection 6.2. the discussion of discretisation sensitivity is not entirely clear. In particular, the relationship between the number of collocation intervals N (global time discretisation) and the influence window parameter N_f would benefit from clarification. Please explicitly state the value of N used and clarify how it relates to the window size $2N_f + 1$.

Response:
Thank you for this comment. We had introduced some confusion by equating the number of intervals used by the finite window approach with the number of OCP intervals. There is no fundamental relationship between the two or between the parameters N and N_{f}. They can be chosen completely independently. We now explicitly differentiate between the number of intervals for the window approach N_{w} and the number of OCP intervals N. We have added the values used for both parameters in Table 3.