Answer to reviews: Integer programming for optimal yaw control of wind farms

We would like to thank both reviewers for their detailed and constructive remarks. We revised the manuscript carefully to accommodate as many suggestions and requests as possible, while keeping an eye on the page count. Please see below for our detailed replies to the review reports.

Superordinate changes

As we have made many changes to figures and tables that cannot be meaningfully represented by the requested comparison via latexdiff, we list the fundamental changes here (apart from updated results from computational experiments in tables and updated visualizations):

- updated Fig. 3: thinned out by removing "downstream sections"
- added Fig. 6: overall process
- added Table 3: comparison of simulations with WinFaST and FLORIS
- added Table 4: comparison of full enumeration, serial-refine, and covering approach
- added Fig. 7: comparison of baseline and optimized results (each with WinFaST and FLORIS simulation)
- previous Table 3 is now Table 5 in the revision
- previous Table 4: removed (as the power outputs of individual turbines are not required)
- previous Fig. 6: removed (a comparison of baseline and optimized result is now in Fig. 7 (for another experiment); the optimized result of the previous experiment (case 1.4) is now in Fig. 9 (b) (as case 3.2 repeats case 1.4) instead of Fig. 6 (b))
- previous Table 5 is now Table 6 in the revision
- previous Fig. 7 is now Fig. 8 in the revision
- previous Table 6 is now Table 7
- previous Fig. 8 is now Fig. 9
- added Table 8: results of series 4

The following changes have been made to the structure:

- Sect. 1.3 (Paper outline): it is now a paragraph instead a separate section; we moved it to the end of Sect. 1 (before Sect. 1.1, cf. line 28 f. beginning with "The remaining paper is structured as follows: ...") to introduce abbreviations (WFYP, CA, IP, WT) early on.
- Sects. 2.1.1 and 2.1.2 have been merged into Sect. 2.1.

Answers to Review 1

- "In Sec. 2.3.3 it is mentioned that this new formulation of the yaw optimization problem based on this "covering approach" is equivalent..."—As announced in the forum answer, we rephrased the sentence to make clear that a marginal difference is possible. In particular, we added that equivalence needs a larger upstream section than used in practice ("...strictly speaking, this is only true if..." in line 436). We already discussed the size of the upstream section in Sect. 4.1, which is now in Sect. 4 (extended by comparisons with full enumeration and serial-refine method), cf. line 532 f.: "The deviations ... are due to the size of the upstream section ...". In addition, this dependency is now emphasized by naming it as an assumption in both Abstract and Concluding remarks (cf. lines 6 and 630 f.) as suggested in the last advice of this review.

  To your objection: "However, the exact optimal solution cannot be in principle guaranteed (especially is we decrease the number of yaw values allowed)."—Yes and no: on the one hand the expression *the same problem* in Sect. 2.3.3 refers to problems (2) to (4) and (8) to (11), which are already formulated with discrete yaw offsets (so, a finer discretization is only a practical run time problem, not a theoretical one); on the other hand the global optimal solution of the *discretized* yaw offset problem can indeed be different from the global optimal solution of the *continuous* formulated problem; however, as it is now clearly stated in the Abstract (and the Concluding remarks) that we use discretized yaw offsets, a comparison to the continuous yaw offset problem seems to be unnecessarily complicated to us in Sect. 2.3.3.

- "Moreover, it would be nice to highlight qualitatively or quantitatively the advantages..."—As announced in the forum answer, we incorporated comparison computations in Sect. 4, see Table 4, namely full enumeration, serial-refine, and covering approach (in each case with FLORIS simulation), which actually result in the same improvement of the total power output, cf. the paragraph "For the mentioned CA validation..." beginning in line 523. Further, we would like to present a case where our covering approach is more advantageous in terms of accuracy than serial-refine, but we have not yet found such an example, see the paragraph "We would like to present a case... In theory, there could be such a case as our CA is equivalent to full enumeration for sufficiently large upstream sections, see Sect. 2.3.3. ..." beginning in line 541. We also mention there that the serial-refine method is more advantageous in terms of run time. (In addition, we included the serial-refine method in Sect. 1.1 (Related work), cf. line 101 f.) Finally, we do not yet know whether the theoretical advantage of our covering approach will pay off in practice.
- "In most of the figures ...a vector indicating the wind direction ..."—realized.
- "A table summarizing the results presented in Sec. 4.4..."—realized (by adding Table 8 and cross-referencing in the text, cf. line 610).
- "The presence of some mild assumption is mentioned both in the Abstract and the Conclusion. ..."—realized by mentioning the most important ones, namely "(under some mild assumptions like discretized yaw offsets, chosen size of upstream section, and homogeneous layout structure)" in line 6 and "like discretized yaw offsets, chosen size of so-called upstream section, and homogeneous layout structure" in line 630 f.

Answers to Major Comments of Review 2

- 1. a. Reproducibility:
  - i. "There is insufficient description of the overall process, and the paper would benefit from a flow-chart or pseudo-code of the entire process."—We added Fig. 6 in Sect. 4 to describe the overall process. There is not enough space for detailed pseudo-code. We think that it is crucial to describe the order in which the individual steps are performed as this is actually a combination of the covering approach, see Sect. 2.2, its formulation as integer program, see Sect. 2.3, and simulation, see Sect. 3. This is now also described in the first paragraph of Sect. 4, cf. line 497 f.
  - ii. "How is the constraint of having consistent yaw offset angles..."—As announced in the forum answer, it is done in two steps: first, we compute the valid yaw configurations $\tilde{L}_{k+1,\ell_k}$ relative to the upstream section S_k and the chosen yaw configuration ℓ_k, which is done sequentially; second, the compatibility in the wind farm is finally ensured in Eq. (10) (previously Eq. (9)), which is not sequentially. In particular, we reformulated two sentences in Sect. 2.3.2 (cf. line 393 ff.: "In fact, it suffices to enforce these conditions...Then, establishing consistency of the respective overlaps...") by adding the phrases "(which explicitly excludes an arbitrary order)" and "which has to overcome the sequential order and is realized in Eq. (10)". By the way, we have not emphasized the sequential computation of the first part additionally as it is already there with the phrase "in ascending sequence", cf. line 395 (previously line 410). As announced in the forum answer, we omit the very detailed answer (given there regarding the order of the sections and the constructed example that violates the requirements) to not digress and distract from the main thread.
- 1. b. Solver: "...descriptions of both the SCIP and Gurobi solvers..."—As explained in the forum answer, it is not possible to answer this question (in short: as we cannot access the source code of Gurobi), so there were no changes: in Sect. 4 we already stated that SCIP is an open-source academic software and Gurobi is proprietary, cf. line 501 ff. (previously line 524 ff.).
- 1. c. Trapezoidal Sections:

  i. and ii. "sensitivity studies on the trapezoid slope" and "multiple yaw configurations"—As outlined in the forum answer, we already have carefully formulated that the used method does not guarantee the threshold (using the words "based on" and "simply" in Sect. 2.2.1, cf. lines 269 and 270, previously line 281). To emphasize this we added the phrase "(without guaranteeing it)", cf. line 270. Moreover, we provide more details on the method with the phrase "(from simulations with 0° and extreme yaw offsets, e.g., ±15°)", see line 270 f., and the placement in the overall process in Fig. 6. On the one hand the added comparison (with FLORIS as underlying simulation software) of our covering approach with full enumeration and serial-refine in Sect. 4 supports that the size of the upstream section was chosen reasonably, on the other hand we discuss in the revision their impact and the resulting model inaccuracy in detail there: "The deviations of the predicted power by CA and full farm simulation (in baseline and/or optimization) are limited and are due to the size of the upstream section, ...",

cf. line 532 ff., which also refers to the first comment of review 1. Further, we also list the size of the upstream section in Abstract and Conclusion among the assumptions, cf. line 6 and 630 f. (here, again we refer to the answers to review 1, namely the first and the last one).

A very detailed answer to the trapezoid is in the forum answer. In addition, we checked—just to report it here—whether a turbine with a different (lower) wind speed could need a different (larger) section (as downstream turbines can have a decreased speed): to that end, we used the simulation software FLORIS and set turbines to positions $(0, 0)$ (for WT 1), $(3D, 0.5D + 0.15 \cdot 3D)$ (for WT 2), and $(3D, -0.5D - 0.15 \cdot 3D)$ (for WT 3), remembering that the slope was at least 0.15, i.e., we placed them on the edge of the trapezoid with a distance of $3D$ in wind direction (which is from west to east)—again, as explained in the forum answer, if the distance is $3D$ in wind direction, turbines are directly behind each other with the chosen spacing of $5D \times 3D$, i.e., in fact, this placement is worse than the worst case scenario. As a result WTs 2 and 3 reach (compared to WT 1) 88% power at 4 m/s, 92% at 6 m/s, 94% at 10 m/s. Finally, there is an effect but as we already described to accept small model inaccuracies, we think that this is not worth a discussion in the manuscript.

Again, the choice of the size of the upstream section is not the main aspect in the paper, so we have not expanded the discussion there, but generally emphasized the impact of its choice in the revision.

iii. "choosing to use the "upstream" sections"—As announced in the forum answer, we only focus on the used "upstream sections" (to avoid confusion) in the revised manuscript, see Sect. 2.2.1 (Upstream sections), i.e., we removed "downstream section" definition as well as their visualizations.

iv. "... how well ... not laid on a grid pattern"—This comment has not led to any changes. This was discussed in detail in the forum answer (in short: it is possible to "choose the grid resolution as fine as needed to allow representing any layout", but may of course be detrimental to the run time).

- 1. d. "Combinations: I believe the formula on line 355 should be $2^{n_{\mathrm{WT,u}}} - 1 \ldots$'"—We corrected the formula, cf. now line 342, and afterwards (cf. line 344) as well as the first occurring in Sect. 2.2.2 (cf. line 297).

- 1. e. "Loading Components"—In contrast to the submission, we now present experiments with weightings of the tower and pitch activity penalties, namely in series 0, case 1, part e), cf. Table 4, (to show the effect) and at the end of Sect. 4.3 (series 3) (at 15 m/s where the maximum total output is already reached in the baseline), cf. line 604 ff.

- 2. a. "Farm-Wide Optimisation"—Indeed, the previous line 545 (in Sect. 4.1) was probably misleading. We moved the discussion to Sect. 4 and now distinguish between "predicted power" and "full farm simulation". Again, this is the already mentioned paragraph beginning with "The deviations of the predicted power by CA and full farm simulation ..." (cf. line 532 ff.); in particular, we describe that we "run full simulation with optimized yaw offsets at the end" in line 537. Moreover, in the revision we actually compare our covering approach with a farm-wide optimization (full enumeration), see Table 4, as announced in the forum answer, an additionally with the serial-refine method (with FLORIS as underlying simulation); so, finally, see Table 4 parts a) to c) as well as the already mentioned paragraph beginning in line 523 ("For the mentioned CA validation ...").

- 2. b. "Timings"—The run times of the methods mentioned in 2. a. with FLORIS as underlying simulation partially address this question. In addition, for simulation with WinFaST, we now predict the run times of full enumeration for series 3, see Table 6, using the number of combinations to answer the actual question.

- 3. a. ""Deeper" Wind Farms"—A detailed explanation of the corresponding idea is in the forum answer. We added a short explanation in Sect. 1.2 ("...scaling in wind direction is possible by the following idea: ...", cf. line 130 ff.).

- 3. b. Steady-State Models:

  - i. As in the answer to review 1, we used the software FLORIS for a comparison with steady-state (see also the answers to 2. a. and b. with regard to the run time): the results are in Table 4, which contains a comparison of the simulations FLORIS and WinFaST. The results suggest that a steady-state replacement of the dynamic model is not sufficient. So, we deleted the guess that a static model would suffice, namely "a static model would suffice", cf. previously line 40 (in Sect. 1.1), and "dynamic models (for wakes and turbines) possibly made the precomputation more expensive than necessary", cf. previously line 633 ff. (in the conclusion). Instead, we close the simulation discussion in Sect. 4 with (cf. line 522): "A comprehensive comparison of the simulations (and which one is preferable) is out of scope of the present paper." Indeed, the steady-state cuts down precomputation time, but we cannot judge the practical suitability.

  - ii. "sufficient time for changes to propagate through the farm"—Indeed, your first question pointed to what turned out to be a mistake in our data generation, which we fixed and subsequently ran *all* simulations again (resulting essentially in updates in Tables, Figures, and text in Sect. 4 and its subsections, but also, e.g., in an updated footnote on page 16). In more detail, our forum answer argued with the initial simulation time; however, as we recycled the steady state simulation for several yaw configurations, indeed, there was not enough time in the overall simulation for changes to propagate through the farm. So, a revision was really necessary. We have now determined the begin of the observation time interval individually (i.e., for each wind farm and wind condition scenario), see Sect. 3.2 for a detailed workflow ("In practice, we round the minimal wind speed...", cf. line 469 ff.) and Table 3 for the results. Moreover, the duration of the observation time interval is now 10 min (instead of 5 min): "(to obtain roughly the specified wind speed on a mean at WT 3...)", cf. line 472.

    The question regarding the steady-state replacement is answered in i.

- 4. "Paper Length"—On the one hand we incorporated most of the proposed suggestions to shorten the manuscript; in particular, we deleted figures and tables as listed above as superordinate changes, we removed the suggested sentence "Furthermore, in the case of specific atmospheric conditions..." (cf. previously line 48) as well as the wind rose explanation (cf. previously line 375), we shortened or removed the summaries at the beginning of the sections, we shortened the appendix (but have left it as an appendix); moreover, we have shortened wherever possible without compromising clarity. On the other hand the addressing of the suggested and valuable improvements required space. So, all in all, we end up with 33 pages again.
- 5. "Contributions"—We reformulated Abstract and Conclusion to highlight: "...that covers the farm by patterns ...for faster ..." (cf. line 4 ff.) and "...that the farm can be covered py patterns ...to solve it faster" (cf. line 624 f.).
- 6. a. "Table 2: ...A diagram of the example described in the text of the chosen consistent yaw offsets may help."—We agree that a diagram would be a good idea, but did not include it in the revision due to space concerns.
- 6. b. "Figure 6"—As announced in the forum answer we understood this remark (to visualize wakes instead of values) also with regard to other illustrations and improved the presentation of all corresponding figures; in particular, we now use a visualization with the software FLORIS, see Figures 7 to 9.

Answers to Minor Comments of Review 2

- 1. "acronyms"—Acronyms are now defined only once (most of them in Sect. 1 at the paper overview, cf. line 28 ff., namely WFYP, CA, IP, WT, and TI) and then used consistently with few exceptions, e.g., in headings, the beginning of the conclusion, and the word "turbine" for WT. In addition, SC is only defined and used in Fig. 5 as not to overload the reader.
- 2. "Do not need to capitalise words after a colon mid-sentence."—done.
- 3. "Please ensure that the angle is referred to as..."—We now consistently write "yaw offset" instead of "yaw angle". Further, we replaced "nacelle orientation towards the wind" by "nacelle direction relative to the wind" (cf. previously lines 15 and 24, now only in line 15).
- 4. As suggested, we have replaced "in form" with "in the form" (cf. previously line 4, now line 5).
- 5. As proposed, we have replaced "green house" with "greenhouse" (cf. previously line 12, now line 13).
- 6. "Calling the total power the "most important objective" is subjective, ..."—We now describe it more carefully as "a primary objective" (cf. line 25 in previous and revised version).
- 7. "Section 1: The description of wakes and yaw control is not very clear, ..."—To address this we have added that yaw offsets "deflect and deform wakes" (cf. line 24).
- 8. "Figure 1"—We realized the suggestions, i.e., we added axis titles and a color bar title (including units) and made the line, that represents the rotor, thicker and black.
- 9. "Section 1.1: There are other papers that have "split up" wind farms into ..."—We included the mentioned papers in Sect. 1.1, namely Bernardoni et al. 2021 Journal of Renewable and Sustainable Energy and Dong & Zhao 2023 IEEE Transactions on Industrial Informatics, cf. lines 88 f.
- 10. "Line 45: Could the reference to a figure in a different paper be removed and the concept explained instead."—We removed the reference to the 3rd party figure and moved the wind turbulence discussion to Sect. 1.2.1 (Assumptions), see the paragraph in line 162 ff. ("In general, wind turbulence depends..."). There, we now refer to Fig. 1 "for an impression of the locally strong speed fluctuations".

- 11. “accuracy of 1%”—The details are in the forum answer. We reformulated this phrase using the so-called *wind farm efficiency* (cf. line 85 f.): “. . .reduces the gap between its so-called wind farm efficiency and its globally optimal quantity (obtained by full enumeration) to up to 1%.” We think that this makes the meaning clear enough (for the section on related work, i.e., Sect. 1.1), even if we do not define the term “wind farm efficiency” (due to space).

- 12. “Line 110: “as a wind farm yaw problem””—We added the missing “a”, see line 111 now.

- 13. “Line 113: “we provide to include””—As suggested we removed “provide to”, see line 113 now.

- 14. “Line 139: “respectively””—We removed the word, cf. line 141 now.

- 15. “Lines 150 & 279/280: The phrasing is difficult to understand, suggest changing to “...downstream turbine is 5% or less””—We realized this by using the phrase “is 5% or less” in line 152 and the phrase “is at a threshold or below—throughout this paper, 5%, . . .” in line 268 f.

- 16. “Lines 178, 346 & 534: . . .”—We have replaced “we remind that” by “we remind the reader that” (cf. lines 267 and 559) or “remembering that” (cf. line 333).

- 17. “Line 180: Is there a reference or calculation for the 15 minute time scale. . .”—We answered in the forum that it was an example to make the order of magnitude clear (indeed, we had the propagation time in mind). We have removed the 15 min time scale and instead discuss the available run time in detail in the rewritten last paragraph of Sect. 1.2.2, line 179 ff.: in particular, we mention a time indication from Fleming et al. (2022) (seconds to minutes) and argue why we assume that a run time of less than one minute is sufficient (the already mentioned mechanical loads and gyroscopic moments and—as new aspect—a possible time advantage due to a LIDAR system). However, we now also clarify that we do not consider the transient phase as “the practical realization of a short real-time control scale would require a time-delayed yaw offset adaption of each turbine”.

- 18. “Section 2.1.2: The “curse of dimensionality” is well known. . .”—As announced in the forum answer, we merged Sections 2.1.1 and 2.1.2 and shortened wherever possible. However, the content of the previous section 2.1.2 cannot be completely omitted: e.g., we refer to “Example 2.2 (A 3×2 farm)” (cf. line 240) several times.

- 19. “Line 234: Surely standard IP solvers can handle black-box problems?”—As we have explained in detail in the forum answer, they do, in fact, not. So, this comment has not led to any changes.

- 20. “Line 262: . . .”—We corrected the typo (by replacing “preservers” with “preserves” in line 255).

- 21. “Line 270: . . .”—We corrected the typo (by replacing “reoccuring patterns” with “recurring patterns” in line 260).

- 22. “Figure 5: It would be helpful to see the effective farm depth and the section depths, . . .Also, the double meaning of plot 5c is quite confusing . . .”—We realized both improvement suggestions: first, we have drawn the “farm depth” and the “effective farm depth” (i.e., the section depth) and second, updated Fig. 5 (c) such that it is sufficient to specify only the wind direction in the description.

- 23. “Line 465: What does “generally possible” mean here? . . .”—We have reformulated this to (cf. line 449): “Axial induction and yaw offsets can be set time-dependent.”

- 24. “Section 3.2: Could you clarify how the different time scales are related . . .?”—We remind the reader that we updated some time scales in the revision as discussed in the answers to major comment 3. b. ii. and minor comment 17. In particular, we now use individually observation time intervals. Therefore, we do not include them in the overall process, Fig. 5, but refer to Table 3 for them. In addition, we reworked the (previously confusing) first paragraph of Sect. 3.2 to describe the simulation time interval as well as the observation time interval and their choice: in particular, we make clear that the observation time beginning takes the propagation time robustly into account and that the duration of the observation time interval is always 10 min (instead of 5 min before).

- 25. “Line 495: Equation should have a reference number.”—We added a reference number, cf. Eq. (12). In addition, Eq. (1) has a reference number now.

- 26. “Line 505: Are the turbines individually weighted in the reward? . . .”—We do not use individual weights in the computational results, but we implemented them. For clarity of notation, individual weights are not shown. We reformulated (cf. line 492 f.): “In fact, they could even do this individually for each turbine.”

- 27. “Section 4.1: There is not enough discussion on the effect of wind directions here given it is the section title.”—The used section titles are based on the experimental series. We reworked Sect. 4.1 but there is finally not enough space for a really detailed discussion.
- 28. “Line 593: It should be acknowledged that the turbines do not have a wide choice of yaw offset angles, so it is not too surprising that many are identical.”—We included this in Sect. 4.3 (cf. line 597): “Further, we observe . . .whereby on the one hand we also have to acknowledge that our number of admissible yaw offsets is not large, on the other hand we refer to Sect. 4.4 for the effect and question of usefulness of a finer discretization.”
- 29. “Line 599: For certain wind directions (rather than speeds) there may be no / very few turbine-wake interactions so the optimal may be 0° for all / most turbines.”—Yes, we should not generalize this. We carefully reformulated this part (cf. 600 ff.) to: “In addition, the optimal yaw offsets of all controlled WTs differ from 0°: on the one hand this may provide opportunities to reduce the running times of both precomputations and IP solving, on the other hand it might be different for other wind directions, although there are already relatively few wake interactions for the wind direction of 20° (which is reflected in the relatively small potential for total power improvement).”
- 30. “Line 625: Months of pre-computation would be a very long time, [. . .]. The authors mention that using a dynamic model was perhaps not needed, would using a steady-state model instead cut down the expected pre-computation times?”—This comment is addressed in the answer to major comment 3. b. i.