

Reply to Reviewers' Comments for wes-2025-67

November 15, 2025

The authors would like to thank the two reviewers for their time and for the useful feedback. All the inputs that they provided have contributed to the improvement of the paper.

A list of point-by-point replies to the reviewers' comments is detailed in the following. We have taken the opportunity of this deep revision to improve readability and understanding by expanding the text in various places and by adding some bibliographical references.

A revised version of the manuscript is attached to the present reply, with additions highlighted in [blue](#) and corrections/deletions marked in [red](#).

The authors,
November 15, 2025

Reviewer 1

The paper "Economic lifetime-aware wind farm control" from Anand et al. presents an approach to optimizing wind farm control setpoints for profit optimization, considering component fatigue and related maintenance costs. The Approach is demonstrated in simplified case studies, showcasing the optimization using only yaw steering as a control mode.

1.1. While the paper is well written overall, it could benefit from some restructuring to improve understandability. I suggest a more rigorous differentiation between methodology, case study definition, and results. At times, these aspects are mixed.

We thank the reviewer for this remark. We agree that the previous version mixed methods with case studies and results. We have now deeply revised the organization of all material, moving what was previously in the methods section into the results or in the appendix.

1.2. It would be helpful to differentiate between the general discussion of control approaches (induction, derating, boosting) and the ones used in this work. I was confused about which ones are actually implemented in this work.

We have clarified in various places that the method is general, but the applications considered here only use wake steering.

1.3. The introduction could benefit from a more in-depth discussion of previous work on the topic and the relation to / contribution of the present work. Relevant references include:

<https://wes.copernicus.org/articles/8/1727/2023/>
<http://dx.doi.org/10.18419/opus-13959>

We have significantly expanded the introduction. We have also added the suggested references to the bibliography, and explained their differences with respect to the present work.

1.4. Line 50: Control modes could also include de-rating, boosting etc.

Please see our response to question 1.2.

1.5. Line 61: Isn't economic optimization equivalent to production maximization if damage is not a major concern (except in negative price scenarios)?

Indeed, when component damage is completely disregarded over the operation lifetime, production maximization is equivalent to economic optimization, except in case of negative price scenarios. However, this is only the optimization of revenue, which is a simplification: damage can increase the costs of operation and maintenance and, by reducing lifetime, will also affect power production and therefore revenue. In the end, what counts is profit, not revenue. Removing this limitation is one of the key innovations of this paper. We have rephrased the text in order to make this point clearer.

1.6. Section 2: It is not entirely clear to me if lifetime is subject to optimization or a prescribed value.

We have more clearly explained that the lifetime duration is a prescribed value. The entire section on the formulation has been deeply rewritten and greatly expanded. We hope that the new version presents the concept in a clearer way than before.

1.7. Line 120: How is lost lifetime cost defined? What does chosen lifetime mean here? Isn't that an optimization parameter?

The new writing hopefully clarifies these points. In a nutshell, the user specifies a goal desired period of time when the farm can operate without exceeding the damage limit of its components. If the farm reaches this time having still a damage margin, then the farm can keep running for longer. On the other hand, if the damage limit is met before the end of the desired time period, then a fraction of the desired lifetime duration has been lost. The reduced lifetime will incur in a loss of profit, whereas the increased lifetime will generate an increase of profit. An extensive description of these situations has now been provided.

1.8. Section 3: Does the optimization take place over a (conditional) pdf of wind conditions with global profit objectives, or over a time series of input conditions sampled from a distribution.

Thank you for this comment, indeed the previous version was quite unclear on this important point. We have greatly expanded Sect. 2 to more precisely explain the formulation. In short, the problem is first posed as an optimal control problem in the time domain. Then we explain that the direct solution of such a problem in time would be unfeasible because of the resulting overwhelming computational cost. To address this issue, the problem is reformulated in terms of probability distributions of discrete ambient conditions and associated control actions. Hopefully the very deep revision of this part of the paper provides now for a clear description of the formulation.

1.9. Line 183: What does sufficiently accurate mean here?

This phrase does not appear anymore in the revised manuscript.

1.10. Line 184: Are the parameters tuned based on the experimental measurements in the wind tunnel?

No, tuning was based on extensive FAST.Farm simulations. This part has been rewritten and better explained.

1.11. Line 200: Please elaborate on this. How are control effects captured “inherently”? Also, consider removing the mention of down-regulation, as you are not using this control mode.

This phrase does not appear anymore in the revised manuscript. The new text also explains clearly that only wake steering is used in the examples, while in principle the method is generally applicable to any wake control technique.

1.12. Section 3.2.1/ 3.3 I think understandability can be improved by explaining the path from DEL à Component damage à Lifetime à Farm lifetime in one section. This information is distributed over several sections, so it is hard to understand the approach in detail. Consider restructuring.

We believe that the introduction at the beginning of section 3 and Figure 1 already serve the purpose of giving an overview of the path from rotor inflow to lifetime. The text has also been deeply revised in this part of the paper. In particular we have moved all plots and discussions on the performance of the surrogate to the results section and to a new appendix. We believe that the new version is much clearer than the previous one.

1.13. Equation 10: I understand that the electricity price model details are given in the reference, but I think it would improve understandability if you could include a plot showing the relationship between wind and price signal. Are there negative prices?

Thank you, we agree and we have now added a new figure showing the correlation between price and wind speed.

1.14. Equation 10: U is not defined, I guess it is the wind speed?

Yes, it was already defined in section 2, but we have now recalled the meaning of that symbol here again for clarity. All symbols are consistent in their meaning throughout the paper. Additionally, all main symbols and acronyms are listed and explained in the nomenclature section at the end of the paper.

1.15. Figure 5: The caption is confusing to me. What are the DELs on the x-axis? Also it is a prediction of the load surrogate not of FLORIS?

Thank you, we agree with this remark. We have revised and improved the description in this and most other captions in the paper. There is also a nomenclature section at the end, which we have updated based on the present deep revision.

1.16. Section 4.1: Some content here is not really "Result and Discussion". I suggest integrating this in the methodology sections.

Thank you for the suggestion, which has now been implemented: the alternative reference methods are now explained in the methodological part of the paper.

1.17. Line 426: This strategy is not clear to me. What is the lifetime constraint here? In Figure 11 it is shown that this strategy ends up having a lifetime greater than the nominal lifetime. Wouldn't a lifetime constraint of 20 years just maximize the production, while ensuring that the 20 years lifetime constraint is met?

The description of this and the other formulations has been improved for clarity, so hopefully this doubt has already been addressed. In any case, as for the proposed formulation, a lifetime constraint of 20 years (imposed as the maximum allowable damage threshold) will ensure that the wind farm operates for at least 20 years. Since this condition is enforced as an inequality constraint, an operational lifetime of more than 20 years is possible, for example because the ambient conditions are not excessively harsh and demanding in terms of loading.

1.18. Section 4.2.3: You mention how lifetime is calculated for one component, but I can't find how you define lifetime for the farm (which you use as a metric later on).

Thanks, this is now explained in Sect. 2.1.

1.19. Figure 12: What are the extended columns without outline? The extended lifetime?

Yes, this is correct. The new text and captions are now clearer on this point.

1.20. Section 4.2.5: A table (or a more detailed description) of the parameters you vary would be helpful.

Thank you. The parameters of the cost model are described in Table 1.

1.21. Line 623: Why did you choose 8% as a variation of input parameters? I would expect some input parameters to have uncertainty higher than 8% easily?

Thank you for the comment. Indeed, depending on the modeling assumptions of the parameters, one could expect that some of them may vary more than 8%. However, the goal behind this assessment was

- to check whether the uncertainty bounds change drastically for a small change in parameters, and
- to understand the interplay of uncertainties of individual parameters on the output uncertainty.

As a consequence, the parameters were varied over values ranging from 4% to 12%, to see the trends in the output uncertainty variation through

Monte Carlo simulations. Figure 1 shows the output uncertainty bounds (on the y axis) for the lifetime profit of the proposed formulation obtained for different input parameter uncertainties (on the x axis). As seen in the figure, the bounds increase only linearly with a small slope for a linear increase in input parameter uncertainties. The choice of 8% was made based on the fact that the slope is essentially constant around that value.

A sentence was added to the text to explain the linear behavior that we observed.




Figure 1: Output uncertainty bounds for the lifetime profit of the proposed formulation obtained for different input parameter uncertainties.

Reviewer 2

This work considers wind farm controls (in particular, wake steering via yaw misalignment) using a cost function that accounts for operation and maintenance costs arising from turbine damage alongside revenues generated by the wind farm from energy sales. Overall, I found the core ideas to be interesting, the work to be well-motivated, and the description of the load and cost model to be well done. I also appreciated the sensitivity study carried out on the cost model parameters and the general approach to tuning and initial comparisons using FAST.Farm.

2.1. 1. The first is, I believe, not a concern about the core approach but simply how it is presented. In the initial portions of the paper, the authors describe optimizing the control actions over the turbines' lifetime T . For instance, the net revenue sum in Eq. (1a) is given over " T ", which is subsequently described as the "lifetime". This, as well as the surrounding text, gave me the impression that the control set points u were being optimized for each "time step" in the turbine's lifetime. After some further reading, and considerably further on in the manuscript, I realized (I think!) that this is not true—rather, the sum should be over the binned wind conditions W_a that the turbine will experience in its lifetime, presumably weighted by the frequency occurrence of that wind

condition based on the wind rose. If I'm correct that that is what is happening, please consider revising the formulation of problem (1) to be clearer. Consider including: the actual variable that is being summed over in the main sum; where that variable appears in the cost function; whether (bold) u represents the control actions at all turbines -for a single wind condition- or whether u represents the control actions at all turbines -for all wind conditions (time steps?)-; whether u -max and u -min are simply N -vectors ($N \times T$ vectors?) of the same values repeated, or if u -max and u -min may be different for different turbines; how weights corresponding to the frequency of occurrence of a given wind condition appear in the cost function (currently, it seems that all time steps appear in the cost function, but this seems unlikely, or that all wind conditions are given equal weighting of 1 in the cost function regardless of their probability of occurrence). I would also recommend updating the text of Section 2 and possibly other sections to clarify whether the optimization happens over wind conditions or time steps.

Thank you for the comments. We recognize that the previous version was not at all clear, and we have greatly expanded this part of the paper. Please see also the reply to a similar request from Reviewer 1.

2.2. 2. I am also concerned about whether there is any time-discounting taking place in the optimization. I don't see any mention of using a discount rate when deciding on what control actions to take, but discounting is a major consideration in any years-long cost/revenue analysis and may have a major effect on your results. For instance, discounting will likely make lifetime extensions considerably less valuable than it appears if a dollar in 10 years' time is considered to have the same value as a dollar today; but if a discount rate of 8% is used (which I understand to be fairly typical in economics, but I am no economist so please feel free to disagree on that front), a dollar in 10 years time is worth less than 50 cents today. Do your results in Section 4 consider discounting? If so, please make this clear. If not, I think this should be included.

Possibly more difficult: can your optimization problem take time value of money into account? When I at first thought that you were optimizing for each time step in the lifetime, it seemed straightforward to put a discount rate into the cost function; but now that I believe you are optimizing over each condition without considering "time" in the cost function, it's not totally clear to me how this would be done.

Finally, the only mention of how value changes in time that I saw in the paper was in Section 4.3.1, "An annual inflation rate of 8% is applied for the economic evaluation.". This needs to be explained further. Are you saying that costs and revenue -increase- at a rate of 8% per year (extremely high inflation if applied year over year for ten years) so that power generated in year 10 is actually worth more than power generated in year 1? I would strongly disagree with this approach, since it seems to be the inverse of discounting. Again, not being an economist I'm not sure how inflation should be handled, but my understanding with a quick search is that inflation is usually left out of discounting calculation under the expectation that while revenue will be higher due to inflation, so will all costs (salaries, normal operating costs, etc) so there is a net zero

effect. I may be misunderstanding what you mean in the inflation sentence, so please explain this better so that readers understand the approach taken.

Once again, we recognize the lack of clarity of our initial submission. We have now introduced time discounting already in the formulation of the problem in Section 2, and we have added a new dedicated section (3.4.3) in the description of the economic model. To summarize the results of the discussion developed in the revised manuscript, time discounting is not used in the calculation of the setpoints, but it is used when analyzing the performance of the controller in a given scenario. A detailed explanation and reasoning is given in Sect. 3.4.3. In the results section, we repeat this for extra clarity.

In the pre-print, the term inflation was mistakenly used, as discount rate was actually meant there. For the source of a discount rate of 8%, see NREL Cost of Wind Energy reports 2024 [5].

2.3. 3. The authors argue that the 100-hour optimization time is acceptable since the optimization does not need to take place online. Still, looking at the curve given in Figure 17, I would argue that some (future) improvements will be needed to make this feasible. Very roughly extrapolating, it looks like a 100-turbine farm, which is pretty typical, would take around 1000 hours to optimize, and that’s using a limited wind rose with only 3 wind speed bins. 1000 hours is over a month, which I feel is likely not acceptable. Although the authors suggest the optimization may not need to be repeated, it seems to me likely that a wind farm operator may want to rerun the optimization say, once per year, to update the turbine damage information and replan based on actual observed damage. It is clearly out of scope to design a new optimization approach, but do the authors have any thoughts about how the optimization might be sped up in future work? Could the problem be parallelized in any way? FLORIS itself has some inbuilt parallelization that may be useful for this. In fact, I didn’t understand what approach/algorithm is taken for optimization—the authors mention a “one-shot” approach, but I’m not familiar with this term so it would be good to explain at least briefly in words how this works.

We agree that in the scenario described by the reviewer, the computational cost is substantial.

The present code was not yet optimized for execution speed. The implemented version is based on FLORIS V3. The current version 4 has increased computational speed, so the code could be adapted to the most recent FLORIS version. Furthermore, the calculation of loads via the surrogate model offers some extra potential. In the present code, the damage equivalent loads (DELs) for all turbines and their components are calculated sequentially, which could be parallelized.

The term “one shot” optimization in the present paper means that the

optimized yaw setpoints must be determined for all wind conditions simultaneously. In a conventional max. power optimization, the optimal control action can be determined individually for every environmental condition (e.g. separately for every wind direction). In the new formulation, the calculation of damage couples all control actions and ambient conditions together. The optimization problem becomes therefore much larger. We have greatly expanded the discussion on this point, which is addressed in multiple places, including the introduction, the methodology section 2, the results and discussion section 4, and the conclusions. We hope that this point is now clearly described.

There are tools available to deal with this increase of complexity. FLORIS already offers options to reduce the number of variables (e.g. discarding symmetric situations). It offers also the "serial refine" method which first probes possible yaw offsets to begin the optimization at a promising starting point. These tools have not been leveraged in the present study.

The optimization algorithm deployed is the SciPy SLSQP algorithm. Alternatives, which scale better with increasing turbine count, could be tested as well.

We have expanded the discussion on the computational cost in the revised version, and we have clearly pointed out that this is probably the main challenge that needs to be address to make this formulation more widely applicable.

2.4. Line 65: What does "(future better)" mean? I'm unfamiliar with this phrase. Consider revising.

The sentence does not appear anymore in the revised manuscript.

2.5. Section 3.2.2: Horizontal shear is included in the load model, but not vertical shear, if I'm understanding correctly. I would have thought that vertical shear would also be an important driver of fatigue loads (especially blade loads), even if it is not controllable using wake steering. Do the authors have any comments on this?

We clearly agree with the reviewers comment that the vertical shear does have an influence on loads. A load surrogate model trained considering vertical shear can produce more accurate predictions. However, the vertical shear in FLORIS does not change inside the wake region. Furthermore, FLORIS does not take into account the vertical wake deflection due to rotor tilt, which –on the other hand– is observed in the FAST.Farm training data.

In the end, the selected design was the best compromise for the current study. Our tests showed that the overall performance of the combined

models FLORIS + load surrogate is better without vertical shear, rather than with an inadequate shear predicted by FLORIS.

2.6. Line 318: Although I agree that the trend between wind power output and electricity price should be cited, having 6 citations for this seems excessive. Perhaps one or two citations would be enough here? Just a suggestion, feel free to ignore.

We have reduced the number of references to Kölle et. al. 2022, which describes the relation between power output and electricity price.

2.7. Line 360: It might be interesting to help readers with intuition to demonstrate in a figure how the cost model (11) changes with respect to (non-optimized) control inputs. I personally would find this helpful. Again, just a suggestion, not critical.

Thank you for the suggestion. Figures 9 and 10 already show the trends of turbine loading with respect to wind direction when the control inputs are not optimized. The resulting economic costs are proportional to loading. We believe that the new figure will not bring any new information that is not already provided in the results. Thus, we have decided not to include the suggested figure, also because the paper is already very long and with many figures.

2.8. Eq. (12): Do the reference optimization problems (12)–(15) still consider the input constraint $u_{\min} \leq u \leq u_{\max}$ (1c)? If not, I would suggest adding this to make it more consistent with the max profit optimization approach (or justify not including this constraint). If so, please make this clear.

Yes, all the reference optimization formulations consider the bound constraints on the decision variable. The formulations in Sect. 4.1 have been updated to state this explicitly.

2.9. On several occasions the authors mention that the "max power" approach takes more "aggressive" offsets than the proposed "max profit" approach. Can the authors explain what they mean by this? Looking at Figure 7, it's not at all clear to me that the "max power" approach takes aggressive offsets (on the contrary, they seem to be smaller than the offsets selected by the "max profit" approach).

Thank you for this remark, and apologies because the wording was in fact not precise. We were referring to the more aggressive wake steering of T2 in the sector $\Gamma > 270$. In contrast, the other strategies deploy a one-sided yaw-offset schedule for turbine T2 and avoid wake steering for this sector. This is the reason for the premature end of lifetime for the Max. Power strategy. The text has been updated to clarify this

point.

2.10. Line 511: "Since FLORIS does not account for vertical shear...". I'm not sure what the authors mean by this—FLORIS does account for vertical shear across the rotor points when computing a rotor-average inflow, although not in the evolution of the wake. Perhaps this is what the authors are referring to? This sentence should be updated.

Thanks for this remark, and yes, that's what was meant. The text has been updated to make this point clearer.

2.11. Line 654: "It is assumed that these wind conditions remain constant over the course of the year". What do the authors mean by this?

The sentence simply points to the assumption throughout the case-studies that the considered wind rose, denoting the annual wind conditions, does not change from one year to another. The text has been modified to bring more clarity to this point.

2.12. Line 657: Inflation of 8% year over year seems to be very high. I would recommend something much lower than this if inflation is still included after addressing my main comment 2. above.

Our apologies again, we mistakenly wrote 'inflation rate', instead of 'discount rate', which is the correct term used to describe time-discounting in the calculation of the present value of future profit. The text has been corrected.

A discount rate of 8% is found to be a typical value for wind energy projects [5]. Please also see our response to your comment **2.2** for more details.

Additionally, we have taken the opportunity of this very deep revision for making numerous editorial changes to the text, in order to improve readability.

A revised version of the manuscript is attached to the present reply, with the main changes highlighted in red (deletions) and blue (additions).

Best regards,
The authors

References

- [1] Harrison, M., Bossanyi, E., Ruisi, R., and Skeen, N.: An initial study into the potential of wind farm control to reduce fatigue loads and extend asset life, *J. Phys.: Conf. Ser.*, 1618, 022007, <https://doi.org/10.1088/1742-6596/1618/2/022007>, 2020.
- [2] Canet, H., Guilloiré, A., and Bottasso, C. L.: The eco-conscious wind turbine: design beyond purely economic metrics, *Wind Energ. Sci.*, 8, 1029–1047, <https://doi.org/10.5194/wes-8-1029-2023>, 2023
- [3] Shou, T.: A Literature Review on the Net Present Value (NPV) Valuation Method, *Advances in Economics, Business and Management Research*, <https://doi.org/10.2991/aebmr.k.220603.135>, 2022.
- [4] Short, W., Packey, Daniel J., and Holt, T.: A Manual for the Economic A Manual for the Economic Evaluation of Energy Efficiency and Renewable Energy Technologies, <https://docs.nrel.gov/docs/legosti/old/5173.pdf>, 1995.
- [5] Stehly, T., Duffy, P., and Hernando, Daniel M.: Cost of Wind Energy Review: 2024 Edition, <https://docs.nrel.gov/docs/fy25osti/91775.pdf>, 2024.

Economic lifetime-aware wind farm control

Abhinav Anand, Robert Braunbehrens, Adrien Guilloiré, and Carlo L. Bottasso

Technical University of Munich, Wind Energy Institute, Boltzmannstraße 15, 85748 Garching b. München, Germany

Correspondence: Carlo L. Bottasso (carlo.bottasso@tum.de)

Abstract.

We present the economic lifetime-aware formulation of wind farm control (WFC), a novel approach ~~that incorporates the current state of damage of turbines—either provided by a digital twin or estimated from operational data—to compute control policies that optimize the long-term life management of a wind farm.~~ The ~~characterized by two key innovations with respect~~ to the state of the art. First, ~~the~~ optimization is guided by an economic ~~value function that balances power capture (taking into account variable spot market prices) with~~ model of profit. By targeting profit instead of the more widely adopted power yield, the operational optimization considers not only the revenue generated by selling electricity (possibly at a variable price when operating on spot markets) but also the costs due to operation and maintenance (O&M) ~~costs, including potential revenue losses due to turbine downtime.~~ The optimization process is subject to constraints ensuring that the target lifetime of selected ~~turbine components is met, considering varying inflow conditions. This results in~~ . Second, since fatigue-induced damage is a finite-budget quantity, we explicitly consider that any effect on damage generated by a control action will leave more or less room to other control actions, a coupling that is neglected by current approaches that arbitrarily limit or penalize fatigue-inducing loads. To achieve this goal, the optimization is subjected to explicit constraints on lifetime, which ensure that the turbines in a farm and their components can indeed reach a desired goal duration. As for other optimization-based WFC ~~approaches, the proposed method results in the offline calculation of~~ optimal control setpoints for each turbine, which are scheduled and dispatched at run-time in open loop as functions of ambient conditions.

After introducing the economic lifetime-aware WFC approach, the paper analyzes two case studies. The first ~~is a synthetic scenario with~~ represents an idealized scenario characterized by a simplified set of environmental conditions, designed to demonstrate the behavior of the new control strategy and its impact on fatigue loads in conditions of very frequent waking. The second case features ~~the more realistic setup of a~~ a more realistic small wind farm, with wind climate data derived from ~~a real-an~~ existing meteorological mast. The benefits of the proposed methodology are highlighted by comparing it to ~~a range of reference~~ alternative previously-published control strategies, which either ignore component damage or address fatigue loads in alternative simplified ways. ~~When compared to greedy and power-maximizing WFC strategies, results show that only~~ Results show that the lifetime-aware formulation always ensures the achievement of the desired lifetime targets, whereas methods that do not ~~enforce such a condition may fail to do so.~~ Additionally, ~~it the proposed method~~ also results in the best economic performance, i.e. in the highest profit.

1 Introduction

1.1 Wind farm control beyond power

Most current wind farm control (WFC) strategies focus on increasing power generation (?). However, this is only one of
30 ~~several goals of~~ the several goals that may be targeted by the more general concept of optimal operation. In fact, a wind energy project sits in a complex dynamic environment and has various stakeholders. External factors, such as electricity market price or grid loads and transmission constraints, continuously evolve over time. ~~This variation has a dynamic~~ These dynamics have an impact on the value of the wind farm power output being fed into the grid. For example, in some cases, it pays off to extract maximum yield from the farm. In contrast, in other cases, generating less will be more convenient because ~~of the~~ low electricity
35 prices ~~that~~ cannot offset the incurred costs. ~~Within the farm itself,~~

Flow control strategies are designed to modify the inflow to the turbines in a farm. This can be obtained either by derating the turbines (called induction control), by misaligning them with respect to the wind (called wake steering), or by applying excitations meant to accelerate wake recovery (called wake mixing) (?). However, every control action will affect the internal response of the ~~system, e.g. by changing farm,~~ not only in terms of power but also in terms of the loading on the turbines ~~and~~
40 ~~their components~~. In turn, this affects ~~the turbine their~~ “health”, ~~directly influenceing with a cascading direct influence on~~ the overall asset lifetime and ~~the its~~ repair and maintenance activities ~~needed during its operation~~. Considering these factors and couplings, it seems hardly optimal to feed in maximum power to the grid at all times. ~~To successfully optimize a wind energy project, the scope needs to be expanded towards a more holistic view. New formulations of WFC need to be able to incorporate, on the one hand, all relevant external requirements and conditions and, on the other, the internal response of the system, and~~
45 the successful optimization of a wind energy project requires a more holistic approach.

~~During operation, wind turbines incur fatigue damage, i.e. structural degradation in their components such as tower, blades, drive-train, and others (?). After accruing a limiting damage value, the asset is considered to have reached its end of life (EOL). Flow control strategies modify the inflow conditions within the farm, either by derating the turbines, by misaligning them with respect to the wind, or by applying excitations meant to accelerate wake recovery (?). Therefore, a farm control policy~~
50 ~~affects the rate at which damage accrues in the turbine components. This, in turn, has a direct effect on the asset lifetime, potentially limiting or extending the project duration and affecting the economic profitability of the wind farm. Thus, wind turbine fatigue damage plays a crucial role in determining the performance of a farm control policy. Some initial efforts at including fatigue in wind farm control have already been documented (???). Such control strategies typically try to restrict damage equivalent loads (DELs) by either setting a limiting value or by~~ considering them directly in penalizing them through
55 ~~an additional objective. However, such choices can hardly lead to a clear understanding of the effects of the resulting control policies on lifetime. In fact, damage accrues over time, and therefore, the control actions taken today – through the loading of the turbines – directly affect the room left for control actions that can be taken in the future. Thus, optimizing the operation of a wind farm over its lifetime cannot only rest on decisions based on present conditions or over short time horizons. Instead, one needs to consider the couplings between present and future actions.~~

60 ~~This work presents~~ More sophisticated approaches have been recently described that attempt at better expressing the trade-offs between fatigue and economic objectives that occur throughout the lifetime of an asset. ? introduced a framework for the optimal long-term planning of turbine operation that treats fatigue damage as a consumable budget. This approach actively distributes the damage budget along the operational life of the turbine to balance load reduction and energy output, demonstrating promise in reducing fatigue while maintaining performance. Similarly, ? developed an adaptive controller for wind turbines, 65 tailored to reconcile fatigue and revenue goals. This approach uses variable power modes – including down-regulation and power boosting – alongside individual pitch control. A Gaussian Process Regression surrogate model is employed to capture trade-offs, enabling an optimization of the operational modes based on wind and price dynamics to maximize lifetime economic benefits.

Although clearly going in the direction of modeling the lifetime economics of wind assets, these studies are limited at the 70 turbine level, and do not consider the effects caused by wakes in a wind farm. As wake effects have been shown to significantly affect not only power capture but also structural loading (???), the present work considers the problem at the farm – instead of the turbine – level.

1.2 Maximizing profit and ensuring a desired lifetime

Here we introduces a new form of WFC, termed economic lifetime-aware WFC. The formulation incorporates the current state 75 of damage of the turbines – either provided by a digital twin or estimated from operational data – to compute control policies that optimize the long-term life-management of the wind farm. The ~~optimization is~~ problem is formulated as an optimization guided by an economic value function that ~~balances power capture (taking into account variable spot market prices) with~~ expresses profit, i.e. the difference between revenue and cost.

Revenue is gained by producing power and selling it on the spot market at variable prices. However, in order to target profit, 80 which is clearly a more complete economic indicator than revenue alone, it is necessary to account for the way turbines are operated. In fact, different operational strategies affect – through loads – the costs of operation and maintenance (O&M) ~~costs,~~ including potential revenue losses due to turbine downtime. The ~~optimization process,~~ in turn leading to loss of revenue due to downtime. Additionally, damage due to loads has also another effect on revenue and cost: when the lifetime of a component is affected, either the component must be replaced (which implies a cost for the spare part and its installation, and a loss of 85 revenue due to associated downtime) or the whole turbine must be dismantled, which leads again to a loss of revenue due to lost production caused by the shortened life. To account for such long-term effects on profitability, the optimization process proposed here is subject to constraints ~~ensuring that the target lifetime that ensure meeting a target lifetime duration~~ of selected turbine components ~~is met, considering varying inflow conditions. The optimization variables are the control setpoints of.~~

The lifetime-aware economic optimization problem is solved in terms of control setpoints for each turbine. The ~~selection of~~ the ~~control variables type of control setpoints~~ depends on the ~~specific WFC implementation chosen : induction-based control to reduce wake intensity, yaw offset-based control for~~ chosen operational modes, which may include induction control, wake steering, or a combination of both. Although mixing is yet another flow control method, it was not considered in the scope of the present work. The proposed methodology is versatile and can be applied to any WFC implementation, allowing the end 90

~~user to choose based on their preferences and the trade-offs offered by different control strategies. The outputs of the WFC optimization problem are the optimal~~ mixing, and possible combinations thereof. The proposed approach is demonstrated here using wake steering alone, but the method is general and all other operational modes could be readily included.

Current WFC methods are based on the offline calculation of steady-state control setpoints, which are typically scheduled as functions of ambient conditions, such as wind speed, direction, and possibly other parameters in order to capture additional effects (?). Since each ambient condition with its associated control action is independent from the others, a separate optimization is solved for each combination of scheduling parameters. The total number of necessary optimizations depends on the desired resolution (i.e. on the size of the bins used to discretize the parameters and the number of considered combinations). At run-time, the current ambient conditions are detected and used to gauge from the precomputed setpoints the ones that need to be dispatched to the turbines in open loop (?). This method has the advantage of moving any computationally-heavy optimization to the offline phase, while the run-time execution is extremely fast and safe, because setpoints are sourced from a precomputed and verified pool. Dynamic (model predictive) approaches optimize performance over a prediction horizon, which might lead to improved performance compared to the just-described steady offline case (?). However, such methods require more complex dynamic flow models, shift the computationally-intensive phase to the execution time, and generate new solutions at each call, which renders their verification for safety an additional challenge.

With our new formulation we aim to capture the long-term effects of control policies. In principle, this means that the economic optimization problem should be formulated over a potentially very long time horizon (the entire lifetime of the farm, or a fraction of it). This would require the solution of a time-marching problem, which however would come at a potentially overwhelming computational cost. To make the problem tractable, we use an approach similar to the offline one described above. First, we bin the ambient parameters and we count the frequencies of their combinations. Next, we solve the optimization problem that yields the control setpoints for each ~~turbine in the farm as functions of ambient condition~~ combination of binned parameters. The crucial distinction with the standard offline method described earlier is that now the solution is performed in one single shot, i.e. simultaneously considering all binned combinations with their respective frequencies. This clearly increases the computational cost, because instead of many decoupled optimizations we now solve one single large one. However, this approach has the key advantage of coupling the effects of the control actions together. In fact, as we wrote earlier, a control action taken in some operating condition consumes some fraction of the lifetime, and therefore changes the room left to other control actions that will be used in other operating conditions. By taking this coupling into account, we can include lifetime directly into the formulation (as a constraint), instead of working with DEL reductions or other unclear proxies.

Other advantages of the standard offline approach are retained here. In fact, while the solution may be expensive, it is still performed completely offline, and at run-time setpoints are readily obtained from the precomputed ones based on the current ambient conditions. And here again, setpoints can be verified before they are actually deployed to the turbines, thereby guaranteeing their safety. ~~Although sophisticated dynamic (model predictive) approaches are possible and indeed promising (?), the present relatively simpler implementation based on offline-computed look-up tables is very close to the power-boosting WFC technology already successfully implemented by industry.~~

Compared to existing approaches, the novelty of the proposed WFC methodology lies in incorporating the desired wind farm lifetime as an optimization constraint and using an economic merit function as the optimization objective

130 1.3 Application scenarios

As explained, our proposed WFC method departs from the current literature by (i) using a model of profit as economic optimization objective, and (ii) incorporating a desired lifetime as a constraint, thereby taking into account the room left by each control action to the others. While this paper highlights the combined use of these two innovations, it is important to realize that they can also be applied separately. If In fact, if component damage is not a major concern (for instance, at a site with a mild environment relative to the turbine design limits), the lifetime constraint can be omitted, while still using the economic value function to drive the optimization. This towards a maximization of profit. Eliminating the lifetime constraint reduces the computational cost (because independent – instead of coupled – optimizations can be performed for every combination of binned ambient parameters), while still offering a clearer optimization goal over existing formulations that blend power production and loads with arbitrary weighting factors (?). Alternatively, a different (future better) economic merit function can be used, than the simple one adopted here, while still maintaining the fatigue constraint to ensure that the wind farm meets the desired target lifetime. (?).

This new economic lifetime-aware control framework allows wind farm developers in a can be used in different phases of a project. In a pre-deployment phase, a wind farm developer could use it to make informed decisions on the economic viability of various alternative operational scenarios. For instance, a developer could compare a shorter lifetime duration of the plant due to higher power capture (but also higher loading and hence higher O&M costs) followed by decommissioning vs. an extension of the plant lifetime beyond its initial intended duration thanks to a less aggressive capture and less loading. Moreover In a second phase, once an optimal life management strategy has been selected, the wind farm operator can deploy the corresponding lifetime-aware controller to operate the turbines in a the chosen profit-optimal way, while ensuring that the specified lifetime goals for the turbine components are met.

150 1.4 Scope and organization of the work

The economic model of profit used here is simple, one of its main limitations being the crude link between the effects of loads and O&M costs. While this is due to the limited open literature existing on this important – although little investigated – topic, it is clear that better models, as those surely available at several major companies, could be readily used within this approach without significant modifications.

155 An experimental validation of the proposed WFC methodology was recently conducted by wind tunnel tests (?). The experiment involved three model turbines in a boundary layer wind tunnel, where a turntable was used to simulate dynamic variations in wind direction. The control strategy, based on wake steering, was synthesized using power and damage estimators. The proposed strategy was compared to conventional greedy operation and the classical maximum power strategy. Results showed how the new control approach not only improved the economic performance of the wind farm, but also ensured

160 operation over the desired lifetime. In this paper, we develop a more detailed description of the economic lifetime-aware WFC formulation than the one of ~~-,~~ and we demonstrate the new method in a simulation environment.

The rest of this paper is organized as follows. The proposed ~~economic lifetime-aware WFC approach~~ is presented in Sect. 2. This is followed in Sect. 3 by a detailed description of the optimization framework ~~in Sect. 3, together that,~~ with its supporting ~~tools~~models and tools, enables the calculation of the profit-optimal control laws. Next, ~~the proposed strategy and a few reference WFC formulations are~~ in Sect. 4 the economic lifetime-aware WFC method is applied to two case studies ~~in Sect. 4,~~ considering various performance indicators, including power, loads, ~~and lifetime duration~~ lifetime duration and profit. Results are compared to some alternative WFC formulation sourced from the recent literature. ~~The section includes a discussion on the results, comparisons with the reference alternative formulations, and an evaluation of the computational burden.~~ Finally, conclusions and key findings are discussed in Sect. 5.

170 2 ~~Optimization problem~~Lifetime-aware wind farm operational optimization

~~A standard power-capture optimal wind farm controller is typically designed by independently optimizing each ambient condition (???). The assumption behind this approach is that optimizing each condition will ultimately maximize the overall performance over time. In the case of power, this is straight-forward as the produced energy over time can easily be summed up. However, designing a lifetime fatigue-optimal wind farm controller requires a different approach. Fatigue damage in a turbine component accumulates in a complex way over its lifespan, influenced by all ambient conditions encountered. Consequently,~~

2.1 Operational optimization for profitability

We consider the problem of optimizing the way a wind farm is operated in order to maximize its profitability.

Profit expresses the trade-off between control actions aimed at increasing the revenue generated by power yield, and those designed to mitigate the economic cost of fatigue damage (??). In order to optimize for profit, a suitable economic model is required. The economic model builds on a physical model expressing the effects of the control actions on the flow within the farm, and hence on power capture. In turn, power yields revenue, gained by selling electricity. The same flow model informs a load model that provides the fatigue loads on the turbines, whose accumulation finally yields damage. In turn, damage enables the calculation of O&M costs and lifetime duration, the latter affecting again revenue.

Section 3 explain how these various models can be formulated and connected together in order to enable the operational optimization of the farm for maximum profitability.

2.2 Operational time horizon

The operational optimization is performed over a certain time horizon.

In a first scenario, this temporal horizon is unknown. In other words, the solution of the ~~damage accumulated under various wind conditions directly affects the remaining operational life of the turbine. As a result, a lifetime fatigue-aware wind farm controller must be optimized holistically, taking into account the entire spectrum of ambient conditions throughout the lifetime~~

of the turbines. Additionally, fatigue damage can significantly impact both the failure rate and the frequency of maintenance activities (??). This not only increases maintenance costs but also leads to economic losses because of the reduced power production during repair and maintenance periods. operational optimization yields the way the turbines should be operated, but also the length of time over which the wind farm should be run. Extending the operation beyond that optimal duration would result in a smaller profit, for example because of an excessive costs due to the aged turbines, or because replacing them with more modern machines would be more economically profitable, or various other reasons.

The proposed WFC formulation is designed to optimize wind farm operations over its entire lifetime. The consideration of lifetime effects is integrated into the optimization problem in two ways. First, the proposed optimization ensures that the turbines operate for a specified duration. In a second scenario, the operational optimization of the wind farm is defined over a known desired temporal horizon. In this case, for example, a developer might want to ensure that a new wind farm is operated with maximum profit for a duration of at least X years. Another example might be the one of an operator that owns a farm that has been running for Y years, and now wants to optimize the economic performance of that farm for the next Z years.

This work considers this second scenario: we want to ensure that turbines in a farm are operated in a profit optimal way for a desired period of time without exceeding a predefined fatigue damage threshold. This threshold may represent the design lifetime of the turbine turbines or a desired fraction of the remaining useful life, depending on the goals of the wind farm operator. This step is crucial, scenario (see later in the results section for an example). Clearly, the ability to reach a desired duration has economic implications, as failure to meet the threshold do so would result in revenue losses due to the reduced turbine lifespan. The second lifespan. Published WFC formulations can not guarantee that this objective is met. The present formulation tries to fill this gap.

2.3 Accounting for lifetime in operational optimization

Although ensuring a desired lifetime has a clear economic advantage, it also drastically complicates the problem of designing a controller capable of achieving such a goal.

To understand this aspect of the proposed optimization focuses on maximizing the overall economic profit by accounting for the impact of fatigue damage on operation and maintenance costs. This is achieved by finding trade-off between control actions aimed at increasing the economic value of power yield and those designed to mitigate the economic cost of fatigue damage. The resulting optimization provides an optimal schedule of setpoints for each turbine, tailored to the site ambient wind conditions.

This control formulation is implemented in an open-loop fashion, similarly to widely used wake-control-based power capture optimization strategies (???), enabling the offline application of computationally intensive one-shot optimizations through problem, consider a standard wind farm controller based on the offline optimization of steady-state simulations setpoints and their subsequent online open-loop dispatch to the turbines (???). Ambient conditions are characterized by combinations of parameters such as wind speed U , direction Γ , turbulence intensity (TI) I , and possibly others. Such conditions have an effect on the wakes that couple the turbines together in a farm. To account for such effects, the atmospheric parameters are first discretized into bins and combined to create a finite set of discrete ambient conditions. Next, setpoints are calculated for each

225 one of these discrete ambient conditions. Since there are no couplings among the different conditions and among the associated control actions, multiple independent calculations of the setpoints can be performed.

~~The optimization problem can be written as~~ On the other hand, a different approach is necessary when lifetime needs to be precisely accounted for. This is due to the fact that fatigue damage in a turbine component accumulates in a complex way over its lifespan, influenced by all encountered ambient conditions and all performed control actions. Consequently, the damage accumulated at a given time in some given wind condition directly affects the remaining operational life of the turbine.

2.4 Lifetime-aware operational optimization in the time domain

As a result of this coupling among actions and over time, the design of a fatigue-aware wind farm control policy can not be performed for separate conditions. Instead, it is formulated as an optimal control problem (?) in the time domain t defined over a desired lifetime horizon T , which writes

$$235 \quad \max_{\mathbf{u}} \sum_{\substack{T \\ t \in T}} \mathbf{u}(t) \int \left(\text{Revenue}(\mathbf{W}_a, \mathbf{u}(t), t) - \text{Cost}(\mathbf{W}_a, \mathbf{u}(t), t) \right) dt, \quad (1a)$$

subject to

$$D_i^c(\mathbf{W}_a, \mathbf{u}) \leq D_{\text{ref}}^c \quad \forall i \in N \text{ and } \forall c \in C,$$

$$D_i^c(T) \leq D_{\text{ref}}^c \quad \forall i \in N, \forall c \in C, \quad (1b)$$

$$240 \quad \mathbf{u}_i^{\min} \leq \mathbf{u}_i(t) \leq \mathbf{u}_i^{\max} \quad \forall i \in N. \quad (1c)$$

and

$$\mathbf{u}^{\min} \leq \mathbf{u} \leq \mathbf{u}^{\max}.$$

For a given set of ambient wind conditions $\mathbf{W}_a$, control set-points $\mathbf{u} = [\mathbf{u}_1, \dots, \mathbf{u}_i, \dots, \mathbf{u}_N]$ for N turbines in the farm, and C components in each turbine, the optimization balances the accrued economic value. In the open-horizon scenario described before, the duration T is an unknown of the problem, together with the control actions $\mathbf{u}(t)$. Although this problem is clearly very relevant for developers and/or operators, here we study the simpler problem where T is given, and we call it the goal (desired) duration. As it will become apparent in the following, even this simpler problem is complex and expensive, and therefore it will be the sole focus of this work.

2.4.1 Optimization objective

250 The optimization problem expressed by Eq. (1) accounts for the economic value accrued through power generation, as expressed by function $\text{Revenue}(\cdot)$ and the incurred economic cost due to, but also for the economic cost incurred through maintenance activities and lost lifetime $\text{Cost}(\cdot)$, over a chosen lifetime T . The vector $\mathbf{W}_a$ denotes the set of all wind conditions through time

series of wind speeds, wind directions, and possibly other atmospheric parameters (turbulence intensity, shear), that occur over the lifetime T . In the present work, the considered as quantified by function $\text{Cost}(\cdot)$. Profit is the difference between these two quantities.

Notice that in Eq. (1a) revenue and cost depend explicitly on time t because of time-discounting, which is used to estimate the present worth of future cash flows. The role of discounting in economic WFC is discussed in Sect. 3.4.3.

2.4.2 Optimization unknowns

Revenue and cost depend on the control actions $\mathbf{u}(t) = [\mathbf{u}_1(t), \dots, \mathbf{u}_i(t), \dots, \mathbf{u}_N(t)]$, where $\mathbf{u}_i(t)$ is the vector containing the control setpoints (yaw, induction, etc.) of turbine i of the N turbines in the farm.

The control actions $\mathbf{u}(t)$ represent the functional unknowns of problem (1), and they depend on the time-variable ambient conditions $\mathbf{W}_a(t)$, i.e. $\mathbf{u}(t) = \mathbf{u}(\mathbf{W}_a(t))$. In this work, ambient conditions are the hub height wind speed U , wind direction Γ represented by the combination of wind speed, direction, and turbulence intensity (TI) denoted as I , resulting in $\mathbf{W}_a = [U, \Gamma, I]$. Assuming the wind climate does not show any inter-annual variability, the time series can be approximated through a corresponding frequency distribution. Depending on the chosen WFC technique, the control set-point vector $\mathbf{u}_i$ can consist of induction factor α_i for induction control, yaw-offset γ_i for wake steering, or a combination of the two: $\mathbf{W}_a(t) = \{U(t), \Gamma(t), I(t)\}$.

The bound constraint Eq. (1c) ensure that the setpoints $\mathbf{u}_i(t)$ for each turbine i remain within desired lower ($\mathbf{u}_i^{\min}$) and upper ($\mathbf{u}_i^{\max}$) bounds.

2.5 Lifetime constraint

The optimization problem is subjected to the lifetime constraint expressed by Eq. (1b), which ensures that the resulting lifetime damage D damage $D_i^c(T)$ of component c in the set of components C for turbine i at the end of the desired lifetime T does not exceed the damage threshold D_{ref} . This parameter D_{ref}^c . This threshold is an appropriate damage limit, which includes the necessary safety factors. Clearly, there are several options for determining the this limit. For instance, $D_{\text{ref}} D_{\text{ref}}^c$ could be set to the design limit, allowing the component to be utilized up to its maximum designed capacity. Alternatively, one might leave a damage margin, ensuring that the component has a second life and extending its overall usage. Despite the various possible scenarios, the underlying concept of this approach is to maintain control over damage accumulation while pursuing a specific goal. Achieving this goal over a desired duration would not be possible without an explicit constraint to guide the process. In this work, for a given component c , D_{ref}^c is chosen as that ensures that the duration will be met.

If the constraint expressed by Eq. (1b) is violated, the problem is infeasible. This means that it is not possible to operate the farm for the time duration T without exceeding the damage threshold. On the contrary, if the constraint expressed by Eq. (1b) is satisfied and

$$D_i^c(T) - D_{\text{ref}}^c > 0 \quad \forall i \in N, \forall c \in C, \quad (2)$$

then this means that the farm can be operated over the entire duration T , and some damage margin is still left. Therefore the farm could in principle be run for a longer period of time than the required T . By knowing the initial damage at time

285 $t = 0$ and the remaining margin at time $t = T$, the point in time when the damage threshold is reached can be estimated by linear extrapolation. It should be noted that the optimization is not performed over this extra time duration. However, the quantification of the extra profit that in principle could be achieved can be readily obtained by computing the extra revenue and cost accrued over this extended lifetime.

290 To conclude the discussion on the lifetime constraint, a distinction should be made between a replaceable and an irreplaceable component (e.g., the tower). A replaceable component requires some cost and downtime, after which the turbine goes back into operation. An irreplaceable component, on the other hand, implies that the turbine lifetime is exhausted, and the machine has to be dismantled. Although both situations should be taken into account for a general formulation, only the case of irreplaceable components is considered here to shorten the notation and to ease the implementation. In the applications presented later, the blades and the tower are treated as components, even the blades, for simplicity, being considered as irreplaceable.

295 2.6 Reformulation from the time domain to probability distributions

A time-series implementation of the problem expressed by Eq. (1) would require (i) providing a time series of atmospheric parameters over the lifetime T , and (ii) solving the optimization problem for the corresponding time series of control setpoints $\mathbf{u}(t)$. Today this is clearly unfeasible, as it would imply overwhelming computational costs. In addition, this approach would also be wasteful, as similar combinations of atmospheric parameters would appear over and over again during the optimization horizon.

300 To address this hurdle, the atmospheric parameters (speed, direction, TI, etc.) are first discretized into bins. Next, time series of the parameters are converted into discrete sets of ambient conditions $\mathbf{W}_a^{\text{bin}}$, each being represented by a combinations of the binned parameters. Counting its occurrences, each ambient condition is associated to its probability f_a^{bin} . To account for the temporal variability of the wind climate, different probabilities $f_{a,y}^{\text{bin}}$ can be used for each epoch y (an epoch being a certain duration of time, for example one year), with $y \in Y$, Y being the number of epochs in the time horizon T .

305 Next, the control setpoints of each turbine $\mathbf{u}_i$, instead of depending on time t , are assumed to depend on the discrete ambient conditions, i.e. $\mathbf{u}_i(\mathbf{W}_a^{\text{bin}})$ and therefore for the whole farm $\mathbf{u}(\mathbf{W}_a^{\text{bin}}) = [\mathbf{u}_1(\mathbf{W}_a^{\text{bin}}), \dots, \mathbf{u}_N(\mathbf{W}_a^{\text{bin}})]$. Finally, the time-dependent cost function appearing in Eq. (1a) can be approximated using combinations of binned parameters as

$$\underline{D_{\text{ref}}^c} = \max_i \int_{t \in T} \left(\text{Revenue}(\mathbf{u}(t), t) - \text{Cost}(\mathbf{u}(t), t) \right) dt \approx \sum_{y \in Y} T D_i^c \sum_{\text{bin} \in \text{Bins}} f_{a,y}^{\text{bin}} \left(\text{Revenue}(\mathbf{u}(\mathbf{W}_a^{\text{bin}}) \quad \forall i \in N, y) - \text{Cost}(\mathbf{u}(\mathbf{W}_a^{\text{bin}}), y) \right) \quad (3)$$

310 ~~which corresponds to the maximum accumulated damage over all the turbines under the standard greedy control strategy, where no wake steering or other WFC strategies are applied. Furthermore, the~~ Notice that here we are assuming that the same ambient condition will always trigger the same control response. While this seems to be reasonable, such assumption does not necessarily hold for the time domain problem Eq. (1). While it would be possible to allow for this extra freedom even here, by assuming a different set of control setpoints for each epoch, this approach is not pursued further because of its extra cost. A further discussion on this aspect of the formulation is developed in Sect. 3.4.3 with reference to time discounting.

By moving from the time domain to probability distributions of discrete ambient conditions, the problem expressed by Eq. (1) is reformulated as

$$\max_{\mathbf{u}} \sum_{y \in Y} \sum_{\text{bin} \in \text{Bins}} f_{a,y}^{\text{bin}} (\text{Revenue}(\mathbf{u}(\mathbf{W}_a^{\text{bin}}), y) - \text{Cost}(\mathbf{u}(\mathbf{W}_a^{\text{bin}}), y)), \quad (4a)$$

subject to

$$D_i^c(T) \leq D_{\text{ref}}^c \quad \forall i \in N, \forall c \in C, \quad (4b)$$

$$\mathbf{u}_i^{\min} \leq \mathbf{u}_i(\mathbf{W}_a^{\text{bin}}) \leq \mathbf{u}_i^{\max} \quad \forall i \in N, \forall \text{bin} \in \text{Bins}. \quad (4c)$$

The unknowns of the reformulated problem are the turbine discrete setpoints for each ambient condition $\mathbf{u}(\mathbf{W}_a^{\text{bin}})$. As for the original problem, even the new one is subjected to damage (Eq. 4b) and bound (Eq. 4c) constraints.

This reformulation drastically decreases the computational cost of the problem, as one does not solve anymore for time series of control setpoints, but for control setpoints for each combination of binned atmospheric parameters, as commonly done in standard power-boosting WFC.

As for this standard approach, even here the optimization is performed offline, while online – once the current ambient conditions have been determined – the corresponding setpoints are first sourced from the pre-computed solution and then dispatched in open loop to the turbines in the farm (???). This means that all expensive calculations are relegated to the offline case, and that the setpoints can be verified for safety before being implemented in the field.

2.7 Reference control formulations

In addition to the proposed formulation, in the following referred to as “*Max-profit*”, another five WFC strategies are used in Sect. 4 for comparing in various scenarios the performance of alternative control strategies. We consider greedy control, power-boosting control, and three load-aware formulations that account for component fatigue. Except for the greedy case, all formulations use wake steering, subject to the same yaw misalignment bounds. A brief description of each of these WFC strategies is provided below.

1. “*Greedy*”. This formulation represents the standard operational mode for most existing wind farms. Following this approach, each turbine in the farm operates independently from the others to maximize its own power capture, by orienting its rotor to face the local wind direction. This formulation serves as a baseline for comparison.
2. “*Max-power*”. This formulation corresponds to the classical AEP maximization (or power boosting) by wake steering (?). The optimization problem is ~~subjected to the bound constraint Eq. (1c) on the decision variable $\mathbf{u}$, where $\mathbf{u}^{\min}$ and $\mathbf{u}^{\max}$ denote the lower and upper bounds, respectively, formulated as~~

$$\max_{\mathbf{u}} \sum_{y \in Y} \sum_{\text{bin} \in \text{Bins}} f_{a,y}^{\text{bin}} \sum_{i=1}^N P_i(\mathbf{u}(\mathbf{W}_a^{\text{bin}})), \quad (5a)$$

subject to

$$\mathbf{u}_i^{\min} \leq \mathbf{u}_i(\mathbf{W}_a^{\text{bin}}) \leq \mathbf{u}_i^{\max} \quad \forall i \in N, \forall \text{bin} \in \text{Bins}. \quad (5b)$$

The cost function represents the overall farm power output for all ambient conditions with their probabilities during the lifetime duration. Maximizing this quantity corresponds to maximizing revenue for a fixed price of electricity.

3. “*Load-constrained*”. The formulation targets revenue while limiting loading on the turbine components, and it is adapted from the discussions in ? and ?. The optimization problem is expressed as

$$\max_{\mathbf{u}} \sum_{y \in Y} \sum_{\text{bin} \in \text{Bins}} f_{a,y}^{\text{bin}} \text{Revenue}(\mathbf{u}(\mathbf{W}_a^{\text{bin}}), y), \quad (6a)$$

subject to

$$D_i^c(\mathbf{u}(\mathbf{W}_a^{\text{bin}})) \leq \mathbb{K} \max_i D_i^c(\mathbf{W}_a^{\text{bin}}) \quad \forall i \in N, \forall c \in C, \forall \text{bin} \in \text{Bins}, \quad (6b)$$

$$\mathbf{u}_i^{\min} \leq \mathbf{u}_i(\mathbf{W}_a^{\text{bin}}) \leq \mathbf{u}_i^{\max} \quad \forall i \in N, \forall \text{bin} \in \text{Bins}. \quad (6c)$$

The damage limit for a given component c is defined as the maximum damage observed across all turbines in the baseline greedy case, scaled by a factor $\mathbb{K}$. This parameter represents the safety factor typically applied during the turbine design phase to account for additional loading during operation and all other uncertainties. The damage limit is set individually for each ambient wind condition bin $\mathbf{W}_a^{\text{bin}}$ within the full set of wind conditions, $\mathbf{W}_a$.

4. “*Load-balance*”. This formulation balances power production and loading, following the discussions in ?, ?, and ?. The optimization problem is expressed as

$$\max_{\mathbf{u}} \left(\mathbb{W}_P \sum_{y \in Y} \sum_{\text{bin} \in \text{Bins}} f_{a,y}^{\text{bin}} \sum_{i=1}^N P_i(\mathbf{u}(\mathbf{W}_a^{\text{bin}})) - \sum_{c \in C} \mathbb{W}_D^c \sum_{\text{bin} \in \text{Bins}} \max_i D_i^c(\mathbf{u}(\mathbf{W}_a^{\text{bin}})) \right), \quad (7a)$$

subject to

$$\mathbf{u}_i^{\min} \leq \mathbf{u}_i(\mathbf{W}_a^{\text{bin}}) \leq \mathbf{u}_i^{\max} \quad \forall i \in N, \forall \text{bin} \in \text{Bins}. \quad (7b)$$

The cost function is composed of two terms. The first represents the total power generated by the wind farm. The second term is designed to prioritize turbines with a higher available fatigue budget, stressing them more than those with greater accumulated damage, thereby ensuring a more balanced load distribution throughout the farm. To achieve this effect, function $\max_i D_i^c$ selects the maximum damage accrued across all turbines for a given component c . The weight parameters $\mathbb{W}_P$ and $\mathbb{W}_D^c$ are used to adjust the relative importance of the two terms in the objective function.

5. “Lifetime-revenue”. This formulation is a modification of the proposed one, where the optimization objective targets only revenue instead of profit, and writes

$$\max_{\mathbf{u}} \sum_{y \in Y} \sum_{\text{bin} \in \text{Bins}} f_{a,y}^{\text{bin}} \text{Revenue}(\mathbf{u}(\mathbf{W}_a^{\text{bin}}), y), \quad (8a)$$

subject to

$$D_i^c(T) \leq D_{\text{ref}}^c \quad \forall i \in N, \forall c \in C, \quad (8b)$$

$$\mathbf{u}_i^{\min} \leq \mathbf{u}_i(\mathbf{W}_a^{\text{bin}}) \leq \mathbf{u}_i^{\max} \quad \forall i \in N, \forall \text{bin} \in \text{Bins}. \quad (8c)$$

3 Optimization framework

375 ~~Optimization framework with block chart of the model toolchain.~~

This section describes the optimization framework ~~designed that was developed~~ to compute the optimal control setpoints for a wind farm ~~using according to~~ the proposed economic lifetime-aware WFC strategy. ~~Elements of this framework are also used for implementing the alternative strategies described in Sect. 2.7.~~ An overview of ~~this the~~ framework, including its underlying models, is depicted in Fig. 1. The framework takes ~~as input binned~~ site-specific wind conditions ~~$\mathbf{W}_a$ as input $\mathbf{W}_a^{\text{bin}}$ with their~~ ~~probabilities of occurrence f_a^{bin}~~ and generates optimal control schedules ~~$\mathbf{u}(\mathbf{W}_a)$ $\mathbf{u}(\mathbf{W}_a^{\text{bin}})$~~ for all turbines in the form of look-up tables.

~~The equations in the grey-shaded region at the top of Fig. 1 highlight two key aspects of the proposed strategy: (i) the use of an economic merit function (Eq. 1a), and (ii) the lifetime constraint (Eq. 1b). Additionally, bounds on the control variables are enforced (Eq. 1c).~~

385 The ~~underlying model set is represented by boxes with different color outlines in Fig. 1. The model model~~ toolchain consists of several components: a ~~wind farm fast wind farm engineering~~ flow model (Sect. 3.1), a fatigue load surrogate model ~~-(Sect. 3.2), a damage estimation model -(Sect. 3.3), and an economic model -of profit~~ (Sect. 3.4).

The upper branch of the toolchain follows established research and industry practices for estimating the power production of ~~the wind farm a wind farm~~ (?), which is then converted into ~~economic~~-revenue by the economic model. The ~~lower~~ ~~central~~ branch tackles the more complex task of estimating the impact of the operational strategy on ~~the fatigue of the~~ turbine components.

~~In this process, To this end, for each turbine, its~~ rotor inflow conditions are ~~translated first mapped~~ into DELs, which are then aggregated to estimate ~~the damage to each specific component. The accumulated damage serves two purposes. First, the lifetime constraint evaluates whether the component damage is within the chosen threshold. Second, the economic model uses~~ ~~the accumulated damage damage. In turn, the accumulated damage is used~~ to calculate the associated maintenance costs.

By combining the upper branch (revenue) and the ~~lower branch (costs central branch (cost))~~, the framework generates the ~~economic profit metric, which is the objective of the optimization. The following subsections provide a detailed discussion~~




Figure 1. Graphical visualization of the computational framework supporting the implementation of the economic lifetime-aware WFC strategy.

of each merit function and constraint that define the optimization problem. An optimization algorithm is used to solve the resulting constrained maximization, which finally yields the control setpoints, as shown in the lower branch of the figure.

Each individual model block is discussed in the following.

3.1 Wind farm flow model

The flow model simulates wake interactions within the wind farm and maps, first mapping the ambient site inflow $\mathbf{W}_a$ to turbine-specific inflow conditions and power production inflows, and then estimating the production of each turbine by a power model. Steady-state engineering wake models offer a suitable compromise between model accuracy and computational expensecost, which is the reason for one of the main reasons behind their widespread use within industry and research in industry and academia (?). Furthermore, their modular architecture and tunable parameters allow for calibrating models to better match the site conditions site-specific calibration. The present study uses the open-source toolbox FLORIS-Floris (?). Real-world wind farm flow fields are influenced by additional factors such as terrain and farm boundary layer interactions terrain-induced effects and interactions between the farm and the atmospheric boundary layer. Models of these effects are being developed -(?), or can be learned-learnt, augmented and/or tuned from operational data -(?).

Figure 2a depicts the hub height flow field for an example exemplary two-turbine cluster for $\mathbf{W}_a = [8 \text{ ms}^{-1}, 275^\circ, 12\%]$. The two branches in Fig. 1 show that the flow model has two distinct outputs. First, the aggregated farm power production, a routine output of engineering wake models, is required to estimate economic revenue. Second, in $\mathbf{W}_a = \{8 \text{ ms}^{-1}, 275^\circ, 12\%\}$. The flow field predicted by the model at each turbine rotor disk is first used, through a turbine-specific power model, to compute the power output of each turbine and then of the whole farm. As shown in the upper branch of Fig. 1, power is then fed into the economic model to estimate revenue. In the lower branch, the surrogate load model requires local wind speed and tur-

bulence intensity conditions generated by the flow model at each specific turbine ~~as input~~. While the rotor inflow field is also available from engineering wake models, it is typically not used directly to compute loads and, therefore, needs further processing. Figure 1 are provided as input to a surrogate load model (see Sect. 3.2).




Figure 2. Streamwise velocity field for a two-turbine cluster, modelled with Floris. (a) Horizontal plane at hub height. The indicated sector visualizes the wind direction sweep used later in Sect. A1. (b) Distribution of the longitudinal flow speed just in front of the downstream turbine (see dashed line in panel a). The figure includes the outline of the rotor, and its left and right quadrants used by the surrogate load model (see Sect. 3.2).

420 Figure 2b illustrates the streamwise velocity field just upstream of the rotor of the second turbine, as extracted from **FLORIS**.

The Gaussian wake model used in this study captures the trend of decreasing velocity on the right side of the rotor during half-wake impingement (?). Floris. In this work we use the Gaussian wake deficit model (?) because of its ability to capture the non-uniform distribution of the flow speed in partial wake impingements, as the one shown in the figure. The model includes the tuning parameters k_α , k_β , k_a , and k_b that, in pairs, control the far-wake onset and decay, respectively. In contrast, simpler submodels, such as the the simpler Jensen velocity model, would only provide a crude provides only for a cruder representation of the inflow due to their the assumption of a simple top-hat profile (?). In combination with the load model, it was observed that the baseline configuration of FLORIS does not fully capture all the necessary flow effects. The baseline wake-added turbulence by ? are found to be inaccurate (?); it was found that such level of approximation is not accurate enough when

425 loading needs to be computed.

Wake-added turbulence is another important parameter for an accurate estimation of loading. In particular, the model by ? was found not to be sufficiently accurate in partial wake situations. Similarly to the Jensen model in the just-discussed case of wake deficit, this is primarily due to the overly simplified, top-hat-like turbulence profile profile of turbulence assumed by the Crespo-Hernandez model. To address this problem, the model was replaced with the one of limitation, we implemented the model proposed by ?, which features an axisymmetric double-Gaussian TI profile.

435

Additionally, the model equations were enhanced by making the parameters k_G , k_σ , and k_{r0} tunable, which respectively allow for the calibration of the maximum added TI, of the profile width, and of the peak distance, respectively. ~~For the velocity deficit, the Gaussian wake model by ? provided a sufficiently accurate baseline for calibration. This model includes tuning parameters k_α , k_β , k_a , and k_b which, in pairs, control the far-wake onset and decay, respectively.~~

440 The “wind farm as a sensor” ~~methodology~~ method was employed for the simultaneous calibration of ~~these submodel parameters (??)~~. The effects of these improvements are ~~discussed next in conjunction with the load surrogate model in Sect. A~~. the parameters of the flow model (??), as explained more in detail in the next section.

Streamwise velocity field for a two-turbine cluster, modelled with FLORIS. (a) horizontal plane at hub height for a wind direction $\Gamma = 275^\circ$. The indicated sector visualizes the wind direction sweep of Fig. A1. The dashed line represents the position
445 of the vertical plane in front of the downstream turbine in panel (b). The figure includes the outline of the rotor, and its left and right quadrants used by the surrogate load model.

3.2 Load ~~surrogate~~ model

Following the lower branch in Fig. 1, the load model ~~block is the second step in estimating~~ is tasked with mapping turbine-specific inflow conditions to fatigue loads, which in turn enable the estimation of turbine lifetime and maintenance costs. As surrogate
450 models are relatively new, the present section gives a brief review of the formulation that was used here.

The load surrogate model predicts short-term DELs for monitored components (load channels), which are then used to calculate ~~the damage metric~~ damage. Loads depend primarily on the local inflow conditions at the rotor disk of a turbine. As shown in Fig. 2b, impinging wakes significantly alter the characteristics of the incoming flow field. Additionally, ~~loads~~ self-loads are influenced by the control actions ~~of the~~ performed by a turbine. For example, ~~performing wake steering by~~
455 ~~misaligning the upstream turbine~~ misaligning a turbine to steer its wake can affect the loading of its various components to a different degree ~~(?)~~.

~~Typically, complex aero-servo-elastic~~ Aeroservoelastic simulations are required to estimate DELs at the component level. ~~Running~~ To include the effects of wakes, here we use FAST.Farm (?), a tool that computes the generation of wakes with their meandering and evolution (?) in a turbulent flow field, here based on the Kaimal model (?), together with the impingement
460 of the wakes on downstream wind turbines. We use the extension proposed by ? to provide a more accurate description of wake-added turbulence. Both wake-shedding and wake-impinged wind turbines are modelled with the aeroservoelastic simulator OpenFAST (?).

However, running these simulations within ~~the optimization framework is an~~ optimization, as in the present case, is clearly computationally infeasible. ~~This limitation~~

465 In fact, the complexity and cost of the simulations needed to predict DELs has led to ~~increasing interest~~ an increasing interest of the recent literature in load surrogates ~~in both research and industry. Several approaches.~~ These are fast models that predict DELs by “interpolating” within a large database of pre-computed results. While the generation of the database may require a very substantial number of complex simulations and therefore a high computational effort, the sourcing from this substantial data pool of the DELs that are relevant to a given condition can be done very efficiently at run time. Hence, load surrogates are

470 key for the inclusion of loads in any situation where one needs to compute DELs in a large number of different conditions, as
for the present control optimization.

Several approaches for load surrogates have been proposed in the literature ~~-(????)~~. However, these methods are often
~~site-specific, requiring retraining when the layout changes~~layout-specific, requiring a dedicated training for any given wind
farm, or fail to account for the ~~complex effects of yaw misalignment~~effects of control actions on self-induced loads and those
475 experienced by downstream turbines.

The present work adopts the load surrogate approach proposed by ~~?~~. This surrogate is location-agnostic, in the sense that a
single model can predict loads at any turbine pad within an arbitrary wind farm. This ~~surrogate is also~~ is obtained by mapping,
through a neural network, some selected features of the inflow at the turbine rotor disk with the DELs experienced by chosen
components. Since the mapping is between the inflow at a turbine and the DELs of that same turbine, it is not necessary to know
480 where the turbine is located, but only its wind inflow. The present surrogate is also control-oriented, ~~in the sense that the effect~~
~~of control actions (yaw misalignment or~~ in the sense that the effects of control actions (yaw misalignment or down-regulation)
~~on fatigue is inherently captured.)~~ on self-fatigue are accounted for.

3.2.1 Model design

Given a load time series at a selected component location, DELs are estimated for a 10-minute time window using the Rainflow
485 ~~counting algorithm and the classical formula~~

$$DEL = \left(\frac{\sum_i (n_i L_i^m)}{N_{eq}} \right)^{\frac{1}{m}},$$

~~as described in ? and in ?. Here, n_i is the number of load cycles of amplitude L_i , N_{eq} is the equivalent number of cycles (taken~~
~~conventionally to be equal to $2 \cdot 10^6$ cycles over a lifetime of 20 years), and m is the Wöhler exponent of the S-N curve of the~~
~~component material. In this work, two load channels are considered:-~~

- 490
- ~~– Blade root: resulting from the projection of the blade root in-plane and out-of-plane bending moments along the maximum~~
~~damage direction.-~~
 - ~~– Tower base: resulting from the projection of the tower bottom fore-aft and side-side bending moments along the maximum~~
~~damage direction.-~~

The material constants were chosen as $m = 10$ for the polymer of the blades and $m = 4$ for the steel of the tower ~~(??)~~. For the
495 ~~detailed methodology on projecting DELs along maximum damage direction, see ?.-~~

~~A large number of wind farm aero-servo-elastic simulations were conducted using FAST.Farm (??), providing the dataset~~
~~for training the machine learning surrogate model. The version of FAST.Farm used here includes an extension proposed by ?,~~
~~which provides a more accurate description of wake-added turbulence.-~~

~~A total of 1,908 FAST.Farm simulations were conducted using an array of three IEA 3.4 MW reference wind turbines (??).~~
500 ~~The simulations consider variations of several parameters, including ambient wind speed, turbulence intensity, wind direction,~~

and yaw misalignment of the first two turbines. Synthetic ambient turbulent wind fields, each 10 minutes in length, were generated using TurbSim with the Kaimal model (?). For each simulation case, 24 turbulent seeds were run to ensure statistical convergence. The resulting dataset was then used to train by machine learning the surrogate model, which establishes a relationship between the relevant turbine inflow and operating conditions and the observed load response.

505 3.2.1 Inflow characterization

The description of the turbine inflow that serves as ~~model input~~ input to the load surrogate has a crucial importance. Simply using the fully resolved velocity field ~~is not an ideal~~ generated in FAST.Farm or other similar mid or high-fidelity models is not a practical choice, as machine learning models tend to “learn the noise” present in the turbulence. Therefore, in addition to requiring a very large number of inputs, the model would perform well on the training ~~data set~~ but poorly on test data (?). To
510 avoid this overfitting problem, it is necessary to define suitable reduced-order quantities that minimize noise while retaining enough ~~information~~ informational content to adequately characterize the flow field from the point of view of its effects on DELs.

Additionally, the inflow description must account for the fact that two flow models of different fidelity are used during two different steps in the process. During ~~dataset generation~~ the generation of the dataset later used to train the surrogate, a mid-
515 fidelity dynamic wake meandering model (, here implemented in FAST.Farm) ~~dynamically simulates the wake displacement (?)~~ , simulates in the time domain the wake evolution caused by ambient eddies. ~~The model also relies on an engineering description of wake decay due to smaller-scale turbulence (?)~~. In contrast, the optimization framework ~~uses a~~ used to synthesize the farm control laws is necessarily based on a computationally cheaper low-fidelity, steady-state engineering ~~wake model~~ (flow model, here implemented in FLORIS) ~~to estimate the turbine inflow, as shown for streamwise velocity in Fig. 2b. Floris (?)~~. Although
520 both models share some similarities, ~~the engineering model cannot produce predictions with the same level of fidelity as the mid-fidelity model,~~ there is an unavoidable mismatch between the flowfields generated by the steady-state engineering model and the mid-fidelity dynamic one.

~~To address this problem, here we used~~ This problem is addressed here as follows.

First, the default Crespo-Hernandez (?) top-hat wake-added turbulence model in Floris was replaced with the ? model,
525 because of the similarity of its double-peaked TI with the model implemented in FAST.Farm (?).

Second, to improve the match between the engineering steady-state model and the dynamic one, all free model parameters of the former were calibrated using the dataset generated with the latter. The calibration of the parameters was obtained with the “wind farm as a sensor” method, which uses a maximum likelihood approach to identify corrections to the baseline values of the parameters (?). This model tuning is obtained as the optimization of a cost function combining the residuals for power and rotor-average TI on each rotor disk in the farm for all ambient conditions of interest (see gray branch in Fig. 1).
530

Third, we use a slightly modified version of the inflow characterization, compared to the one presented in ?. Instead of relying on sector-averaged inflow quantities, this work utilizes three rotor-averaged inflow quantities (RAIQs), namely the rotor-averaged wind speed (RAWS), the rotor-averaged turbulence intensity (RATI), and the rotor-averaged horizontal shear (RAHS). This approach simplifies the complex local inflow across the rotor by reducing it to only three representative variables.

535 Although the sector-averaged inflow quantities of ? offer an improved accuracy, the present simpler approach still contains sufficient information to capture DELs .~~The sketch superimposed to the flow field in Fig. 2b provides a visualization of the definitions of the RAIQs. The rotor-averaged wind speed (RAWS) and their trends.~~

The RAWS is defined as the mean wind speed across the rotor disk, i.e.

$$\text{RAWS} = \int_0^{2\pi} \int_0^R U(r, \psi) dr d\psi \quad (9)$$

540 Here, $U(r, \psi)$ is the local wind speed at a given point, with r and ψ being the radial and azimuthal coordinates, respectively. Similarly, the ~~rotor-averaged turbulence intensity (RATI)~~ RATI is defined as the mean TI across the rotor disk and ~~written as~~ is written as

$$\text{RATI} = \int_0^{2\pi} \int_0^R I(r, \psi) dr d\psi \quad (10)$$

where $I(r, \psi)$ is the local TI at a given point on the inflow plane. Finally, the ~~rotor-averaged horizontal shear (RAHS)~~ RAHS represents the heterogeneity of wind speed across the rotor disk in the lateral direction. RAHS is particularly useful for characterizing cases of full or partial wake overlap. The definition is based on ~~?~~ and ? and ?, using the left and right quadrants of the rotor disk, as visualized in Fig. 2b. The left and right quadrant-equivalent wind speeds can be calculated as

$$U_{\text{right}} = \int_{\pi/4}^{3/4\pi} \int_0^R U(r, \psi) dr d\psi \quad (11a)$$

$$550 \quad U_{\text{left}} = \int_{5/4\pi}^{7/4\pi} \int_0^R U(r, \psi) dr d\psi \quad (11b)$$

The horizontal shear coefficient RAHS is estimated under the assumption of a linear variation in wind speed along the lateral direction. It is further assumed that the quadrants sample the inflow at ~~positions corresponding to~~ $\pm 2/3$ of the rotor radius, as described by ~~?~~ and ? and ?. Thus, RAHS can be calculated as

$$\text{RAHS} = \frac{3}{2} \frac{U_{\text{right}} - U_{\text{left}}}{U_{\text{right}} + U_{\text{left}}} \quad (12)$$

555 An ~~illustrative~~ example of the ~~RAIQs, as estimated by the engineering wake model, is shown in the three panels~~ behavior of these three RAIQs is illustrated in Appendix A1 with reference to a wind direction sweep for the two turbines of Fig. A1. ~~The inflow situation corresponds to the downstream turbine in the example illustrated in Fig. 2, for a wind direction sweep of $250^\circ < \Gamma < 290^\circ$, where the wake of the upstream turbine makes a full pass over the downstream rotor. The ambient wind speed and TI are held constant at 8 ms^{-1} and 12%, respectively.~~ 2.

560 Figure A1 clearly shows that the RAWS is at its minimum for the full wake situation ($\Gamma = 270^\circ$)

3.2.2 Damage equivalent loads

Given a load time series at a selected component location obtained with the FAST.Farm simulations, DELs are estimated for a 10-minute time window using the Rainflow counting algorithm and the classical formula

$$DEL = \left(\frac{\sum_i (n_i L_i^m)}{N_{eq}} \right)^{\frac{1}{m}}, \quad (13)$$

as described for example in ? and in ?. Here, n_i is the number of load cycles of amplitude L_i , N_{eq} is the equivalent number of cycles (taken conventionally to be equal to $2 \cdot 10^6$ cycles over a lifetime of 20 years), and m is the Wöhler exponent of the S-N curve of the component material. Typical values for the material constants are $m = 10$ for the polymer of the blades, and progressively recovers toward the ambient wind speed as the wind direction moves away from this value. In contrast, the RATI is largest in full-wake due to the wake-added turbulence. As the turbine experiences less wake impingement, $m = 4$ for the steel of the tower (??).

In this work, we consider two load channels, one at blade root and the other at tower base. Both are obtained by projecting the bending moment components (flapwise and edgewise for blades, fore-aft and side-side for the tower) along the maximum damage direction. See ? for details on this projection.

3.2.3 Surrogate load model

After having defined the inputs (RAIQs) and outputs (DELs), the RAHS shows the expected asymmetry, characteristic of left and right partial-waking. For both the full-wake ($\Gamma = 270^\circ$) and no-wake ($\Gamma < 255^\circ$ or $\Gamma > 285^\circ$) scenarios, the RAHS is close to zero, as there is no significant wind speed asymmetry across the lateral direction of the rotor disk. surrogate load model can be synthesized from a sufficiently complete pool of data. The dataset is obtained by running wind farm aeroservoelastic simulations for a given turbine type and a reference farm layout, chosen to create single and multiple wake overlaps. Here, these simulations were conducted for a regular three-by-three layout. However, during partial-waking situations, the RAHS is positive for $255^\circ < \Gamma < 270^\circ$ and negative for $270^\circ < \Gamma < 285^\circ$. These trends in RAWS, RATI, and RAHS further contribute to the understanding of the results presented in Sect. 4.1.2. the resulting surrogate is applicable to other layouts, as long as the flow and load conditions experienced by the turbines in these different layouts can be inferred from the ones present in the reference one.

Example of the RAWS (left), RATI (middle), and RAHS (right) for a waked turbine for varying wind direction, for an ambient wind speed of 8 ms^{-1} and TI of 12% at a downstream distance of 6 diameters from the wake-shedding turbine. For a wind direction of 270° , the array of turbines is aligned with the incoming wind. The simulations consider variations of the ambient parameters (here wind speed, direction, and turbulence intensity) of synthetic turbulent wind fields of 10-minute duration, and of the control setpoints (here, yaw misalignment). For each simulation case, enough turbulent seeds (here, 24) should be run to ensure statistical convergence.

3.2.4 Model performance

After establishing the inflow characterization, the surrogate load model can be synthesized from the dataset. For each simulation case and each turbine, the RAIQs were ~~are~~ recorded from the flow field, along with the ~~projected DELs at blade root and tower base.~~

595 Two artificial neural networks (ANNs) were ~~corresponding DELs at the components of interest. An artificial neural network (ANN) is~~ then trained to predict ~~the blade root DELs and tower base DELs, respectively, using DELs from~~ the RAWs, RATI, RAHS, and control ~~setpoint variables~~ ~~setpoints~~ as inputs. ~~The Matlab Deep Learning Toolbox was used to create and train the ANNs (?). The loss function is defined as the mean squared error, which is minimized using Bayesian regularization. The~~ Each load channel is associated to a separate ANN. In this work, ANNs consist of a single hidden layer with 30 neurons, using
600 hyperbolic tangent sigmoid activation functions, and are implemented and trained using the Matlab Deep Learning Toolbox (?). Training is based on the minimization of mean squared errors, using Bayesian regularization.

Performance of the load surrogate models for the projected blade root DELs (a) and projected tower base DELs (b) when the input RAIQs are estimated with FAST.Farm. Root-mean-square percentage errors (RMSPE) are used as performance metrics. Half of the dataset was used for the training of the ANN (in blue), while the other half was used for testing the
605 trained ANN (in orange). Figure A2 illustrates the performance of the load surrogate model when the inputs are directly estimated from FAST.Farm for both the training and testing datasets. The root-mean-square percentage error (RMSPE) is used as performance metric. The scatter plots for both the training and testing datasets are similar, indicating that the ANN training database was sufficiently large to ensure a correct generalization. An example of the quality of the load surrogate is illustrated in Appendix A1. Results indicate that DELs can be estimated with good accuracy when RAIQs are computed from the inflows
610 generated by the fast engineering wake model (Floris) (which is indispensable for the application pursued here, because of computational cost reasons), notwithstanding the unavoidable differences with the inflows produced by the higher-fidelity dynamic model (FAST.Farm).

Predicted DELs when the RAIQs are computed from the FLORIS baseline (in black) and the calibrated wind farm flow models (in red). The plots include both the training and testing datasets, filtered for cases with $I = 12\%$. (a) Projected blade
615 root DEL; (b) projected tower base DEL.

Figure A3 visualizes the performance of the combined engineering wake and load surrogate models. As discussed in Sect. 3.1, load estimates based on RAIQs from FLORIS in its baseline configuration were not satisfactory. To address this problem, the wake-added turbulence model in FLORIS was replaced with the ? model, and the “wind farm as a sensor” was used for calibrating all model parameters. The training data for this calibration came from the previously described set of
620 FAST.Farm simulations. A maximum likelihood approach was used to identify the correction of wake model parameters, based on a cost function combining the residuals for power and RATI. To illustrate the tuning improvements, Fig. A3 compares the estimated DELs using the load surrogate linked with the baseline version of FLORIS (shown as black dots) with the ones obtained with the calibrated flow model (shown as red dots). The results show considerable improvements. In particular, panel

(b) highlights the impact of the flow model adjustments, as the baseline model significantly underpredicts the tower base loads. This is primarily due to the crude estimation of RATI in the baseline model, which uses the ? model.

3.3 Damage model

The load model provides a set of DELs for each site ambient condition in the set $\mathbf{W}_a$. However, DELs are primarily representative quantities used to compare the effects of different operational conditions. To quantify the actual impact of fatigue on each component, DELs must be converted into mechanical fatigue damage, which must then be aggregated across the range of ambient wind and operating conditions experienced over the lifetime of the turbine.

Accurately estimating the actual damage suffered by a component requires detailed information about its shape, geometry, and structural characteristics. However, such level of detail is beyond the scope of this work, and – more in general – it is typically only achievable by the designer of that component. Therefore, here we adopt the approach outlined in ?, where the fatigue damage $D_{\Delta t}^c$ for a given equivalent load $DEL_{\Delta t}^c$ over the time duration Δt can be calculated as is calculated as

$$D_{\Delta t}^c = \frac{(DEL_{\Delta t}^c)^m \cdot N_{\Delta t}^{eq}}{(2 \cdot L_u)^m} \quad \forall c \in C, \quad (14)$$

where L_u represents the ultimate load obtained using the material stress curve (?).

3.4 Economic model

The economic model estimates the overall (lifetime) net economic profit by combining revenue from power production with cost from fatigue damage, yielding the optimization objective in Eq. (1a) (and hence also in Eq. 4a). It should be noted that the economic merit function used in this study is a simplified version sourced from the open literature. While this provides for a useful simple approximation, it is essential to acknowledge that wind farm operators and developers, with access to more detailed operational data, are in a much better position to create more precise economic models than the one used here. The net economic profit is calculated as the difference between the accrued revenue and the incurred cost over the wind farm lifetime T .

3.4.1 Revenue model

The revenue model estimates the economic value of supplying the wind farm power to the electricity grid as

$$\text{Revenue}(\mathbf{W}_a, \mathbf{u}) = \mathbb{R}_1(\mathbf{W}_a) \cdot \sum_{i=1}^N P_i(\mathbf{u}(\mathbf{W}_a, \mathbf{u})), \quad (15)$$

where $\mathbb{R}_1$ is the electricity market price, the P_i denotes the power generated by i -th turbine, while the total power yield of the wind farm is $\sum_{i=1}^N P_i$, and P_i denotes the power generated by i -th turbine. As explained, yield depends on ambient conditions $\mathbf{W}_a$ and the control actions $\mathbf{u}$ that are triggered in response to them.

Electricity market prices are volatile and vary across years and regions. The power output of the wind farm directly influences a wind farm may directly influence these pricing patterns, especially as the share of wind power in total generation

increases ~~(???????)~~(?). During periods of high wind generation, energy prices typically decrease. Any additional power generated by the wind farm during these hours results in lower revenue, creating a negative correlation between wind farm yield and market price.

While the relationship between wind conditions and spot-market prices is complex ~~as discussed in ?~~(?), a simplified model is used in this work. This model captures the dependency of market price on ambient wind speed through a polynomial equation, which is sufficient to represent the observed negative correlation ~~(?)~~. The ~~following expression is considered in this work~~, model writes

$$\mathbb{R}_1(\mathbf{W}_a) = \mathbb{R}_{\text{avg}} \cdot (\mathbb{r}_1 \cdot U^3 + \mathbb{r}_2 \cdot U^2 + \mathbb{r}_3 \cdot U + \mathbb{r}_4), \quad (16)$$

where ~~$\mathbb{r}_1, \mathbb{r}_2, \mathbb{r}_3, \mathbb{r}_4$~~ denote the parameters of the third order polynomial, and U is the wind speed. A best fit was obtained using the MATLAB Curve Fitting Toolbox ~~(?)~~(?). For improved accuracy, the electricity market price is normalized by the annual average price $\mathbb{R}_{\text{avg}}$, as discussed in ?. For details regarding the performance of this regression model, refer to ?.

Historical spot-market prices from the German energy market, along with the ambient wind speed data from a wind farm in the same market region, were used as input for the regression. ~~For improved accuracy, the electricity market price was normalized by the annual average price $\mathbb{R}_{\text{avg}}$, as discussed in ?~~. ~~For details regarding regression performance, refer to ?~~. The resulting correlations between spot market price and wind speed at two locations in Germany are shown in Fig. 3, which was adapted from Fig. 8 in ?. In the examples considered later in this work, the correlation for the site in the north of Germany was used.

~~The second component~~




Figure 3. Spot market price vs. wind speed at 50 m for a site in the north (a) and in the south (b) of Germany. In the boxplots, the central red mark indicates the median, and the bottom and top blue edges of the box indicate the 25th and 75th percentiles, respectively. Whiskers extend to the most extreme data points not considered as outliers, whereas outliers are plotted individually using a red '+' symbol. Above 15.52 ms^{-1} , the curve is extrapolated. Adapted from Fig. 8 of ?.

3.4.2 Cost model

The second element of the profit equation is the cost model, which estimates the maintenance expenses arising from accumulated damage to the components of each turbine in the wind farm. Maintenance costs are generally divided into two categories: scheduled (preventive) maintenance, and unscheduled (corrective) maintenance (?).

675 Scheduled maintenance involves routine tasks carried out according to a predetermined time plan, often through continuous or periodic monitoring of turbine components. Unscheduled maintenance, on the other hand, is performed after a failure occurs and the operational readiness of the component needs to be restored. Unscheduled maintenance can also be opportunistic, prompted by condition monitoring to prevent imminent failure.

? provides a summary of both maintenance categories. Maintenance costs are primarily driven by factors such as equipment repair or replacement costs, the number of technicians required for repairs, failure rates, and operational downtime losses. Scheduled maintenance costs are generally easier to estimate, as they are based on a predefined maintenance plan. In contrast, unscheduled maintenance costs depend on stochastic factors such as mean time between failures, component reliability, and repair time per failure (?).

This work considers a simplified maintenance cost model, which aims to capture only the fundamental aspects of the complex, variable, and uncertain annual nature of O&M costs. The formulation is inspired by the discussions in ?, ?, ?, ?, and ?. While it is clear that industry has certainly access to more precise cost models, the present one is deemed sufficient as a proof of concept and could be readily replaced with a more accurate one without affecting the general formulation proposed here.

The incurred maintenance cost can be written as costs are written as

$$690 \text{ Cost}(\mathbf{W}_a, \mathbf{u}) = \sum_{c \in C} \frac{\mathbb{C}_{\text{opex}}^c(\mathbf{W}_a)}{D_{\text{ref}}^c} \cdot \sum_{i=1}^N D_i^c(\mathbf{u}(\mathbf{W}_a, \mathbf{u})), \quad (17)$$

where

$$\mathbb{C}_{\text{opex}}^c(\mathbf{W}_a) = \mathbb{C}_1^c \cdot [(\mathbb{R}_1(\mathbf{W}_a) \cdot \text{CF} \cdot \sum_{i=1}^N \mathbb{C}_{3,i}^c \cdot P_i^r) + \mathbb{C}_5^c] + \mathbb{C}_2^c \cdot [(\mathbb{R}_1(\mathbf{W}_a) \cdot \text{CF} \cdot \sum_{i=1}^N \mathbb{C}_{4,i}^c \cdot P_i^r) + \mathbb{C}_6^c].$$

where $\mathbb{C}_{\text{opex}}^c(\mathbf{W}_a)$ are the operation and maintenance costs for component c , while $D_i^c(\mathbf{u}(\mathbf{W}_a))$ is the damage in that component for turbine i , which in turn depends on the control actions $\mathbf{u}$ triggered on all turbines by the ambient conditions $\mathbf{W}_a$.

695 For a given component c , the operation and maintenance expense $\mathbb{C}_{\text{opex}}^c$ is composed of costs are formulated as

$$\begin{aligned} \mathbb{C}_{\text{opex}}^c(\mathbf{W}_a) = & \mathbb{C}_{\text{sch lost}}^c \left(\left(\mathbb{R}(\mathbf{W}_a) \text{CF} \sum_{i=1}^N (\mathbb{C}_{\text{sch down},i}^c P_i^r) \right) + \mathbb{C}_{\text{sch fixed}}^c P^r \right) + \\ & \mathbb{C}_{\text{unsch lost}}^c \left(\left(\mathbb{R}(\mathbf{W}_a) \text{CF} \sum_{i=1}^N (\mathbb{C}_{\text{unsch down},i}^c P_i^r) \right) + \mathbb{C}_{\text{unsch fixed}}^c P^r \right). \end{aligned} \quad (18)$$

The first cost item, depending on coefficients $\mathbb{C}_{\text{sch}}^c$, corresponds to the scheduled maintenance expenses. The second cost item, depending on coefficients $\mathbb{C}_{\text{unsch}}^c$, corresponds to the unscheduled maintenance expenses.

700 Both scheduled and unscheduled expenses include fixed and variable cost items. The fixed costs, denoted as $\mathbb{C}_{5-6}^c$, account for the average repair cost, including are assumed to be proportional to the average turbine rating P^r through the

coefficients $C_{sch, fixed}^c$ (for scheduled maintenance) and $C_{unsh, fixed}^c$ (for unscheduled maintenance). These repair costs account for equipment, labor, and operational personnel expenses.

705 The variable costs represent the opportunity cost incurred due to because of production losses during turbine maintenance. In Eq. (18b), the first cost item, consisting of parameters C_1^c caused by component c . For scheduled maintenance, C_3^c , and C_5^c , corresponds to the scheduled maintenance expenses. The second cost item, consisting of parameters C_2^c , C_4^c , and C_6^c , corresponds to the unscheduled maintenance expenses.

710 The opportunity cost in the variable component is estimated these are proportional through coefficient $C_{sch, lost}^c$ to the revenue lost due to downtime, which is computed as the product of the electricity market price R_T , the average downtime for the component C_{3-4}^c due to turbine maintenance, and the average R and the corresponding lost wind farm yield. The average yield latter quantity is calculated as the product of the sum of the products of the rated power output P_i^r of all turbines outputs P_i^r of each turbine and the turbine component unavailability caused by downtime $C_{sch, down, i}^c$, multiplied by the wind farm capacity factor CF. A similar formula applies to the non-scheduled maintenance, with coefficients $C_{unsh, lost}^c$ and $C_{unsh, down}^c$.

715 The O&M expense is then linearly amortized in Eq. (17) over the reference damage D_{ref}^c to determine the maintenance cost due to the incurred damage D_i^c . The total maintenance cost is obtained by summing the costs for each individual component.

The final estimated net economic profit represents the optimization objective, as shown in. Among other simplifications, the model does not account for the teething and ageing effects that typically lead to higher costs at the beginning and end of the lifetime of an asset.

720 3.4.3 Time discounting

When performing an economic analysis over years-long horizons, it is customary to discount the value of future cash flows, the underlying idea being that the value of one euro in the far future is less than the value of one euro today. Following this approach, the net present value (NPV) of economic profit is calculated as

$$\text{Profit}(\cdot) = \sum_{t=0}^T \frac{\text{Revenue}_t(\cdot) - \text{Cost}_t(\cdot)}{(1 + \mathbb{D})^t}, \quad (19)$$

725 where $\mathbb{D}$ denotes the discount rate (??).

When evaluating a given control policy, it is necessary to use time discounting to reflect risk, inflation, and opportunity cost. However, the situation is different when designing a control policy, for example when solving an optimal control problem such as the one of Eq. (1a). It should be noted that the economic merit function used in this study is a simplified version sourced from the open literature. While this provides for a useful simple approximation, it is essential to acknowledge that wind farm operators and developers, with access to more detailed operational data, are in a much better position to create more precise cost models(1).

In fact, consider a situation where control actions not only depend on the ambient conditions, but are also allowed to change every year; in other words, the same ambient conditions trigger different control actions depending on the year of operation.

Since, by using the NPV, profit is discounted over time, the effects of the control actions tends to fade away as time progresses, which means that their performance can decrease in time without significantly affecting the overall lifetime profit optimization objective. However, when that later time finally comes, one would be obliged to dispatch those less performing control actions. In fact, failure to do so would violate the lifetime constraint, as damage had been computed by using those same control actions. But dispatching the less performing actions would clearly be a suboptimal choice, as better ones had been computed and used in earlier years. This seems to discourage the use of time discounting in the solution of lifetime-aware optimal control problems.

Consider now the situation where discounting is used, but the same better performing control actions are used for each ambient condition, no matter the year of operation (here we are assuming that no other external factors are present, such as for example a change in the state of health of a turbine, which might force one to change the way it is operated). In this case, if every year the same climate is present at the site and the same control actions are used for the each ambient condition, then the same revenue and cost are generated. This means that in Eq. (19) the numerator on the right hand side does not depend on time, and the overall discounted profit is equal to the yearly profit multiplied by the factor $\sum_{t=0}^T 1/(1 + \mathbb{D})^t$. Since multiplying an objective function by a factor does not change the value of the solution (i.e., the control actions), discounting is irrelevant.

Yet another way of looking at the problem of time discounting is to consider that control actions are coupled to each other over time in the optimal control problem of Eq. (1), in the sense that an action performed at a given time affects the actions that can be performed at other times. As a consequence, all time instants are equally important, and time discounting should not be applied when computing control actions.

Because of these reasons, we do not include time-discounting in the economic model when solving the control synthesis optimization problem expressed by Eq. (1). However, we do include time-discounting when evaluating a posteriori the economic performance of a given control policy, as shown in the next section.

4 Results and discussion

~~To evaluate the benefits of the proposed formulation through an application scenario, WFC is deployed on an existing wind farm.~~ The economic lifetime-aware formulation is demonstrated in two exemplary scenarios: a first idealized case characterized by a simplified set of environmental conditions causing very frequent waking, and a second case featuring a more realistic small wind farm exposed to a wind climate derived from an existing meteorological mast. For both scenarios, the wind farm

- ~~is considered to have~~ has already operated for half of its target operational life T of 20 years, and
- it has accrued additional unexpected damage over this duration, which would lead to ~~premature component failure if the~~ a premature failure of a component if its operation continued as usual.

~~In this approach, the lifetime-aware formulation~~ The proposed approach is applied to the remaining lifespan of the assets, taking into account the current state of health resulting from accumulated damage. The ~~proposed formulation focuses on~~ component failure due to accumulated fatigue at the blade root fatigue constraints target the root of the blades and the base

of the tower to avoid their premature failure. These components were chosen because of their critical nature and tower-base. These locations were chosen based on the critical nature of the components and the significant maintenance costs costs that the wind farm operator would face in the event of failure (???). Additionally, in the their experimental validation of the lifetime-aware WFC formulation, limited the fatigue damage assessment to these two components due to technical limits in the sealed turbine models. Results revealed indicate that both the tower base and blade root exhibit distinct responses to varying inflow and control conditions, making them interesting targets also for the present study.

Any component failure leads to maintenance and downtime, as outlined discussed in Sect. 3.4. For simplicity, it is assumed that the fixed costs associated with maintenance, downtime, and the number of maintenance activities remain constant over time, although this is not really true according to a more realistic assumption would be to use the classical bathtub curve. To evaluate lifetime quantities, rather than simulating the entire operational life of the wind farm using time series of wind conditions, only a single year is considered. It is behavior. It is further assumed that the ambient wind conditions for this year are representative of the conditions for the remaining years do not change over time. The set of ambient wind conditions $\mathbf{W}_a$ over the operational lifespan T comprising individual wind condition bins, $\mathbf{W}_a^{\text{bin}}$, and their corresponding and their probabilities of occurrence, f_a^{bin} is are derived from the site wind rose.

In the case studies presented here, only yaw misalignment-based wake steering WFC is considered. Therefore, Wake control is based on steering, and therefore the decision variables in the proposed optimization problem (Eq. 1)-(4) are limited to the yaw-offset set points, $\mathbf{u} = \boldsymbol{\gamma} = [\gamma_1, \dots, \gamma_i, \dots, \gamma_N]$ setpoints, $\mathbf{u} = [\gamma_1, \dots, \gamma_i, \dots, \gamma_N]$, where γ_i is the misalignment of the i th turbine with respect to the incoming wind direction. The bounds on the yaw-offsets, as defined yaw offsets appearing in Eq. (4c), are constrained to a limiting value of (4c) are $\pm 20^\circ$, for all the turbines throughout their lifetime.

The relevant model parameters, including those for the economic model, are detailed load surrogate (see Sect. 3.2 and Appendix A) was trained on the results of 1,908 FAST.Farm simulations, conducted for an array of three IEA 3.4 MW reference wind turbines (?). The damage limit D_{ref}^c appearing in Eq. (4b) is chosen as

$$D_{\text{ref}}^c = \max_i \sum_T D_{i,\text{greedy}}^c(\mathbf{W}_a) \quad \forall i \in N, \quad (20)$$

which corresponds to the maximum accumulated damage over all the turbines under standard greedy control, where no wake steering is used.

The parameters of the cost model expressed by Eq. (18) are reported in Tab. 1. The maintenance cost parameters are derived from the values reported in ?????, which and are based on various databases also referenced in ?.

4.1 Reference control formulations

In addition to the proposed formulation, referred to as “Max-profit”, five alternative WFC strategies are used for comparison. These include greedy control, power-boosting control, and three load-aware formulations that account for component damage. A brief description of each of the considered WFC strategies is provided below.

Table 1. Overview Coefficients of the cost ~~items for the maintenance expense model of Eq.~~ Cost parameters are derived from values reported in the following references:

Cost item	Description	Category	Unit	Value <u>for tower</u>	Value ~~blade Scheduled for blades~~
~~C_1~~ <u>$C_{sch\ lost}$</u>	Number of downtimes	Variable	-	5	5
~~C_2~~ <u>$C_{unsch\ lost}$</u>				0.02	0.05
~~C_3~~ <u>$C_{sch\ down}$</u>	Downtime duration	Variable	hours	96	192
~~C_4~~ <u>$C_{unsch\ down}$</u>					
~~C_5~~ <u>$C_{sch\ fixed}$</u>	Average cost of repair	Fixed	€ kW ⁻¹	3	3
~~C_6~~ <u>$C_{unsch\ fixed}$</u>				5	5

1. ~~The first reference formulation represents the state-of-the-art operation for most existing wind farms. In this approach, each turbine within the farm operates independently from the others by orienting its nacelle to face the local wind direction. This formulation serves as a baseline for comparison and will henceforth be referred to as “Greedy”.~~

2. ~~The second reference formulation corresponds to the classical wake steering strategy to maximize AEP (??). The unconstrained optimization problem~~

$$\max_{\mathbf{u}} \sum_T \sum_{i=1}^N P_i(\mathbf{W}_a, \mathbf{u})$$

~~maximizes the overall farm power output and, in turn, the accrued revenue. This formulation will henceforth be called “Max-power”.~~

3. ~~The third reference formulation builds on the second one by limiting the loading on the turbine components. The strategy is adapted from the formulations discussed in ? and ?.~~

~~The optimization problem can be expressed as~~

$$\max_{\mathbf{u}} \sum_T \text{Revenue}(\mathbf{W}_a, \mathbf{u}),$$

~~subject to~~

$$\underline{D_i^c(\mathbf{W}_a^{\text{bin}}, \mathbf{u}) \leq \mathbb{K}_1 \cdot \max_i D_i^c(\mathbf{W}_a^{\text{bin}}) \ \forall i \in N \ \forall \text{bin} \in \mathbf{W}_a \text{ and } \forall c \in C.}$$

~~The damage limit for a given component c is defined as the maximum damage observed across all turbines in the baseline formulation, scaled by a factor $\mathbb{K}_1$. Parameter $\mathbb{K}_1$ represents the safety factor typically applied during the turbine design phase to account for additional loading during operation and all other uncertainties. The damage limit is set individually for each ambient wind condition bin $\mathbf{W}_a^{\text{bin}}$ within the full set of wind conditions, $\mathbf{W}_a$. This formulation will henceforth be referred to as “Load-constrained”.~~

4. The fourth reference formulation balances the power production and the incurred loading of different components. The formulation has been adopted from the discussions in ?, ?, and ?. The resulting optimization problem can be expressed as-

$$\max_{\mathbf{u}} \sum_T \left(\mathbb{W}_P \cdot \sum_{i=1}^N P_i(\mathbf{W}_a, \mathbf{u}) - \sum_{c=1}^C \mathbb{W}_D^c \cdot \max_i D_i^c(\mathbf{W}_a, \mathbf{u}) \right).$$

In this formulation, the optimization problem maximizes the total power generated by the wind farm, while minimizing the distribution of loading across individual turbines. The approach prioritizes turbines with a higher available fatigue budget, using them more than those with greater accumulated damage, thereby ensuring a more balanced load distribution throughout the farm. The function $\max_i D_i^c$ represents the maximum damage accrued across all turbines for a given component c . The weight parameters, $\mathbb{W}_P$ and $\mathbb{W}_D^c$, are used to adjust the relative importance of the different components in the objective function. These parameters are fine-tuned through a parameter-sweep optimization, yielding the most optimal solution for the given set of ambient wind conditions, $\mathbf{W}_a$. This formulation will henceforth be referred to as “Load-balance”.

5. The fifth reference formulation is a modification of the proposed formulation, where though the optimization problem is subjected to a lifetime constraint, the optimization objective only aims to maximize revenue i.e. it simply writes-

$$\max_{\mathbf{u}} \sum_T \text{Revenue}(\mathbf{W}_a, \mathbf{u}),$$

subject to Eq. (1b). This formulation will henceforth be referred to as “Lifetime-revenue”.

4.1 Case study I: Simple wind farm layout and synthetic wind rose

The first case study evaluates the performance characteristics of the proposed WFC formulation in comparison to reference control strategies, including the baseline greedy one, in a simplified case consisting of three turbines and a synthetic wind rose that generates frequent waking. Results are analyzed in terms of yaw offsets, within a simplified setup. This study provides a comprehensive assessment of optimal yaw offset schedules, power output, component loading, and their impact on lifetime economic performance in terms of revenue, cost, and profit across different WFC formulations. For this purpose, a simple wind farm configuration consisting of three turbines is considered, along with a simplified ambient wind condition subset derived from a synthetic wind rose.

The performance of the resulting optimal yaw offset schedule is validated using FAST-Farm. As for the generation of the load surrogate model database, the turbulent ambient wind field is generated using the Kaimal model in TurbSim tool (?), with 24 turbulent seeds employed to ensure statistical convergence of the results. The yaw angles for each turbine are prescribed as fixed values from the look-up tables in the ElastoDyn module (?). One of the main goals of this first example is to quantify the accuracy of the computational framework that supports the synthesis of the control setpoints. In fact, as discussed in Sect. 3, the framework is based on a fast but low-fidelity engineering flow model (FLORIS), which in general will match with only a

limited accuracy the ground truth. In the absence of an even higher fidelity model of the plant, here the ground truth is assumed to be represented by the mid-fidelity model FAST.Farm. The setpoints computed by the framework for each WFC formulation are dispatched to the FAST.Farm model, and then various quantities of interest are measured based on its transient response. A quantification of the accuracy of the computational framework is finally obtained by comparing predicted and measured quantities of interest.

4.1.1 Layout Farm layout and input wind conditions

Layout and input wind conditions for the simplified case study setup. Wind rose showing ambient input wind speeds and directions (a). Wind farm layout and wakes at hub height for wind blowing from the western direction (b). The turbine layout resembles a typical WFC demonstrator case for wind tunnel experiments and simulations (?). Aligned layouts are also studied in field tests (?). Here we consider three IEA 3.4 MW reference wind turbines, with a rated wind speed of 9.8 ms^{-1} (?). Figure 4b shows the turbines T1, T2, and T3 in the (?) in an aligned layout with 6-diameter-spacing between the turbines. This figure also shows the wakes in the fully aligned layout for a representative ambient wind condition for the baseline. This setup is illustrated in Fig. 4b for the baseline greedy case, where all turbines point straight into the incoming wind. The turbine layout resembles a typical WFC demonstrator case for wind tunnel experiments and simulations (?). Aligned layouts are also studied in field tests (?).




Figure 4. Layout and ambient wind conditions for the simplified case study. Wind rose showing ambient input wind speeds and directions (a). Wind farm layout and wakes at hub height for a wind blowing from the western direction and greedy control (b).

The synthetic wind rose is shown in Fig. 4a, where the individual wind conditions, $\mathbf{W}_a^{\text{bin}}$, are assumed to have an equal probability of occurrence, f_a^{bin} . The purpose of selecting this wind rose is to evaluate the performance of different WFC strategies without the influence of the frequency distribution of the individual wind condition bins. The synthetic wind rose consists of wind direction sectors. The wind direction spans a sector ranging from 250° to 290° , with where the 270° direction aligning is aligned with the turbine layout. Each wind direction sector includes wind speeds two wind speed bins of 8 and

9 ms^{-1} , each with equal probability an equal probability of occurrence. Under these conditions, the baseline *Greedy* formulation results in an overall wake loss of 11%. The purpose of selecting this wind rose is to evaluate the performance of different WFC strategies without the influence of the frequency distribution of the individual wind condition bins.

For all WFC approaches, T1 and T2 perform wake steering according to their respective control policies, whereas T3 always operates in *Greedy control* mode. Moreover, T2 is assumed to have accumulated in the first 10 years of operation 25% more damage to its components than anticipated.

4.1.2 Control schedule, power, and load distribution

Figure 5 displays the yaw-offset schedules derived from the different WFC formulations for an ambient wind speed of 8 ms^{-1} for turbines T1, T2, and T3, respectively from left to right. The x -axis represents axis reports the ambient wind direction sector Γ in degrees, while the y -axis axis indicates the optimal yaw-offset yaw offsets in degrees.




Figure 5. Yaw-offset schedules for T1, T2, and T3 as functions of all the wind direction sectors Γ , at an ambient wind speed of 8 ms^{-1} , for all the formulated-considered WFC strategies. As wake steering is only implemented on T1 and T2, the yaw-offset yaw offset of T3 is constant at 0° .

To evaluate the performance of the various wind farm controllers, transient simulations in turbulent winds are conducted in FAST.Farm where, for each binned wind condition, the yaw offsets are dispatched to the turbines according to the schedules shown in Fig. 5.

Figures 6, 7, and 8 present the resulting power, blade root DELs, and tower base DELs, respectively, obtained by applying the optimal yaw-offset schedules of each WFC formulation shown in Fig. 5. In each figure, the outputs are displayed as panels, with rows corresponding to the six WFC strategies and columns representing the three turbines. The y -axis shows axis reports the absolute value of the output in the respective units, and the x -axis axis shows the ambient wind direction sector Γ in degrees. Solid lines represents the measured-outputs outputs measured in the FAST.Farm simulations (labeled ‘Meas.’), while dashed lines show the corresponding predictions (labeled ‘Pred.’) obtained during the synthesis of the optimal yaw-offsets yaw offsets using the steady-state engineering model Floris.

In all panels of Fig. 5, ~~6~~, ~~6~~, 7, and 8, different colors are used to distinguish the outputs ~~for~~ of the six WFC strategies: the baseline *Greedy* strategy ~~is shown~~ in black, the *Max-power* strategy in orange, the *Load-constrained* strategy in blue, the *Load-balance* strategy in violet, the *Lifetime-revenue* strategy in yellow, and the proposed *Max-profit* strategy in green.




Figure 6. Power output for case study I. The three columns from left to right denote the outputs of the three turbines, T1, T2, and T3, respectively. The six rows from top to bottom denote the six control formulations. In every panel, the x -axis ~~shows~~ ~~axis indicates~~ the ~~input~~-wind direction Γ , the y -axis ~~axis~~ shows the generated power, ~~the solid~~. Solid curves ~~show the represent~~ measured ~~value values~~ from FAST.Farm, ~~and the~~ ~~while~~ dashed curves show the corresponding value predicted by the optimization framework ~~based on Floris.s~~

For all WFC strategies, turbines T2 and T3 operate within the wake of the upstream turbines, resulting in a reduced power output compared to T1, as shown in Fig. 6. In the baseline strategy, the power of the downstream turbines decreases as they enter partial wake conditions ~~for $\Gamma \neq 270^\circ$~~ , with the largest power loss occurring under full waking at $\Gamma = 270^\circ$. ~~The power loss is~~ ~~Power losses are~~ most significant for T3, as it operates under the combined wakes of both upstream turbines. The power produced by the downstream turbines T2 and T3 is influenced by the available wind speed in the flow, which follows the trend of the RAWS, as shown in Fig. A1a.

~~Additionally, the~~ ~~The~~ blade root DELs for the baseline formulation show an asymmetric behavior with respect to the ambient wind direction, as seen in Fig. 7. For $\Gamma < 270^\circ$, the DELs of T2 and T3 are higher than those of T1, while for $\Gamma > 270^\circ$, the DELs of T2 and T3 are ~~lower~~ ~~smaller~~. This asymmetry is primarily driven by the corresponding RAHS, which also exhibits an asymmetric response with respect to Γ (see Fig. A1c).




Figure 7. Blade root projected DELs for case study I. [The organization of the figure and meaning of the line styles are the same as for Fig. 6.](#)




Figure 8. Tower base projected DELs for case study I. [The organization of the figure and meaning of the line styles are the same as for Fig. 6.](#)

In contrast, the tower base DELs for the baseline formulation display a symmetric behavior with respect to Γ , as shown in Fig. 8, with both T2 and T3 experiencing higher DELs than T1. This is because tower base DELs are mainly influenced by the RATI, which is higher for the waked turbines (see Fig. A1b) and exhibits a symmetric pattern with respect to Γ .

For an 8 ms^{-1} ambient wind speed, as shown in Fig. 5, all WFC formulations result in similar yaw-offsets for T1. The yaw-offsets for T1 are generally higher in magnitude than those for T2, as T1 operates in clean inflow conditions, which are typically associated with lower-a smaller total loading compared to the waked inflow conditions experienced by T2. Consequently, T1 has a larger available budget to accommodate additional load increases due to active wake steering, resulting in larger yaw-offsets compared to T2.

For T1, the highest yaw offset for the four damage-aware formulations – *Load-constrained*, *Load-balance*, *Lifetime-revenue*, and *Max-profit* – result in slightly lower yaw-offsets are reached further away from the aligned condition, compared to the *Max-power* strategy. However, The situation is different for T2, these-where the damage-aware strategies lead-to-exhibit a more aggressive wake steering for $\Gamma > 270^\circ$ and a less aggressive steering behavior for $\Gamma < 270^\circ$ compared to the *Max-power* strategy. This is because the yaw schedules are primarily influenced by the-blade loads, which exhibit an asymmetric behavior with respect to Γ . As a result, it is beneficial to employ more aggressive wake steering for $\Gamma > 270^\circ$ as it reduces blade root DELs, while for $\Gamma < 270^\circ$ wake steering increases these same loads.

Compared to the baseline, all other-greedy case, all WFC strategies result in a slight decrease in power for T1 and an increase in power for T2 and T3, as shown in Fig. 6. Consequently, all WFC strategies lead to an overall power gain for the wind farm. Additionally, a comparison between the solid and dashed lines for all strategies reveals that the measured turbine power outputs closely align with the estimated values. This high level of accuracy is due to the precise estimation of RAWs by the engineering flow model used within-by the optimization framework (refer to Sect. 3.1 for details). Since RAWs is the primary driver of power production, the predicted power closely matches the measured one for the yawed turbine T1, the waked turbine T3, and the yawed and waked turbine T2.

Active wake control leads to an increase in blade root DELs for all turbines when $\Gamma < 270^\circ$, and a slight decrease in DELs for $\Gamma > 270^\circ$, as shown in Fig. 7. The variation in blade root DELs is most significant for T2, as it experiences the combined effects of operating in a waked inflow and performing wake steering. Additionally, a comparison of the solid and dashed lines indicates that the blade root DELs are well predicted by the optimization framework (cf. SeeAppendix. A), which effectively captures the trends in load variation with respect to Γ , for T1 (yawed), T3 (waked), and T2 (yawed and waked). This accuracy is especially important for load-aware strategies, where the-optimal-yaw-offset-is-optimal yaw offsets are directly influenced by these variations in loadloading.

For tower-base-loadingthe tower base, as shown in Fig. 8, active wake control leads to a slight increase in DELs for the yawed turbine, while significantly-reducing-it significantly reduces DELs for the downstream turbines. Since tower base DELs are primarily driven by the-RATI, deflections of the upstream wake results-result in lower RATIs for the downstream turbines, leading to reduced loading. Additionally, a comparison of the solid and dashed lines reveals that, for all control strategies, tower loads are consistently under-predicted in partial wake conditions ($\Gamma < 265^\circ$ and $\Gamma > 275^\circ$) and over-predicted in full wake conditions ($265^\circ \leq \Gamma \leq 275^\circ$). However, the optimization framework still captures the trends in tower load variation

935 accurately. This with good accuracy. The observed mismatch is primarily attributed to the presence of rotor up-tilt in FAST.Farm, which leads to vertical wake displacement and vertical shear at the rotor inflow. Since FLORIS does not In Floris, however, the presence of a wake velocity deficit does not change the vertical shear. In addition, Floris does not account for vertical shear, the predictions from wake displacement due to rotor tilt. Therefore, predictions generated by the optimization framework slightly differ from those obtained from FAST.Farm. Despite this, the under- and over-predictions generally cancel each other out for the inflow conditions considered in this case study. The impact of this small net prediction mismatch on lifetime performance is discussed in the following section.

4.1.3 Lifetime consumption

Lifetime consumption of the wind farm for all the WFC strategies for case study I. The solid outline of the horizontal bars in panel (a) denotes the measured value based on FAST.Farm, and the dashed outline in panel (b) denotes the corresponding predicted value from the optimization framework.

945 Figure 9 illustrates the lifetime consumption of the wind farm across for all considered WFC strategies. Lifetime consumption is determined by the accumulated fatigue damage on the components, with. For simplicity, we assume that the wind farm reaching its EOL reaches its end of life (EOL) when any component on of any turbine reaches its EOL. While this is clearly an approximation, it is used here to avoid complicating the analysis further with various possible scenarios. A component reaches its EOL when the accumulated damage D ratio of the accumulated damage $D^c(T)$ (cf. Sect. 3.3) and the reference damage D_{ref}^c equals 1 (cf. Eq. 4b), which corresponds to a lifetime consumption of 100%. Figure ?? displays

950 Figure 9a reports the lifetime consumption resulting from the application of optimal yaw offsets in the validation case study that results from dispatching the optimal yaw offsets to the turbines in the transient FAST.Farm simulations, while Fig. ?? shows the expected lifetime consumption predicted 9b shows the predictions by the optimization framework. In both figures, lifetime consumption is represented as horizontal bars, with each bar corresponding to a specific WFC strategy. The x -axis axis represents time in years, and the y -axis axis lists the six WFC strategies. As detailed in Sect. 4.1.2, the

The performance of the wind farm is assessed over a 20-year operational life (reference EOL), which is the goal desired EOL. The WFC strategies are implemented after the wind farm has operated under the baseline greedy strategy for 10 years (half-time strategy change). Consequently, each bar is divided into shaded regions representing four distinct phases:

- The grey region indicates the lifetime consumed while operating without any optimal yaw offset (baseline strategy) with the baseline greedy strategy.
- The white region represents the lifetime consumed while operating with the different WFC strategies.
- The red region denotes the lost lifetime of the wind farm when the EOL is (i.e., when the accrued damage is equal to the allowed damage limit) is reached before the reference goal EOL.

– The green region represents the gained or extended lifetime, ~~where the consumed life until the reference EOL was less than 100%~~when the accrued damage at the goal EOL is below the allowed damage limit. The remaining useful life is linearly extrapolated to estimate the additional operational years.

A close examination of Fig. ~~??~~9a reveals that, when operating under the *Greedy*, *Max-power*, and *Load-constrained* strategies, the wind farm fails to reach the ~~reference-goal~~ EOL, resulting in a loss of operational years. In contrast, the remaining strategies – *Load-balance*, *Lifetime-revenue*, and *Max-profit* – operate in a more damage-aware manner, leading to a gain in operational years beyond the ~~reference EOL~~goal EOL.

The *Max-power* strategy yields the shortest operational life. This suggests that ~~, while the aggressive wake steering employed in the Max-power strategy (as seen in the orange line plots in the symmetric wake steering that it employed for T2 (cf. Fig. 5),~~ may enhance energy production, ~~it but~~ also accelerates fatigue accumulation~~, even surpassing the accumulation seen in the baseline (no wake steering) case. This outcome is primarily due to the combination of increased fatigue from the pre-damage scenario and the aggressive wake steering implemented on T2 to boost the yield of T3.~~ As a result of its own activity, T2 experiences a significant increase in blade root DELs in the sector $\Gamma < 270^\circ$ (cf. Fig. 7), which is not offset by the modest decrease due to the wake steering of T1, as discussed in Sect. 4.1.2. Consequently, ~~the wind farm suffers this~~ wind turbines suffers a higher fatigue damage, leading to a reduction in operational life.

The *Load-constrained* strategy results in a less aggressive yaw-offset schedule compared to the *Max-power* strategy ~~. This is because the optimization constraints defined in~~ because of the load constraints enforced by Eq. (6b), which account for the loading in individual wind condition bins ~~W_a^{bin} , are rarely active.~~ However, notwithstanding the presence of these load constraints, this strategy still fails to reach the desired lifetime duration. This is because, as discussed in Sect. 1, the fatigue budget for a specific W_a^{bin} bin of ambient conditions does not only depend on the fatigue damage accumulated within that bin, but also on the damage accumulation from ~~other wind condition bins that~~ the other bins the wind farm is exposed to. In other words, from a wind farm lifetime perspective, the control actions for a given wind input condition in load-aware operation are influenced by the relative impact of this condition in relation to others. Since the *Load-constrained* formulation only considers individual bins, it overlooks the cumulative effects of other ~~wind condition~~ bins, which ultimately results in a loss of lifetime.

~~While The situation is different for~~ the *Load-balance* strategy ~~ensures that the wind farm reaches its reference EOL, the optimal balance between fatigue accumulation and energy yield extends the remaining useful life of the wind farm, allowing it to operate beyond its original operational lifetime, which not only reaches but actually extends the intended EOL.~~ However, because the optimization ~~formulation (see Eq. 7) problem used for synthesizing its control setpoints~~ does not explicitly impose constraints to guarantee a desired lifetime ~~, the observed performance can be attributed to the optimal tuning of weights to maximize the overall objective. For more complex (see Eq. 7), the performance that is observed here cannot be expected in general. For other~~ farm layouts and inflow conditions, this formulation may not ensure that the wind farm will operate over a specific lifetime, despite the ~~use of optimized weight tuning tuning of the weights of the two terms of the cost function.~~

As anticipated, the *Lifetime-revenue* strategy ensures that the wind farm reaches the desired operational life. This is because the strategy accounts for the combined loading from all wind conditions the farm experiences throughout its lifetime, while simultaneously optimizing farm power and, consequently, revenue~~, for each individual wind condition bin.~~




Figure 9. Lifetime consumption of the wind farm for all WFC strategies for case study I. The solid outline of the horizontal bars in panel (a) denotes values measured on FAST.Farm simulations, whereas the dashed outline in panel (b) are the corresponding predicted values from the optimization framework.

The proposed *Max-profit* formulation, in addition to explicitly considering lifetime constraints, finds the most profitable way of operating the assets throughout their lifetime. It does so by not only focusing on the power output of the individual turbines, but also considering the total energy yield and balancing it against the ~~operational and maintenance (O&M)~~ costs resulting from accumulated damage. Compared to the *Lifetime-revenue* formulation, the *Max-profit* strategy adopts a more aggressive wake steering in partial wake overlap conditions ($260^\circ < \Gamma < 280^\circ$) for T1, implements one-sided wake steering for $\Gamma > 270^\circ$ to reduce blade root DELs, while avoiding wake steering altogether for $\Gamma < 270^\circ$. This results in an optimal balance between accumulated damage and revenue, enabling the wind farm to operate for the desired duration while maximizing economic profit throughout its operational life.

Figure ~~??~~ 9b shows the estimated lifetime consumption based on the predicted blade root DELs (dashed curves in Fig. 7) and tower base DELs (dashed curves in Fig. 8) for all the WFC strategies. ~~The predicted lifetime consumption is visually represented in the same way as the calculated lifetime consumption in~~ By comparing the predicted quantities of Fig. ??, ~~However, the predicted lifetime consumption features dashed outlines around the horizontal bars, highlighting the expected performance, while the measured lifetime consumption is shown with solid outlines. By comparing the two figures, it is clear~~ 9b with the measured ones of Fig. 9a, it appears that the optimization framework ~~predicts is capable of estimating~~ lifetime consumption with reasonable accuracy. ~~The remaining useful life is only slightly underpredicted compared to the validation case study, with only slight underpredictions.~~ The minor ~~prediction mismatches across different wind direction sectors Γ , as~~ mismatches shown in Fig. 7 and Fig. 8, ~~have have only a~~ minimal impact on the overall lifetime consumption and remaining useful life.

4.1.4 Lifetime economic performance

~~Lifetime profit of the wind farm for all the WFC strategies for case study I. The economic profit is calculated as the difference between accrued revenue and the sum of incurred cost and lost revenue. The solid outline of the vertical bars in panel (a) denotes the measured value based on FAST.Farm, and the dashed outline in panel (b) denotes the corresponding predicted value from the optimization framework.~~

Figure 10 illustrates the economic performance of the wind farm over its lifetime. ~~Fig. Figure~~ 10a shows the economic profit ~~obtained by applying the optimal yaw offsets in the validation case study~~ calculated on the FAST.Farm simulations, while Fig. 10b displays the expected economic profit ~~as~~ predicted by the optimization framework using the engineering model FLORIS. In both figures, the vertical ~~bar plots represent the economic profit. This is calculated as~~ bars represent the time-discounted economic profit, i.e. the difference between the economic revenue from wind farm yield and the sum of economic costs from O&M and ~~the loss of~~ lost lifetime. The ~~x -axis shows the different WFC strategies and the y -axis represents the relative economic profit compared to the baseline. The numbers written in white on the bars indicate the relative values over~~ the economic profit is expressed in relative terms with respect to the greedy baseline over the 20-year operational duration. ~~Additionally, operational duration. The height of the bars (i.e., the relative profit) is also reported in white on each bar to ease comparison.~~




Figure 10. Time-discounted lifetime profit for all WFC strategies for case study I. The solid outline of the vertical bars in panel (a) denotes measured values based on FAST.Farm, while the dashed outline in panel (b) denotes the corresponding values predicted by the optimization framework. The green shaded extensions denote the additional economic gain generated when the farm operates beyond the goal EOL and until its lifetime is finally exhausted.

As discussed in Sect. 3.4.3, while the control actions were computed without considering time discounting, the economic performance is quantified taking time discounting into account. The present value of future profits is actualized using an annual discount rate of 8%, which is a typical value for wind energy projects (?).

1035 Additionally, in Fig. 10 some strategies include stacked green bars, which represent the extra economic gain due to lifetime extension. The corresponding For these strategies, the numbers written in green above the bars denote the additional profit assuming relative profit that is gained if the wind farm continues to operate beyond its reference EOL goal EOL and until its lifetime is finally exhausted.

Figure 10a shows that all WFC strategies result in higher economic profit compared to the baseline formulation. The magnitudes of these profits-greedy formulation. Profits are notably higher for the damage-aware strategies. This demonstrates indicates that incorporating damage considerations into the WFC optimization problem into the design of wind farm controllers – whether accurately or through approximations – has a significant impact on lifetime-economic profit.

Both the Max-power and Load-constrained strategies experience a loss of lifetime, leading to a reduction in which in turn leads to a loss of revenue. The Max-power strategy suffers a larger loss due to its shorter operational life. The Load-constrained strategy, on the other hand, incurred smaller economic-incurs smaller costs due to a slightly lower damage accumulation. As 1045 a result, the Max-power strategy only achieves a 1.7% increase in profit, while the Load-constrained strategy sees a larger increase of 7.7%increase.

The Load-balance strategy, due to its optimal tuning of weights in the thanks to a careful tuning of the weights in its optimization objective, results in an additional 7% profit over the Load-constrained strategy. The Lifetime-revenue strategy

1050 further optimizes loading ~~over~~, by simultaneously considering all various operating conditions ~~simultaneously, maximizing revenue and leading~~, which leads to a profit increase of 16.5% over the baseline.

Finally, the proposed *Max-profit* formulation ensures ~~the wind farm meets the~~ meeting the desired lifetime while maximizing profit. Although it generates less revenue than the *Lifetime-revenue* strategy, it saves significantly on maintenance expenses due to a reduced accumulated damage. This results in an additional 4% profit compared to the *Lifetime-revenue* strategy.

1055 The importance of optimizing profit, rather than solely focusing on revenue ~~maximization~~, is further ~~emphasized~~ stressed by the additional ~~profit earnings that can be~~ gained from the ~~extension of the~~ extended operational life of the wind farm. The *Max-profit* formulation achieves a 30% profit gain over the baseline, and an additional 13% gain compared to the *Lifetime-revenue* strategy.

Similarly to ~~the prediction of~~ lifetime consumption, ~~the expected lifetime profit shown in Fig. 10b even~~ lifetime profit is quite accurately predicted by the optimization framework ~~, as shown in Fig. 10b~~. The dashed outline ~~in of~~ the vertical bars in Fig. 10b represents ~~the expected value~~ estimated values, while the solid outline ~~in of~~ the bars in Fig. 10a corresponds to ~~the~~ measured values. When comparing the two figures, it ~~is evident~~ appears that there is a slight overprediction of lifetime profit for all WFC strategies compared to the respective baseline. This ~~overprediction effect~~ can be attributed to the ~~increased reduction in costs achieved by the optimization framework, which results from the~~ previously observed overprediction of component
1065 loads.

4.1.5 ~~Uncertainty quantification~~ Sensitivity to changes in the economic model parameters

~~Quantification of the uncertainties on lifetime economic performance caused by uncertainties in the economic model parameters of Sect. 3.4. Results of standard Monte Carlo analysis for a uniformly distributed economic parameter space for the proposed Max-profit strategy, where the numbers on the bars in white denote the mean value and the numbers in black on the dashed boxes~~
1070 ~~denote the corresponding uncertainty bounds (a). Results of variance-based parametric sensitivity analysis for the proposed Max-profit strategy, where the vertical bars represent the SOBOL index for different categories (b).~~

~~Due to the uncertainties associated with lifetime fatigue damage estimation, their impact on O&M costs, and the regulations governing wind farm market participation, the~~ The economic lifetime-aware WFC optimization problem ~~(refer to Sect. 2 for details)~~ is inherently stochastic is affected by significant aleatoric and epistemic uncertainties. While the underlying models
1075 ~~in the optimization framework use deterministic inputs and outputs, the resulting lifetime profit from power generation and damage accumulation can vary significantly. This variation may lead to different optimal yaw offset schedules. because of~~ these uncertainties the response that the computed control actions generate in the wind farm and their performance indicators are expected to change to some unknown degree. Clearly all aspects of the formulation are affected by uncertainties. However, the flow and load models – together with their uncertainties – used here are also widely employed in the literature to treat
1080 similar problems. The situation is different for the economic model. For example, spot market prices can vary in a very significant manner in time and space. Additionally, the cost model adopted here is highly simplified and might not fully reflect the complexities of the economics of wind farm O&M. Especially the effects of loads on cost is a problem that is not sufficiently documented in the open literature. Because of these reasons, here we try to quantify the effects that changes in the

parameters of the cost model might have on the results. While this analysis is far from being a comprehensive quantification of all uncertainties, it is hoped that it contributes to shedding some light on the robustness of the conclusions so far with respect to the economic modeling aspects of the problem.

Figure 11a illustrates the uncertainty-variations in lifetime profit resulting from economic-model-uncertainty-changes in the parameters of the economic model (cf. Sect. 3.4). This-uncertainty-is-represented-as-a-vertical-bar-plot. The y -axis shows the relative lifetime profit compared to the baseline greedy case, while the x -axis denotes the different WFC strategies. For each vertical bar, the numbers from top to bottom represent the upperbound, average, and lower bound of the lifetime-profit-uncertaintyprofit values, respectively.

The-uncertainty-is-




Figure 11. Quantification of the uncertainties on lifetime economic performance caused by uncertainties in the economic model parameters of Sect. 3.4. Results of a Monte Carlo analysis for a uniformly distributed economic parameter space, where the black numbers denote the mean value and the grey numbers the corresponding uncertainty bounds (a). Results of a variance-based parametric sensitivity analysis for the *Max-profit* strategy, where the vertical bars represent SOBOL indices for different cost items (b).

Uncertainties are calculated by varying all the economic model parameters, assuming a uniform distribution. These parameters include thirteen economic model factors across two components. Four parameters are related to the mean time to maintenance (both scheduled and unscheduled), four are related to downtime for repair and maintenance, four pertain to personnel/labor costs for maintenance, and one is associated with the electricity market price. Uncertainty uniform distributions for each. On the revenue side, the varied parameter is the electricity market price $\mathbb{R}_{avg}$. On the cost side, for each component the varied parameters include all six maintenance factors appearing in Eq. (18). Since we consider both the blade and tower components,

a total of 13 parameters is varied in the uncertainty analysis. For all parameters, we consider a variation of 8% around its nominal value listed in Tab. 1.

The uncertainty quantification is performed using Monte Carlo analysis. In this approach, numerous evaluations are conducted by randomly sampling, where the parameter space and evaluating the merit function until convergence is achieved is randomly sampled until the mean as well as the 25th and 75th percentiles of profit converge.

The results in Fig. 11a show that, for all WFC strategies, a maximum variation of 8% in the parameters leads to uncertainties varying only between 3% to 4%. This indicates that, for all control strategies, the lifetime profit is only modestly affected by the affected only to a modest degree by uncertainties in the economic parameters across all control strategies. The same analysis was repeated for variations between 4% and 12% in increments of 2%, observing a nearly linear behavior of the upper and lower bounds of profit.

Additionally To better understand the sources of uncertainty in lifetime profit, a variance-based parameter-sensitivity analysis is conducted for the *Max-profit* strategy, where the thirteen economic model parameters are varied to better understand the sources of uncertainty in the lifetime profit by varying the thirteen parameters of the economic model. The results of the sensitivity analysis are presented in terms of first-order sensitivity coefficients, also known as first-order SOBOL indices (?). The sensitivity coefficient S_{X_k} indicates how much the output variance $\text{Var}(Y)$ can be reduced, on average, if the uncertainty associated with the input variable X_k is fully known, where the output is some known function $f(\cdot)$ of the inputs, i.e. $Y = f(X_1, X_2, \dots, X_k, \dots)$.

Figure 11b displays the sensitivity analysis results. The economic model parameters The resulting SOBOL indices are grouped into four qualitative categories, which are shown along the x -axis. The number in parentheses next to each category indicates the number and are reported in Fig. 11b. The number of parameters in that group. The vertical bars represent the SOBOL index for each category, and the y -axis shows the corresponding values.

From the figure, it is clear each category is indicated in parentheses. Results indicate that the uncertainty in electricity market prices contributes most significantly is the most significant contributor to the overall uncertainty. This is because market prices directly affect both the accrued revenue and the lost revenue due to downtime. The second largest source of uncertainty is the mean time to maintenance, which represents the time to failure failure, which directly impacts the frequency of maintenance activities and the associated costs. While the remaining two categories contribute relatively less, their absolute impact is still noteworthy. This further emphasizes the importance of quantifying the uncertainties associated with failure rates and other O&M parameters.

4.2 Case study II: Realistic wind farm layout and wind rose

The goal of the second case study is to apply the proposed WFC approach to a more complex and realistic scenario. In this case, the wind farm is a small offshore cluster operating under realistic field-measured wind conditions. The primary objective is to evaluate how the new method performs for a larger cluster of turbines, focusing on how the increased turbine count and varied inflow conditions affect the solution. Additionally, this case study aims to assess the growth of the computational cost of the proposed method as the wind farm cluster size increases.

For comparison, all reference control formulations presented in Sect. 2.7 are also considered here. The results shown in the following are those obtained by the optimization framework. Since the previous [results section](#) showed that the [optimization framework](#) – notwithstanding its underlying low-fidelity flow model – provides accurate predictions of lifetime quantities, additional validation using FAST.Farm was deemed not [necessary to be necessary in this case](#).

4.2.1 [Layout Farm layout and input ambient wind conditions](#)

[Layout and wind conditions for the realistic case study. Wind rose \(a\). Wind farm layout and wakes at hub height for wind blowing from the southwestern direction \(b\).](#)

This case considers the nine-turbine wind farm layout defined in the CL-Windcon project [\(?\)](#). The turbines are arranged in a regular grid, with a spacing of 7 D in the [x direction](#) and 5 D in the [y direction](#), as shown in Fig. 12b. The wind conditions are based on reanalysis data obtained from the site of the Hollandse Kust Nord offshore wind farm [\(?\)](#), which is also being [assessed as part of](#) [considered in](#) the SUDOCO project [\(?\)](#). The corresponding wind rose is shown in Fig. 12a. The wind speed range is limited between 6 and 10 ms^{-1} . It is assumed that these [wind ambient](#) conditions remain constant over the [course of the year](#) [various years of operation](#).




Figure 12. [Layout and wind rose for the realistic case study. Wind rose \(a\). Wind farm layout and wakes at hub height for wind blowing from the southwestern direction \(b\).](#)

Since the wind inflow comes from all directions, all turbines operate in a combination of waked and free-stream conditions. The [baseline](#)-AEP loss due to wake interaction [in greedy mode](#), as calculated by [FLORIS](#) [Floris](#), is approximately 5%. Similarly to case study I, the market price is assumed to vary around an annual average, depending on the ambient wind speed. An annual [inflation discount](#) rate of 8% is [applied for the economic evaluation](#) [used to estimate the present value of future profit](#)

1150 (?), as discussed in Sect. 3.4.3. For simplicity, the number of maintenance activities, fixed maintenance costs, and maintenance downtime per component are assumed to remain constant over time. ~~Additionally, the lifetime consumption is calculated based on damage accumulation at the~~ Here again, loads, DELs, and damage are computed at tower base and blade root.

This case study assumes a 20-year lifetime. During the first 10 years, the farm operates in greedy control mode, before switching to one of the ~~tested-control~~ WFC formulations for the second half of ~~the-its~~ lifetime. As in case study I, one turbine (T5) is assumed to have accumulated more fatigue damage than the others. T5 operates in waked conditions most of the time, and also ~~wakes-affects with its wake~~ the turbines at the trailing edge of the farm. As a result, it is assumed that after 10 years turbine T5 will have accumulated 75% more damage than expected. This additional damage is ~~greater-larger~~ than in case study I, as the more realistic wind conditions in this study are less challenging than ~~focusing-on-a-the~~ single wind sector considered in that case. This is clearly an idealized exemplary situation. In another scenario, T5 might represent a machine that suffers from higher failure rates than the others, for example because of ~~malfunctioning-manufacturing~~ issues, extra vibrations suffered during ~~erectionsits erection phase~~, or other possible reasons.

4.2.2 Lifetime consumption

~~Lifetime consumption of the wind farm for all the WFC strategies for case study II-~~

Figure 13 shows the lifetime consumption for the wind farm ~~across-for~~ all WFC strategies. Similarly to Fig. ~~??9~~, the horizontal bars represent the operational years of the wind farm, with a ~~reference-goal~~ EOL of 20 years. As previously noted, the various WFC strategies are deployed after the first 10 years of operation.

~~The-~~




Figure 13. Lifetime consumption of the wind farm for case study II for all considered WFC strategies.

The figure adopts the same layout as used in Sect. 4.1.3. The green-shaded region indicates ~~the-extended-lifetime-years beyond the-reference~~ an achieved extended lifetime beyond the desired goal EOL, while the red-shaded region represents ~~the~~

1170 loss of operational years due to lifetime consumption before the reference EOL. For a more detailed interpretation of the layout of the figure, please refer to Sect. 4.1.3. lost operational time due to having exhausted the lifetime before the desired EOL.

Figure 14 shows that ~~, compared to the baseline, all other all~~ WFC strategies result in ~~a larger number of operational years~~ longer operational times than the baseline greedy mode. Strategies that account for component damage ~~over the lifetime~~ – either explicitly or indirectly – ~~also lead to increased operational years compared to the baseline and~~ perform better from
1175 this point of view than the Max-power strategies ~~strategy, which does not account for loads in its formulation~~. However, only the lifetime-aware strategies, specifically *Lifetime-revenue* and *Max-profit*, ensure operation throughout the entire ~~reference~~ goal EOL. This is because the optimal yaw ~~offset~~ offsets for these two strategies ~~is~~ are less aggressive, as ~~they the controllers~~ consider the inflow and operational conditions at all turbines over their entire lifetime. Furthermore, the proposed *Max-profit* formulation provides a slight lifetime extension, allowing the wind farm to operate beyond the ~~reference goal~~ EOL, thus
1180 generating additional revenue.

4.2.3 Lifetime economic performance

~~Lifetime profit of the wind farm for all the WFC strategies for case study II. The economic profit is calculated as the difference between accrued revenue and the sum of incurred cost and lost revenue.~~

Figure 14 ~~presents the reports the time-discounted~~ lifetime economic performance of the wind farm. ~~The results~~ Results are
1185 displayed in the same format ~~as in used for~~ Fig. 10b. The vertical bars show the ~~lifetime~~ profit relative to the baseline profit ~~for all the formulated WFC strategies. The numbers written in white on the bars represent obtained in greedy mode over the 20-year operational duration. The white numbers represent the heights of the bars (i.e., the relative profit~~ values over the 20-year operational duration), to ease comparison among the various WFC strategies.




Figure 14. Time-discounted lifetime profit of the wind farm for case study II for all WFC strategies.

All strategies result in economic gains ~~over the baseline for the considered with respect to the baseline greedy~~ scenario.
1190 However, due to ~~lifetime loss~~ loss of lifetime, the *Max-power*, *Load-constrained*, and *Load-balance* strategies show relatively minor gains, which depend on ~~the corresponding revenue loss~~ their specific revenue losses. In contrast, the two lifetime-aware

strategies – *Lifetime-revenue* and *Max-profit* – yield significantly higher gains of 7.3% and 8.6%, respectively, thanks to their guaranteed operation over the full desired lifetime.

1195 ~~The difference in gains between these two strategies highlights the importance of balancing economic costs and revenue within the optimization objective of the proposed~~ Of these two, the best performer is the *Max-profit* ~~formulation. Additionally, the profit gain magnitude is influenced by wake interactions across the farm over its entire lifetime, because of its ability to include the effects of loads on cost, and hence on profit.~~

4.3 Computational effort

The proposed ~~methodology approach~~ requires a setpoint optimization for all turbines in all possible ambient wind conditions throughout their entire lifetime. The resulting one-shot optimization can become computationally demanding. To assess the growth of the computational cost, ~~the layout shown in Fig. 12b is expanded to include more turbines.~~

~~Starting with farms of varying number of turbines are investigated. Starting from~~ the 3×3 layout, the configuration is first reduced to a 2×2 setup and then extended to 4×4 and 5×5 configurations, all while maintaining the same x - y separation distances. The other problem parameters (wind rose, goal EOL, time of deployment of WFC, etc.) are ~~unvaried~~ unchanged with respect to the 3×3 case.

The optimization problem is solved iteratively until numerical convergence is achieved. All optimizations are performed on a desktop computer with an Intel(R) Core(TM) i7 processor, 8 GB of RAM, and a 64-bit operating system.

~~Increase of wall time with the number of turbines in a regular grid with spacings of 7 D and 5 D in the x - and y -directions, respectively.~~

1210 Figure 15 presents the results of the computational performance test in a log-log plot. The x ~~-axis-axis~~ represents the total number of turbines in the ~~assessed configuration~~ farm, while the y ~~-axis-shows the real~~ -axis shows the wall clock time taken by the optimization problem to converge. Time is normalized with respect to the 2×2 configuration, which required 0.24 hours to complete. This wall clock time is a value with little practical meaning, as hardware performance evolves quickly and no effort was made to optimize the code for speed. Therefore, the more useful piece of information provided by the figure is the slope of the computational cost growth.

1215 As expected, the proposed ~~methodology is computationally~~ method is expensive, with the computational cost increasing significantly as the wind farm size grows. ~~For instance, the optimization on a 5×5 layout takes approximately 100 hours. However, since the optimization is performed offline, the problem only needs to be solved once throughout the lifetime of the wind farm, either before or during the operational phase. As a result, the computational effort, relative to the operational lifetime of the wind farm, is negligible. Even if the wind farm operator decides to update the optimal yaw-offsets based on more accurate data, such as actual damage measurements or digital twins, the computational cost of reevaluating the optimization problem will still be minimal compared to the operational duration~~ This is unfortunately the price that has to be paid for performing a correct calculation of the cumulative effects of fatigue on damage. Clearly, this extra cost can be avoided by resorting to simpler formulations. However, as shown by the results presented earlier, without a correct calculation of damage, in general it is not possible to guarantee a desired duration of the operation of the farm.




Figure 15. Increase of wall-clock time with number of turbines in a regular grid with spacings of 7 D and 5 D in the x and y directions, respectively. Values are normalized with respect to the 2×2 configuration.

Regarding the high computational cost of the formulation, two additional observations are in order.

First, we made no effort at optimizing the efficiency of the software, which is only a proof-of-concept Matlab implementation that was run on a standard desktop computer. Clearly, very different software and hardware resources could be employed for a real industrial implementation. On the software side, we used Floris Vers. 3, while the more recent Vers. 4 (?) offers much improved performance through an extensive reorganization of the code and vectorization of the operations. Additionally, various phases of the calculations could be easily parallelized because they are independent one from the other, including all calculations of loads on the turbines and their components through the surrogate. Additional speedups can be obtained by discarding symmetric conditions, and using the “serial refine” approach available in Floris to enhance convergence by choosing promising starting guesses for the optimization (?). It is also important to note that no special efforts were made to optimize the code for computational efficiency, and only a standard desktop computer was used for the optimization which is not realistic for a large industrial project. unclear at this point what performance increase would actually be achievable by these various methods. However, it is also clear that there is potential for improving – possibly in a very significant manner – the results presented here.

Second, the optimization of the setpoints is performed offline, and therefore the computational cost is less of a concern than in the case of online calculations that need to be executed during operation. It is reasonable to assume that the optimization will be repeated multiple times throughout the lifetime of the farm, for example to update the state of damage of the turbines based on measurements or digital twins, or to progressively improve the wind farm flow model by recalibrating it with measured SCADA data (?). However, as the reoptimization will be probably performed only every several months, there will be ample time to complete the necessary calculations.

More importantly, the large increase in economic profit made possible by the proposed *Max-profit* strategy – compared to the state-of-the-art *Max-power* and other simple strategies (as discussed in Sect. 4.1.4 and Sect. 4.2.3) – justifies the additional computational effort is the only one that can guarantee meeting a desired goal EOL, as it is the only one that explicitly enforces that constraint. If the increased economic performance is significant enough, it could justify the extra computational burden.

5 Conclusions

1250 We have introduced a novel approach: the economic lifetime-aware WFC. The new method integrates the current damage state of turbines – either obtained through a digital twin or estimated from operational data – to derive control policies that optimize the long-term lifetime management of the wind farm. The optimization is driven by an economic value function that balances power generation (considering fluctuating spot market prices) with O&M costs, including potential revenue losses due to turbine downtime. The optimization is constrained to ensure that the target lifetime of selected turbine components is
1255 met while operating under varying inflow conditions. The optimization variables include the control setpoints of each turbine, ~~which can include axial induction, wake steering, mixing, or combinations thereof. In this sense, the proposed method is general, allowing end-users to select based on the availability or technical applicability of the different control approaches. The resulting optimal setpoints depend on the ambient wind conditions and~~. Once computed offline for a set of ambient climate conditions, the optimal setpoints are interpolated at run time based on the detected current conditions, and then dispatches to the turbines, similarly to what is done for open-loop power boosting WFC.

~~To implement our~~ Our new proposed method ~~, first, relies on~~ an engineering flow model ~~is adapted to estimate the for estimating the inflow and~~ power output of each individual ~~turbines within turbine in~~ a wind farm ~~for given wind climate conditions. Next, , and on~~ a load surrogate model ~~uses the estimated rotor inflow to predict the DELs for selected turbine components, which resulting DELs. These~~ are then accumulated into a representative damage value ~~. Finally, an for selected components of interest. An~~ economic model integrates energy production with the cost of component fatigue damage to estimate ~~the net profit, providing a comprehensive view assessment~~ of the economic performance of the wind farm.

The proposed lifetime-aware strategy, termed *Max-profit*, has been compared to five relevant reference formulations. The ~~first two, Greedy and Max-power , approaches~~ do not account for ~~loads. The following two, the effects of loads on fatigue damage. The Load-constrained and Load-balance , methods~~ incorporate loads by considering DELs either as a constraint or
1270 as an objective. The ~~fifth formulation, Lifetime-revenue , is an intermediate formulation that approach is a simplified version of the Max-profit strategy, as it~~ considers the lifetime damage constraint but does not account for load-induced costs. The performance of the proposed *Max-profit* formulation and the reference ones was evaluated ~~through two different case studies in two different scenarios~~ using wake steering. ~~The designed life of the wind farm is considered to be~~ For both cases, the operator would like the farm to reach a goal life of 20 years, and ~~the control strategies are deployed~~ deploys the various control strategies
1275 after 10 years of ~~baseline operation~~ operation in standard greedy mode.

The first case study examined a simple wind farm layout consisting of three wind turbines and a synthetic wind rose ~~. For validation, the mid-fidelity wind farm simulation tool FAST.Farm was used with multiple turbulent seed realizations. Results indicated that produces very frequent and persistent wake interactions. Results indicate~~ that non-load-aware strategies lead to increased damage accumulation, causing the wind farm to prematurely reach its EOL, in turn resulting in reduced revenue. In
1280 contrast, all load-aware strategies are able to decrease lifetime consumption, leading to lower costs, smaller revenue losses, and higher profits. However, only the lifetime-aware strategies, *Lifetime-revenue* and *Max-profit*, ensure full operation throughout

the desired lifetime duration. Additionally, ~~is some ease in some cases~~ we observed an even longer lifetime, ~~allowing for operation which allows the farm to operate~~ beyond the defined EOL, thereby generating additional profit.

1285 ~~Results also show accurate predictions of lifetime performance using the developed optimization framework~~One additional goal of the first simplified problem was to produce a quantification of the accuracy of the computational framework. To this end, predictions generated by the framework low-fidelity and steady-state engineering models are compared to measurements obtained on the mid-fidelity transient simulation tool FAST.Farm, assumed as a ground truth. Results show in general a good agreement between predictions and measurements, thanks also to specific tuning of the models by the “farm as a sensor” approach. Modest mismatches in the prediction of DELs appear to have only a limited impact on the quality of lifetime performance indicators. Furthermore, uncertainties in lifetime profit due to variations in the economic model parameters were assessed ~~using first with a~~ Monte Carlo analysis ~~, and and then with a~~ variance-based parametric sensitivity ~~analysis with study using~~ first-order SOBOL indices. Results indicate that even with an 8% variation in all economic model parameters, lifetime profit uncertainties remain within 3%, while market price uncertainty has a more significant impact.

The second case study expanded the analysis to a larger wind farm cluster of nine ~~turbines, typical of an onshore wind~~ farm, ~~onshore turbines,~~ arranged in a regular grid layout. This study ~~also incorporated a real~~ was based on a realistic wind rose based on measured data. ~~The results~~ Results highlight the clear advantages of explicitly incorporating lifetime constraints ~~into in~~ the WFC optimization problem. Specifically, only the *Lifetime-revenue* and *Max-profit* formulations were able to ensure operation over the entire desired lifetime. As damage accumulates ~~over time~~, the control actions taken ~~today — through the loading on the turbines — in one operating condition~~ directly impact the available ~~capacity for future control decisions~~. Since ~~the other strategies~~ room left to control decisions taken in other conditions, something that cannot be accounted for by control synthesis procedures that consider one ambient condition at a time. In fact, strategies that do not explicitly account for the interplay ~~between present and future and couplings among~~ control actions, ~~they lead to are shown here to generate~~ higher fatigue accumulation and, ultimately, ~~to a reduction in~~ possibly encounter a reduced operational life. In contrast, the proposed *Max-profit* strategy results in approximately 8.6% higher profit compared to the baseline *greedy policy*, with 1.3% of that increase ~~attributed solely to the consideration of~~ being due to having considered maintenance expenses in the optimization process.

1310 ~~Additionally, we~~ We have also assessed the computational burden of solving the proposed optimization problem, which performs a one-shot optimization over the defined lifetime for all inflow conditions ~~;-This simultaneously. This analysis~~ was carried out on a standard desktop computer, rather than using high-performance computing resources, without any attempt at optimizing ~~the software~~ for computational efficiency. ~~For a wind farm with 25 turbines arranged in a regular grid layout, the optimization process took over 100 hours to converge. The computational load~~As expected, accounting for couplings among control actions – which is key for ensuring a desired goal life – comes at the cost of a higher computational load. This cost increases significantly as the size and complexity of the wind farm grows. ~~However, since the optimization is performed offline once and then rarely (if ever) repeated, the computational burden is minimal in the context of the typical operational lifespan of a wind farm, which spans several years.-~~

The proposed formulation can be utilized as a decision support tool during the pre-deployment phase to evaluate the economic feasibility of various wind farm operational scenarios. For instance, an operator could compare the trade-offs between a shorter plant lifetime driven by higher power capture (which also results in greater loading and increased O&M costs) followed by early decommissioning, versus extending the lifetime of the plant beyond its original design by adopting a more conservative power capture strategy and reducing turbine loading. Once an optimal life management strategy has been selected, the wind farm operator can deploy the lifetime-aware controller to maximize profit while ensuring that the specified lifespan targets for turbine components are met. How much this cost can be actually decreased by improving the implementation remains at this point an open question. We have listed a number of possible changes to the software that might lead to substantial improvements, but only a more mature and complete implementation will tell the actual speedups that can be achieved. However, we believe that the high computational costs could be amply justified by the significant improvements in the economic performance of the farm that appear to be possible from the promising results presented here.

Appendix A: Surrogate load model

A1 Inflow characterization by RAIQs

The three panels of Fig. A1 report an illustrative example of RAIQs computed on the downstream turbine shown in Fig. 2. As the wind direction sweeps in the range $250^\circ < \Gamma < 290^\circ$, the wake of the upstream turbine makes a full pass over the downstream rotor. The ambient wind speed and TI are held constant at 8 ms^{-1} and 12%, respectively.




Figure A1. Variation of the RAWS (a), RATI (b), and RAHS (c) on the waked turbine of Fig. 2 for varying wind direction Γ , for an ambient wind speed of 8 ms^{-1} and TI of 12%. For a wind direction of 270° , the array of turbines is aligned with the incoming wind.

Figure A1 clearly shows that the RAWS is at its minimum for the full-wake situation ($\Gamma = 270^\circ$) and progressively recovers toward the ambient wind speed as the wind direction moves away from this value. In contrast, the RATI is largest in full-wake condition due to wake-added turbulence. As the turbine experiences less wake impingement, the RATI decreases toward the

1335 ambient TI value. Finally, the RAHS shows the expected asymmetry, characteristic of left and right partial-waking. For both the full-wake ($\Gamma = 270^\circ$) and no-wake ($\Gamma < 255^\circ$ or $\Gamma > 285^\circ$) scenarios, the RAHS is close to zero, as there is no significant wind speed asymmetry across the lateral direction of the rotor disk. However, during partial-waking situations, the RAHS is positive for $255^\circ < \Gamma < 270^\circ$ and negative for $270^\circ < \Gamma < 285^\circ$. These trends in RAWS, RATI, and RAHS further contribute to the understanding of the results presented in Sect. 4.1.2.

1340 A2 Effect of calibration on surrogate performance

Figure A2 illustrates the performance of the load surrogate model when the input RAIQs are estimated from FAST.Farm for both the training and testing datasets. The root-mean-square percentage error (RMSPE) is used as performance metric. The scatter plots for both the training and testing datasets are similar, indicating that the ANN training data pool was sufficiently large to ensure a correct generalization.




Figure A2. Performance of the load surrogate model for the projected blade root DELs (a) and projected tower base DELs (b). DELs measured on FAST.Farm (the ground truth) are on the x axis, while DELs predicted by the surrogate on the y axis. Root-mean-square percentage errors (RMSPE) are used as performance metrics. Half of the dataset was used for the training of the ANN (in blue), while the other half was used for testing the trained ANN (in orange).

1345 As discussed in Sect. 3.1, load estimates based on RAIQs from Floris in its baseline configuration were found to be not satisfactory, because of the unavoidable mismatch between its steady-state inflows and the ones generated by the dynamic model of FAST.Farm. To address this problem, the wake-added turbulence model in Floris was replaced with the ? model, and the “wind farm as a sensor” method was used for calibrating all Floris model parameters from the FAST.Farm dataset (see Sect. 3.2).

To illustrate the improvements achieved by model tuning, Fig. A3 compares the DELs estimated by the load surrogate when the inflow RAIQs are computed from the non-calibrated Floris (shown as black dots) with the ones obtained with the calibrated Floris model (shown as red dots). The results show considerable improvements. In particular, panel (b) highlights the beneficial effects of flow model tuning, whereas the baseline model significantly underpredicts tower base loads. This is primarily due to the crude estimation of RATI in the baseline model, which uses the model of ?.




Figure A3. Comparison of DELs predicted by the load surrogate model fed by RAIQs measured on the non-calibrated Floris (in black) and on the calibrated version (in red). DELs measured on FAST.Farm (the ground truth) are on the x axis, while surrogate-predicted DELs on the y axis. Data points include both the training and testing datasets, for $I = 12\%$. (a) Projected blade root DELs; (b) projected tower base DELs.

1355 **Appendix B: Nomenclature**

	c	Turbine component
	C	Number of turbine components
	$\mathbb{C}$	Cost item
	D	Damage
1360	D_{ref}	Damage threshold
	$\tilde{\mathbb{D}}$	$\sim$ Discount rate
	f_a	Wind condition frequency
	I	Turbulence intensity
	k_*	Flow model parameter
1365	$\mathbb{K}_T \mathbb{K}$	Damage safety factor
	L	Load cycle amplitude
	L_u	Ultimate load
	m	Wöhler exponent
	n	Number of load cycles
1370	N	Number of turbines
	N_{eq}	Equivalent number of load cycles
	P	Turbine power
	P^r	Turbine rated power
	r	Rotor radial coordinate
1375	$\mathfrak{r}$	Polynomial fit parameter
	R	Rotor radius
	$\mathbb{R}_T \mathbb{R}$	Electricity market price
	$\tilde{t}$	$\sim$ Time instance
	T	~~Turbine~~ Desired lifetime horizon
1380	U	Wind speed
	$\mathfrak{u}$	Control ~~set-points~~ setpoints
	$\mathbb{W}_a$	Set of ambient wind conditions
	$\mathbb{W}$	Weight parameter for <i>Load balance</i> strategy
	$\tilde{y}$	$\sim$ Year
1385	$\tilde{Y}$	$\sim$ Number of years in T
	α	~~Induction factor~~
	γ	~~Yaw~~ Misalignment angle

	Γ	Wind direction
1390	Δt	Time duration
	ψ	Rotor azimuthal coordinate
	AEP	Annual energy production
	ANN	Artificial neural network
1395	CF	Capacity factor
	DEL	Damage equivalent load
	EOL	End of life
	~~FLORIS~~ Floris	FLow Redirection and Induction in Steady State
	<u>NPV</u>	<u>Net present value</u>
1400	O&M	Operation and maintenance
	RAWS	Rotor-average wind speed
	RATI	Rotor-average turbulence intensity
	RAHS	Rotor-average horizontal shear
	RAIQ	Rotor-average inflow quantity
1405	RMSPE	Root mean square percentage error
	TI	Turbulence intensity
	WFC	Wind farm control

Data availability. The content and data of figures 3-5, 7-13, and 15-17 can be retrieved in Python pickle format via the DOI¹: <https://doi.org/10.5281/zenodo.15118524>

1410 *Author contributions.* CLB developed the concept of economic lifetime-aware WFC and supervised the overall research. AA and RB expanded the concept and performed literature research. AA and RB implemented the proposed method using the models discussed in the optimization framework section. AG developed the load surrogate and supported AA and RB with its integration into the framework. AA, RB, and AG carried out the validation studies, where AG performed the FAST.Farm simulations. All authors contributed to the interpretation of the results. AA, RB and CLB wrote the manuscript, with contributions from AG in the sections on load surrogate, results and discussion.

1415 All authors provided valuable input to this research work through discussions, feedback, and improvement of the manuscript.

Competing interests. The authors declare that they have no conflict of interest, except for CLB who is the Editor-in-Chief of the Wind Energy Science journal.

¹Review note: This is a preliminary repository, the final DOI will be generated if the paper is accepted

Financial support. This work has been supported by the SUDOCO and TWAIN projects, which receive funding from the European Union's Horizon Europe Programme under the grant agreements No. 101122256 and 101122194, respectively.

1420 *Acknowledgements.* The authors would like to thank S. Kainz for the valuable support with the uncertainty quantification study.