First experimental results on lifetime-aware wind farm control

R Braunbehrens, A Anand, F Campagnolo, C L Bottasso

Wind Energy Institute, Technische Universität München, Garching bei München, Germany

E-mail: robert.braunbehrens@tum.de, abhinav.anand@tum.de, filippo.campagnolo@tum.de, carlo.bottasso@tum.de

Abstract. The study presents the methodology, implementation, and first results of the experimental test of a lifetime-aware wind farm control (WFC) approach. The novel strategy goes beyond the conventional WFC formulations and aims to maximise profit while ensuring fulfilment of a desired lifetime. Damage due to fatigue is formulated using a cost model and is also limited by a lifetime constraint. The control strategy is based on yaw-induced wake steering and utilises a tool chain of power and damage estimators. The experiment was conducted with three model turbines in a boundary layer wind tunnel, which allowed the simulation of a realistic dynamic variation of the wind direction. The novel strategy is compared to conventional greedy operation, as well as to the classical maximum power strategy. The results show that the novel control strategy can improve the economic performance of the farm.

1. Introduction

Current wind farm control strategies focus mainly on increasing power generation [1]. However, fatigue damage leads to a reduction in wind turbine operational life and an increase in maintenance expenses. Together with yield, it therefore plays a crucial role in determining the economic profitability of a wind farm. Proposed wind farm control methods, which account for fatigue, typically try to limit the damage equivalent loads (DELs), a policy that makes it difficult to understand the effects of such control choices on lifetime [2]. In fact, the lifetime optimisation of the operation of a wind farm cannot only rest on decisions based on present conditions or predictions over short time horizons. Instead, as damage accrues over time, one needs to consider the couplings between present and future actions. The control actions taken today directly affect the room left for control actions that can be taken tomorrow.

This work presents the first experimental testing of a new form of wind farm control, termed lifetime-aware wind farm control. The control strategy aims at improving the economic performance of the wind farm through wake steering. The experiment is conducted with three model turbines, installed on a turntable in the test section of a boundary layer wind tunnel [3]. The turntable allows for the simulation of a dynamically changing wind direction [4].

The paper first introduces the novel formulation of the control problem. It then reports the underlying tool chain followed by the setup of the wind tunnel experiment. Lastly, the evaluation of the first results is presented.

2. Optimisation problem

Differently to conventional power maximisation, lifetime effects are included in the derivation of the optimisation problem through two main aspects. Firstly, the formulation considers the effect of fatigue damage on the economic performance of a project. The strategy finds the optimal compromise between aggressive yield-increasing control actions and protective load-mitigating actions that limit maintenance costs. Secondly, the turbine should still be guaranteed to reach a defined lifetime. This is motivated by the assumption that a premature plant failure leads to lost production, and consequently, a loss of revenue over for the missing years. Optimality is accordingly defined by a value function that includes 1) the effects of more or less aggressive control on revenue through an electricity price, and 2) the negative effect of O&M costs, including loss of revenue due to downtime. The optimisation is further subjected to constraints that enforce a desired target lifetime duration for selected turbine components. The controller is implemented in open-loop configuration, as commonly used in power-maximising control strategies [1]. This has the advantage that the costly optimisation problem can be solved offline, through steady-state simulations. The controller then actuates turbines based on real-time measured environmental inflow conditions, as well as the turbine operating state, on the basis of pre-computed look-up tables.

The optimisation problem can be written as

$$\max_{\boldsymbol{\gamma}} \sum_{T} \text{Revenue}(\boldsymbol{\gamma}, \mathbf{W}_a) - \text{Cost}(\boldsymbol{\gamma}, \mathbf{W}_a), \tag{1}$$

subject to

$$D_i^c(\boldsymbol{\gamma}, \mathbf{W}_a) \leq D_{ref}^c \quad \forall i \in N \text{ and } \quad \forall c \in C, \tag{2}$$

for a given set of ambient wind conditions $\mathbf{W}_a$ over a chosen lifetime T, and yaw-offsets $\boldsymbol{\gamma} = [\gamma_1, ..., \gamma_i, ...\gamma_N]$ for N turbines in the farm and components C in each turbine. The vector $\mathbf{W}_a$ should contain all wind conditions, e.g. a time series of wind direction, wind speeds etc., that occur during the turbine lifetime. For sufficiently long periods, these can be approximated through a representative frequency distribution.

The optimiser aims to balance the accrued revenue through power generation and the incurred cost due to maintenance activities. The optimisation problem is subjected to the lifetime constraint (2) which ensures that the resulting damage D does not exceed the limit damage D_{ref}. This parameter is a desired damage limit, which includes the necessary safety factors. Clearly, various choices of the limit are possible. For example, D_{ref} could be set to the design limit. Alternatively, one could leave a remaining damage margin to ensure a second life to a component. Notwithstanding the fact that various scenarios are possible, the concept behind this approach is to ensure the control of damage accumulation in order to achieve a desired goal. This would not be possible without an explicit constraint.

In the present study

$$D_{ref}^c = \max_i \sum_{T} \; D_i^c(\mathbf{W}_a), \tag{3}$$

is chosen as the maximum of the accumulated damage over all the turbines that occurs under a conventional greedy control strategy for the component c.

3. Model framework

Figure 1 provides an overview of the model tool chain used to derive the lifetime-aware control strategy. The two branches represent the estimation of the two main quantities, the power production, and the component fatigue damage. The individual tool chain elements can be summarised in three categories:

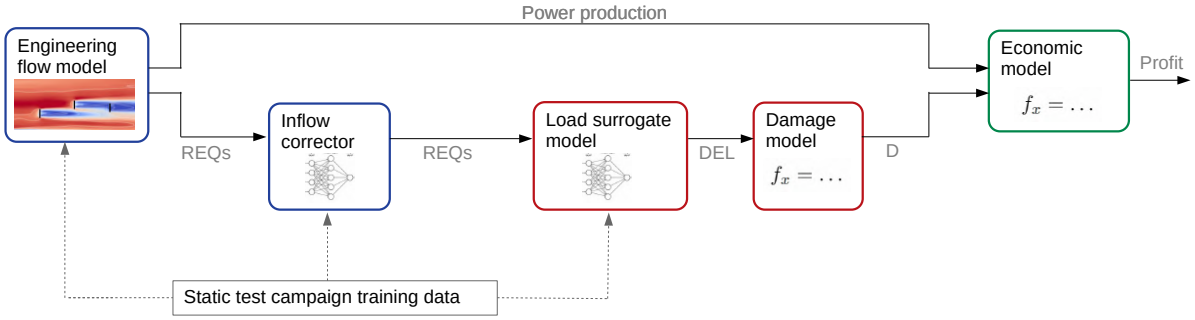

Figure 1: Tool chain used to synthesise the lifetime-aware control strategy: engineering flow model and inflow corrector (blue blocks, section 3.1); fatigue damage estimator (red blocks, section 3.2); and economic model (green block, section 3.3).

- An engineering flow model that serves two purposes: to routinely predict the power production, and to estimate the Rotor-Equivalent inflow Quantities (REQs) for each turbine within the cluster (blue blocks within Fig. 1). REQs consist of: Rotor-Effective Wind Speed (REWS), vertical wind shear (modelled as a power-law profile with exponent α), horizontal shear (assumed to be linear with coefficient κ), and Rotor-Effective Turbulence Intensity (RETI). In order to improve the predictions of the flow model, a data-driven inflow corrector block is used.
- Estimates of the REQs are used to calculate the fatigue damage on the selected components by a load surrogate model, coupled to a damage model (red blocks).
- Power and damage estimates are the inputs to an economic model (green block), which estimates the wind farm profit.

Training data from a previous campaign is used to tune the engineering flow model, inflow corrector, and load surrogate model to the specific wind tunnel case. This previous campaign was conducted with the same experimental setup, however, under steady wind conditions. It featured two different boundary layers: Mod-TI and High-TI, with $\alpha = [0.14, 0.21]$, and RETI $= [0.03, 0.07]$, respectively. The turntable angle ϕ was kept fixed during each run, with its value within the range $\pm 15°$. For each turntable position, the yaw angles of the two upstream turbines were yawed up to $25°$, in settings that steered the wake away from the impinged turbines. The REQs were measured at each turbine with an observer similar to the one described later in the experimental setup (section 4.1). Overall, the training dataset consists of 203 observations. The following sub-sections describe the flow, fatigue damage, and economic models.

3.1. Flow model

In this work, the engineering flow model is FLORIS [5]. The inflow to the wind tunnel test section is not uniform, in the sense that the wind speed shows a lateral gradient [6]. To reflect this inhomogeneity and improve the predictive capabilities, FLORIS was adapted following the "wind farm as sensor" methodology [7]. The model adaption consists of a representation of the background flow as well as a tuning of the wake sub-models. For the former task, following the formulation in [6] (albeit implementing the corrections here for FLORIS V3.2), four speed-up nodes, placed with lateral separation of 1 m, are used to represent the lateral gradient. The background wind speed at the turbine positions is found by interpolating between the nodes. The correction is constant along the streamwise direction. However, it was observed, that the gradient

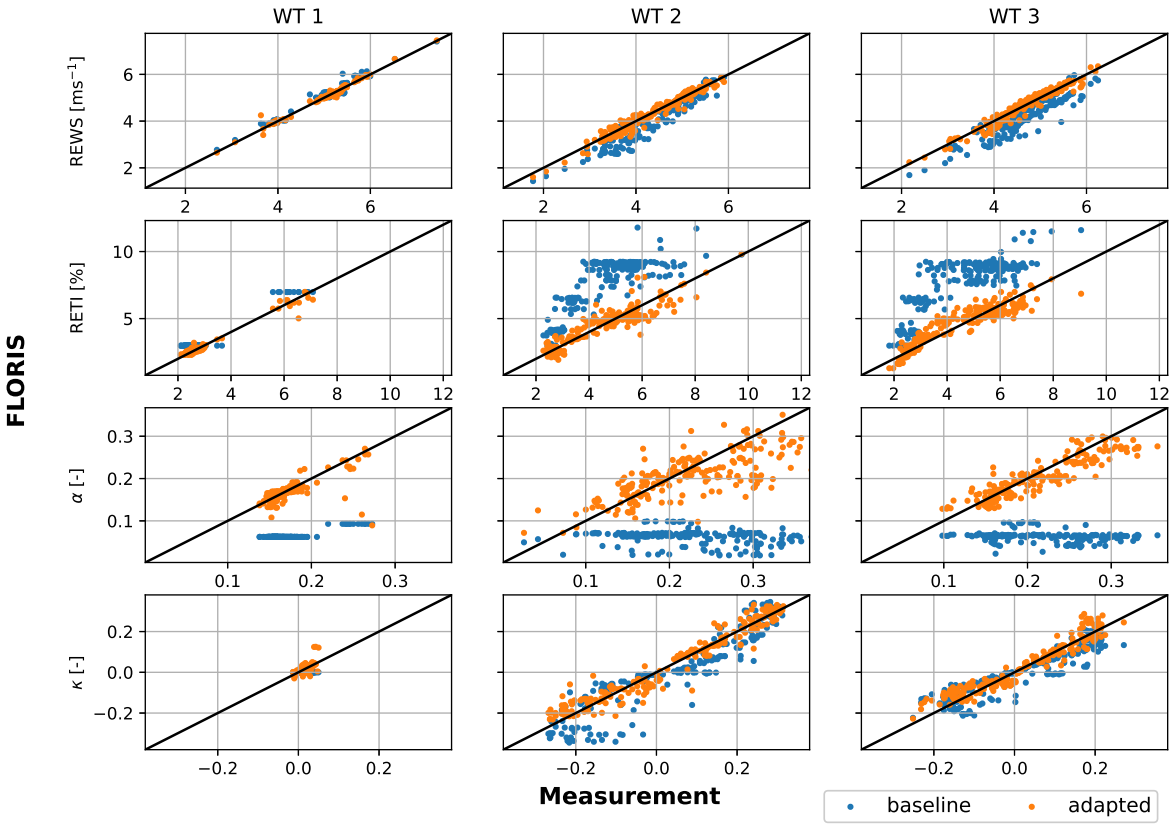

Figure 2: Comparison of estimated and measured power and REQs for the static campaign data set.

changes for different turntable positions, probably as a result of tunnel blockage. Therefore, a new set of nodes was placed for turntable positions $\phi = [-12, -8, -4, 0, 4, 8, 12]^\circ$. Besides identifying the speed-up factors, the adaptation method simultaneously calibrates the wake sub-models. This included the parameters $k_\alpha, k_\beta, k_a, k_b$ of the Gaussian wake model, in pairs responsible for far wake onset and decay, respectively [8], as well as the parameters TI_{constant}, TI_{ai}, TI_{initial}, and $TI_{\text{downstream}}$ of the Crespo-Hernandez wake-added turbulence model [9]. A maximum likelihood estimation identifies the correction of background flow and wake model parameters, based on a cost function that combines the residual for power as well as RETI estimations.

Besides power production, the flow model needs to predict REQs for load and damage estimation. A first analysis revealed that the estimates of some of the REQs provided by the adapted FLORIS model were still not satisfactory. For example, in the adopted implementation of FLORIS V3.2, the estimates of the vertical shear at the downstream turbine do not consider the effect of wakes. The output of the flow model was therefore further improved by a corrective model. This consists of a multi-layer neural network, trained to reduce the mismatch between the FLORIS-estimated REQs and the REQs measured in the static campaign.

Figure 2 reports the predictions of a baseline FLORIS and the final adapted version (that includes the additional corrector), versus the measured quantities of the static test campaign. The first row shows that the REWS of the adapted FLORIS agrees well with the observations, for all three turbines. This was expected and has been the focus of previous work, albeit in the form of power production, which behaves however very similar to the rotor effective wind speed

(see [6]). The remaining rows report the comparison between the other measured and estimated REQs. The estimates of the baseline model show a high discrepancy with respect to the RETI and shears measured at the waked turbines. On the other hand, the estimates greatly benefit from the additional data-driven correcting terms.

3.2. Load surrogate and damage model

The following two elements in the model tool chain estimate the loads and accumulate the fatigue damage as a result of the experienced inflow conditions. First, a load surrogate model calculates the damage equivalent loads $DEL^c_{\Delta t}$, for a duration Δt, corresponding to an $N^{eq}_{\Delta t}$ equivalent number of cycles [10, 11]. The model consists of a multi-layer neural network trained with observations of the same static test campaign used to adapt the flow model (refer to section 3.1). The model maps measurements of the REQs into DELs computed from the measured loads (refer to section 4.1). An additional input parameter to the load surrogate model is the yaw misalignment of the turbine. 70% of the data set was randomly selected for training.

The equivalent loads must then be converted into damage, which is calculated as

$$D^c_{\Delta t} = \frac{(DEL^c_{\Delta t})^m \cdot N^{eq}_{\Delta t}}{(2 \cdot L_u)^m}, \tag{4}$$

using the approach described in detail in [12]. Here, L_u represents the ultimate load and m represents the corresponding Wöhler exponent obtained using the material stress curve [13].

3.3. Economic model

The economic model acts as a combining element of power production and fatigue damage into a common metric. The expected revenue is formulated as a product of wind farm power output $\sum_i P^g_i$, where P^g_i denotes the power generation from turbine i, and the spot-market price $€_1$

$$\text{Revenue}(\boldsymbol{\gamma}, \mathbf{W}_a) = €_1 \cdot \textstyle\sum_i P^g_i(\boldsymbol{\gamma}, \mathbf{W}_a). \tag{5}$$

The formulated operation and maintenance (O&M) cost model is inspired by the formulations and discussions in [14], and writes

$$\text{Cost}(\boldsymbol{\gamma}, \mathbf{W}_a) = \sum_c \lambda^c \cdot \frac{€^c_2 \cdot P^r_i + €_1 \cdot CF \cdot \sum_i P^r_i \cdot t^{d,c}_i}{D^c_{ref}} \cdot \textstyle\sum_i D^c_i(\boldsymbol{\gamma}, \mathbf{W}_a)\ . \tag{6}$$

The formulation consists of two components: a fixed part $€_2$ scaled with the turbine rated power, capturing the fixed cost of operation as well as maintenance personnel, and a variable part, capturing the opportunity cost due to loss of revenue for the maintenance period. The loss of revenue is calculated as a product of spot-market price $€_1$, component specific generation downtime due to maintenance activity $t^{d,c}_i$ for turbine i, and an average generation obtained as the product of rated power output P^r_i of turbine i and wind farm capacity factor CF, summed over all the turbines and components. The cost of maintenance is further scaled with parameter λ, which denotes the average number of maintenance activities over the considered duration T.

4. Experimental setup

The experimental setup consists of a cluster of three G1 scaled wind turbines [3]. Each G1 features a rotor of diameter equal to 1.1 m, and has a rated rotor speed, power, and wind speed of 850 rpm, 46 W, and 5.7 ms^{-1}, respectively. The G1 is also equipped with active pitch, torque, and yaw, which are controlled in a closed-loop as described in [3]. The wake and aerodynamic performance of the G1 are representative of those of a hypothetical multi-MW wind turbine, as

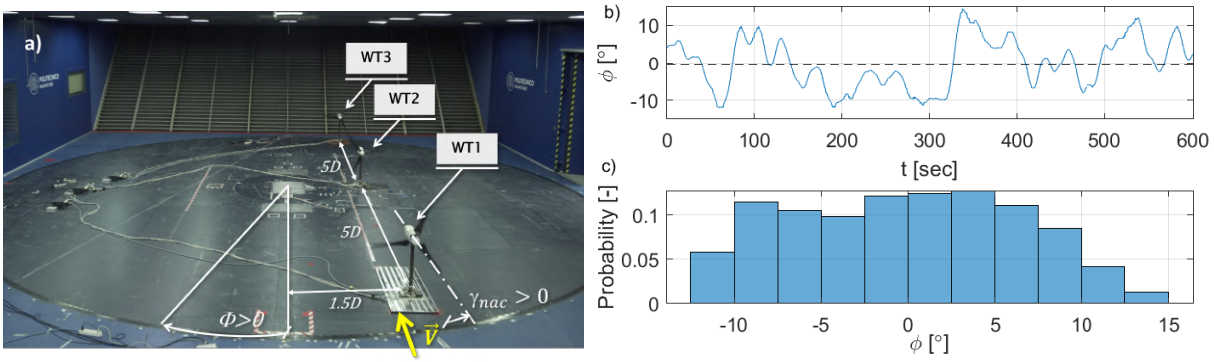

Figure 3: a) Experimental setup within the boundary layer test section of the wind tunnel of the Politecnico di Milano. b) Time history of the wind direction simulated through the rotation of the turntable. c) Corresponding distribution of wind direction.

discussed in [15]. The physical quantities (e.g. power, rotor speed, wake flow, etc) of this multi-MW machine can be derived, as explained in [16], from the length scale factor $n_\mathrm{l} = 1/162.1$ and the time compression ratio $n_\mathrm{t} = 1/82.5$.

The machines are installed on the turntable of the boundary layer test section of the wind tunnel of the Politecnico di Milano [4], and arranged in an aligned configuration with a longitudinal spacing of five rotor diameters, as shown in Fig. 3a. The cluster is embedded within a simulated Atmospheric Boundary Layer (ABL), also used for the experiments described in [4], characterised by a turbulence intensity at hub height of approximately 6%, and a sheared inflow featuring a power-law exponent $\alpha = 0.14$. Through the actuation of the turntable, the dynamic variation of the wind direction ϕ is also simulated, with ϕ null when the turbine row is parallel to the wind tunnel centerline.

In this work, a 10-minute interval of the 90-minute time series of the turntable rotation adopted in [4] is used, which aimed at reproducing the wind direction variability observed at an onshore test site located in northern Germany. The 10-minute interval, depicted in Fig. 3b, was selected to have the same number of occurrences of positive and negative wind directions, as highlighted by the median value of the 10-minute time series (black dashed line in Fig. 3b). This results in a similar distribution of left and right wake-to-rotor overlaps. The resulting distribution of wind direction is depicted in Fig. 3c.

During the experiments, the wind speed V_Pitot in the wind tunnel was measured by a Pitot tube, placed at hub height and 4 rotor diameters in front of the upstream machine (WT1). The power of the wind tunnel fans was kept constant during the experiment, with the Pitot tube measuring an average wind speed of 5.25 ms^{-1}, i.e. lower than the rated one.

4.1. Measured and estimated quantities

Strain gauges were used to measure the two out-of-plane bending moments and the torque on the rotating shaft, midway between the rotor disk and the aft bearing. Similarly, the north-south and east-west bending moments are measured at the tower base. Measurements of the blade pitch, rotor azimuth, and nacelle position were provided by optical encoders, while the power was computed as the product of rotor speed and torque shaft.

The bending moments measured at the tower base are used to estimate the DEL of the combined tower bending moments, obtained by projecting the corresponding two orthogonal components on the direction associated with the maximum DEL.

The inflow observer described in [17] is used to estimate the local wind speed at each turbine.

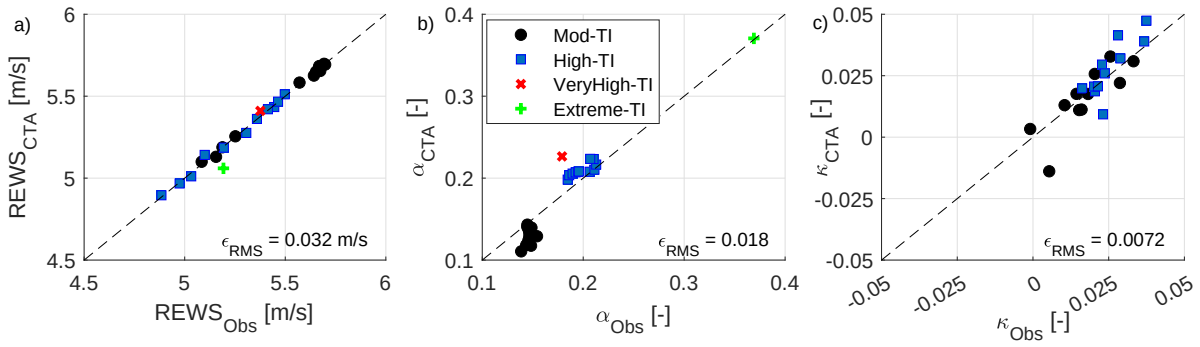

Figure 4: Comparison between REWS (a), power law exponent α (b), and horizontal shear κ (c) measured with CTA hot-wire probes $(\cdot)_{\mathrm{CTA}}$, and estimated through the tuned inflow observer $(\cdot)_{\mathrm{Obs}}$. The plots also include the root-mean-squared error.

The method, which has been validated with wind tunnel [18] and field [19] data, employs the power C_{P} and the cone C_{m} coefficients of the G1, derived with the Blade Element Momentum (BEM) model described in [20]. The C_{P} and C_{m} are used to compute lookup tables (LUTs) that return the local wind speed at the azimuthal position occupied by a blade, given the aerodynamic torque and the two out-of-plane bending moments measured on the shaft, as well as the measured rotor speed, blade pitch, and air density. Proper manipulation of the local wind speed [19] allows for estimating the REQs, as defined in section 3. The RETI is obtained by computing the ratio between the mean and standard deviation of the REWS, assessed over a moving time window of 7.5 sec, which corresponds to approx. 10 minutes at full scale.

4.2. Tuning of the inflow observer

The observer implemented on the G1 was tuned using measurements of the inflow performed 4 rotor diameters upstream of WT1. Specifically, CTA hot-wire probes were used to map 4 different simulated ABLs characterized by approximately 6% (Mod-TI), 13% (High-TI), 15% (VeryHigh-TI), and 22% (Extreme-TI) turbulence intensity at hub height. For the Mod-TI and High-TI ABLs, a mapping was performed over the entire test section, showing some lateral nonuniformity of the inflow [6]. For the other two ABLs, the inflow was instead measured only along a vertical line. Tests were then conducted with WT1 located at different lateral positions along the test section and with the four aforementioned ABLs. The data measured on board, and the inflow mappings, were used for the fine-tuning of the observer parameters, including the azimuth bias ψ_{bias} used to compensate for the unmodelled effects of blade dynamics. A comparison of the average of the observer estimates with the REWS and shears derived from the inflow mappings is shown in Fig. 4. The small root-mean-squared errors ϵ_{RMS} of the estimates, reported within the plots, highlight the proper tuning of the inflow observer.

5. Results

5.1. Scenario definition

Two distinct scenarios were defined for the experiment, which showcase two different possible applications of lifetime-aware wind farm control. The key aspects of these two scenarios are summarised in Tab. 1. The first "pre-construction" scenario considers a case where the farm has not yet commenced operation or is even still in the design phase, and a control strategy for the entire operational life of the farm should be determined. The second scenario "half-time strategy change" considers a case where the farm has been already in operation for half of its target lifetime. Here, the lifetime-aware control formulation is aimed at the remaining

lifespan of the assets, but it also considers the history of damage accumulation up to the point of its activation. Assuming the presence of an anomaly, the middle turbine (WT2) has accrued additional damage, which would lead to premature component failure under continued greedy operation.

Table 1: Key characteristics of the tested scenarios.

Scenario	Pre-construction	Half-time strategy change
Years in operation	0	10
Additional damage	-	+50% on middle turbine

The target operational life is $T = 20$ years. The remaining setup aspects are similar for both scenarios. It is evident, that it is not possible to simulate the entire operational life in the wind tunnel, even considering the time compression ratio achieved during a scaled experiment. Therefore, the experiment is simulating only a representative time period of environmental conditions. The simulated time history, described in section 4, corresponds to 13.3 hours of real time. It is assumed that these conditions are recurring throughout the life of the farm. As described, the inflow wind speed as well as turbulence level in the tunnel are constant during the experiment. The only varying component in $\mathbf{W}_a$ is therefore the wind direction, approximated with the frequency distribution in Fig. 3. In the current study, as component c, only the tower was considered to simplify the analysis. The corresponding load surrogate model for the tower displayed a quality measure of $r = 0.85$ and RMSE= 2.06 Nm for the validation part of the data set. The parameters of the economic model are assumed to stay constant over the lifetime. For revenue calculation, a fixed feed-in tariff was considered. For the O&M cost for the tower component, the parameters were adopted and scaled down from the range of full scale values utilised in [21, 22, 23, 24]: $CF = 0.3$, $€_1 = 0.02$ €/kWh, $€_2 = 3.6$ €/kWh, $t^d = 84$ h, $\lambda = 2$. These are based on various databases mentioned in [25].

Two load-unaware control strategies provide a reference performance used for comparison. Firstly, a greedy controller serves as the baseline case. Secondly, a power-maximising strategy with aggressive wake steering represents a typical, conventional power boosting controller. These strategies do not change depending on the scenario. The lifetime-aware strategies are referred to as "maximum profit". Figure 5 shows the obtained yaw offset schedules. It is visible, that

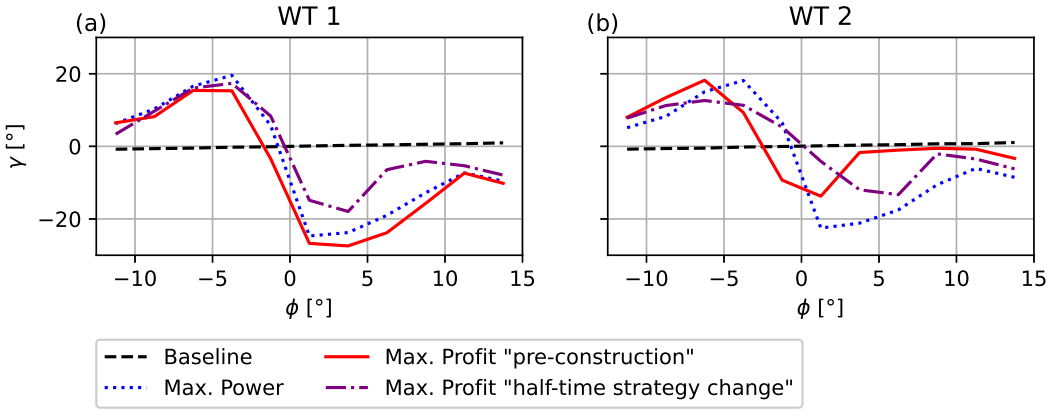

Figure 5: Yaw-offset schedules for the four tested strategies.

the maximum profit strategies are non-symmetric and differ depending on the scenario.

5.2. Economic performance

The test of each control strategy was repeated five times. The recorded tower DELs and power production were averaged in the post-processing over the five runs. The output is converted to a financial outcome by the economic model.

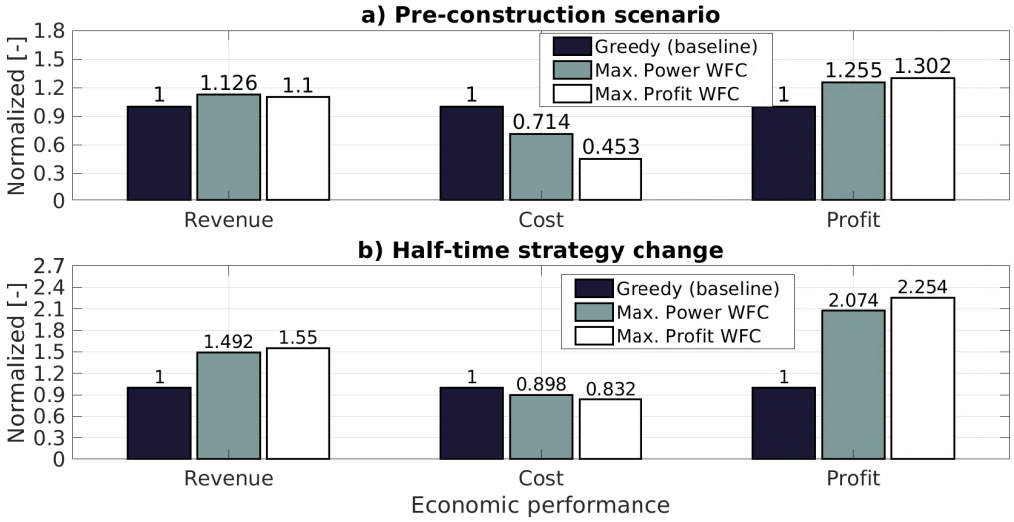

Figure 6: Economic performance over the farm lifetime for the three controllers normalised by the baseline control strategy.

Figure 6**a** reports the economic performance for the three tested control strategies for the "pre-construction" scenario. Results are reported in terms of revenue through power production, cost from accrued damage, and final profit. The respective numbers are normalised by the result of the baseline greedy controller. The results show that if pure revenue was the goal, the maximum power strategy would be naturally most suitable, as it achieves the highest gain with an increase of 12.6% over the baseline strategy. However, the lifetime-aware strategy achieves the highest reduction in maintenance costs. Interestingly, also the aggressive wake steering of the power boosting strategy seems to reduce costs. Ultimately, the highest financial gain is obtained for the maximum profit strategy.

Figure 6**b** reports the economic result for the second scenario "half-time strategy change". The farm is assumed to operate here for the remaining half of its design lifetime starting with an increased tower damage for WT2. As expected, the revenue of the two controllers employing wake steering (maximum power and maximum profit) is increased with respect to the baseline. However, due to the missing damage constraint, not only the greedy baseline, but also the power maximising controller, suffer from of a premature end-of-life of WT2. This leads to lost production and missing revenue. The damage constraint on the maximum profit strategy ensures that all turbines operate for the target 20 years. As a result, the novel control formulation results in a higher profit when applied to the second half of the farm lifetime.

6. Summary and outlook

The proposed lifetime-aware control framework allows owners and operators to make informed decisions on the economic viability of various alternative wind farm operational scenarios, and to implement those decisions by deploying a farm controller that guarantees economic optimality and the achievement of a life duration goal.

The study showcased firstly the design characteristics as well as a possible implementation of a lifetime-aware wind tunnel experiment. The presented model tool chain estimates revenue

and associated costs and is applicable in an optimisation algorithm. The two tested scenarios showcased promising application scenarios for economically optimised control strategies.

7. Acknowledgements

This work has been partially supported by the e-TWINS project (FKZ: 03EI6020), which receives funding from the German Federal Ministry for Economic Affairs and Climate Action (BMWK). This work has also been partially supported by the SUDOCO project, which receives funding from the European Union's Horizon Europe Programme under the grant agreement No. 101122256.

8. Author contributions

RB and AA contributed equally to the implementation of the method, FC conducted the wind tunnel experiments, CLB proposed the idea and supervised the work. All authors provided important input to this research work through discussions and feedback and by improving the manuscript.

References

[1] Ali C Kheirabadi et al. "A quantitative review of wind farm control with the objective of wind farm power maximization". In: *Journal of Wind Engineering and Industrial Aerodynamics* (2019).

[2] Matt Harrison et al. "An initial study into the potential of wind farm control to reduce fatigue loads and extend asset life". In: *Journal of Physics: Conference Series*. 2020.

[3] Carlo L. Bottasso and Filippo Campagnolo. "Wind Tunnel Testing of Wind Turbines and Farms". In: *Handbook of Wind Energy Aerodynamics*. 2022.

[4] Filippo Campagnolo et al. "Wind tunnel testing of wake steering with dynamic wind direction changes". In: *Wind Energy Science* (2020).

[5] NREL. *FLORIS. Version 3.2*. 2022.

[6] Filippo Campagnolo et al. "Further calibration and validation of FLORIS with wind tunnel data". In: *Journal of Physics: Conference Series*. IOP Publishing. 2022.

[7] R. Braunbehrens et al. "The wind farm as a sensor: learning and explaining orographic and plant-induced flow heterogeneities from operational data". In: *Wind Energy Science Discussions* (2022).

[8] Majid Bastankhah et al. "A new analytical model for wind-turbine wakes". In: *Renewable energy* (2014).

[9] Antonio Crespo, J Herna, et al. "Turbulence characteristics in wind-turbine wakes". In: *Journal of wind engineering and industrial aerodynamics* (1996).

[10] Herbert J. Sutherland. "On the Fatigue Analysis of Wind Turbines". In: *Sandia National Lab.* (). DOI: 10.2172/9460.

[11] Adrien Guillore et al. *A load surrogate model based on rotor-equivalent inflow quantities*. 2023.

[12] *MLife User's Guide for Version 1.00*. 2012.

[13] João Pacheco et al. "Fatigue Assessment of Wind Turbine Towers: Review of Processing Strategies with Illustrative Case Study". In: *Energies* (2022).

[14] "D2.1 Cost Model for fatigue degradation and O&M - Advanced integrated supervisory and wind turbine control for optimal operation of large Wind Power Plants". In: *TotalControl* (2022).

[15] Chengyu Wang et al. “How realistic are the wakes of scaled wind turbine models?” In: *Wind Energy Science* (2020).

[16] Helena Canet et al. “On the scaling of wind turbine rotors”. In: *Wind Energy Science* (2021).

[17] CL Bottasso et al. “Local wind speed estimation, with application to wake impingement detection”. In: *Renewable Energy* (2018).

[18] Filippo Campagnolo et al. “Wind tunnel validation of a wind observer for wind farm control”. In: *ISOPE International Ocean and Polar Engineering Conference*. 2017.

[19] Johannes Schreiber et al. “Field testing of a local wind inflow estimator and wake detector”. In: *Wind Energy Science* (2020).

[20] Chengyu Wang et al. “Identification of airfoil polars from uncertain experimental measurements”. In: *Wind Energy Science* (2020).

[21] Sebastian Pfaffel et al. “Performance and Reliability of Wind Turbines: A Review”. In: *Energies* (2017).

[22] Ángel M. Costa et al. “New Tendencies in Wind Energy Operation and Maintenance”. In: *Applied Sciences* (2021).

[23] Daniel Chan and John Mo. “Life Cycle Reliability and Maintenance Analyses of Wind Turbines”. In: *Energy Procedia* (2017).

[24] R. Poore. “Development of an Operations and Maintenance Cost Model to Identify Cost of Energy Savings for Low Wind Speed Turbines”. In: *Tech Report National Renewable Energy Lab.* (2008).

[25] Roozbeh Bakhshi et al. “Overview of Wind Turbine Field Failure Databases: A Discussion of the Requirements for an Analysis”. In: *ASME Power Conference* (2018).