



Delft University of Technology

Closed-Loop Model-Predictive Wind Farm Flow Control Under Time-Varying Inflow Using FLORIDyn

Becker, Marcus; van den Broek, Maarten J; Allaerts, Dries; van Wingerden, Jan Willem

DOI

[10.1002/we.70044](https://doi.org/10.1002/we.70044)

Publication date

2025

Document Version

Final published version

Published in

Wind Energy

Citation (APA)

Becker, M., van den Broek, M.J., Allaerts, D., & van Wingerden, J. W. (2025). Closed-Loop Model-Predictive Wind Farm Flow Control Under Time-Varying Inflow Using FLORIDyn. *Wind Energy*, 28(9), Article e70044. <https://doi.org/10.1002/we.70044>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

RESEARCH ARTICLE OPEN ACCESS

Closed-Loop Model-Predictive Wind Farm Flow Control Under Time-Varying Inflow Using FLORIDyn

Marcus Becker¹  | Maarten J. van den Broek²  | Dries Allaerts³  | Jan-Willem van Wingerden¹ 

¹Delft Centre for Systems and Control, TU Delft, Delft, Netherlands | ²sowento GmbH, Stuttgart, Germany | ³Faculty of Aerospace Engineering, TU Delft, Delft, Netherlands

Correspondence: Marcus Becker (marcus.becker@tudelft.nl)

Received: 21 December 2024 | **Revised:** 9 May 2025 | **Accepted:** 10 June 2025

Funding: This work is part of the research programme “Robust closed-loop wake steering for large densely spaced wind farms” with project number 17512, which is (partly) financed by the Dutch Research Council (NWO). This work has been supported by the SUDOCO project, which receives the funding from the European Union’s Horizon Europe Programme under the grant no. 101122256.

Keywords: closed-loop control | dynamic wake model | wake steering | wind farm flow control

ABSTRACT

Wind farm flow control has been a key research focus in recent years, driven by the idea that a collectively operating wind farm can outperform individually controlled turbines. Control strategies are predominantly applied in an open-loop manner, where the current flow conditions are used to look up precomputed steady-state set points. Closed-loop control approaches, on the other hand, take measurements from the farm into account and optimize their set points online, which makes them more flexible and resilient. This paper introduces a closed-loop model-predictive wind farm controller using the dynamic engineering model FLORIDyn to maximize the energy generated by a ten-turbine wind farm. The framework consists of an Ensemble Kalman Filter to continuously correct the flow field estimate, as well as a novel optimization strategy. To this end, the paper discusses two dynamic ways to maximize the farm energy and compares this to the current look-up table industry standard. The framework relies solely on turbine measurements without using a flow field preview. In a 3-h case study with time-varying conditions, the derived controllers achieve an overall energy gain of 3% to 4.4% with noise-free wind direction measurements. If disturbed and biased measurements are used, this performance decreases to 1.9% to 3% over the greedy control baseline with the same measurements. The comparison to look-up table controllers shows that the closed-loop framework performance is more robust to disturbed measurements but can only match the performance in noise-free conditions.

1 | Introduction

A switch away from fossil fuels to less greenhouse gas-emitting (GHG) sources of energy is necessary to prevent a climate crisis [1]. Wind energy is one alternative that provides energy at a fraction of the GHG emissions. Wind farms are, therefore, an essential part of the energy transition. However, they do not provide the maximum amount of energy they could. How come? This is in part due to the way that turbines interact: As a wind turbine converts kinetic energy from the surrounding airflow

into electricity, it leaves behind an area of low wind speed called a wake. In wind farms, these wakes will likely influence downstream turbines, lowering their power output. Wind farm flow control (WFFC) strategies aim to mitigate this effect by manipulating the wake shape.

In this work, we focus on model-based approaches to WFFC. With this approach, a model of the farm is used as a surrogate for the real wind farm. The model predicts how the turbine wakes behave, given the atmospheric conditions and turbine

Dries Allaerts passed away in October 2024 and greatly contributed to the publication as a supervisor and during the conceptualization.

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2025 The Author(s). *Wind Energy* published by John Wiley & Sons Ltd.

states. Different types of models exist with varying costs and capabilities. On the one end, high-fidelity models like Large Eddy Simulations (LES) provide the most insight into the physical phenomena that take place in a wind farm, for example, Churchfield et al. [2]; Chatelain et al. [3]. These models are typically too demanding for control which has motivated model simplifications to achieve faster computational speeds. Medium-fidelity models capture a coarse image of the flow. Two-dimensional Reynolds-averaged Navier-Stokes solvers like Rott et al. [4]; Boersma et al. [5]; van den Broek and van Wingerden [6] fall in this category. They aim to capture the core wake dynamics at a reduced cost by limiting the dimensionality of the flow. However, later research showed that this simplification can render them useless for wake steering applications [7]. Other models in the category of medium-fidelity models simulate the wake propagation based on synthetic turbulence. The Dynamic Wake Meandering (DWM) model, introduced by Larsen [8], propagates the wake as a series of turbulence boxes. The propagation speed and direction are then determined by the contents of the box. This approach, coupled with an aeroelastic turbine model, can give estimates of loads onto the turbine and its structure. Successor models like FAST.Farm [9] and HAWC2Farm [10] provide a basis to investigate damage equivalent loads on a farm scale at a relatively low computational cost. Another approach to model wind turbine wakes is to simulate free-vortex particles shed by the rotor, for example, Marten [11]; van den Broek et al. [12]. The downside of these models is that they tend to become numerically unstable in the far-wake region and under turbulent conditions. Low-fidelity models provide a flow field prediction at a very low computational cost. This is achieved by modeling the wake shape and wind speed reduction as a set of analytical equations, for example, Bastankhah and Porté-Agel [13]. These models typically predict the steady-state wake shape of a single turbine wake and use superposition methods to combine the effect of multiple wakes in a farm. Within WFFC applications, they are used to test and optimize yaw-angle set points for an entire farm.

Based on the computational speed provided by the steady-state engineering models, paired with the success of the DWM model, an additional group of models has been proposed: These models use passive Lagrangian tracers, called Observation Points, to propagate turbine and flow field states from each turbine downstream. The Flow Redirection and induction Dynamics model FLORIDyn initially proposed this approach [14]. Since its introduction, the model has since been revised [15] and further

developed [16]. Similar modeling approaches exist, for example, Lejeune et al. [17]; Braunbehrens et al. [18]. Their simplicity, simulated wake dynamics, and speed make them attractive for model-based WFFC applications.

Model-based WFFC is predominantly applied in an open-loop configuration [19]. This means that the yaw angle set points are optimized ahead of time; see Figure 1b. This is mainly done by employing a steady-state wake model and selecting a set of ambient conditions. For each ambient condition, the yaw angle set points are then optimized and stored in a Look-up Table (LuT). During operation, the current ambient conditions, such as wind speed and direction, are identified and used to look up the optimal yaw angles. Examples of this strategy can be found in the field experiments conducted by Fleming et al. [20]; Duc et al. [21]; Doekemeijer et al. [22]. The issue with this approach is that the strategy can not react to unforeseen circumstances. Examples of this could be wind farm layout changes due to offline turbines, unmodeled or incorrectly modeled turbine interactions, or heterogeneous and changing flow conditions.

This is addressed by closing the loop, see Figure 1a, where the model is continuously updated by measurements. A recent approach is to close the loop on the parameters of the steady-state model [23–26]. This entails a coupling between the measured flow conditions and the modeled ones. The difference is then fed back to correct aspects like the wake expansion. In Bachant et al. [27], they employ a similar strategy, but rather than correcting the model parameters, they build a corrector for the model output. This version of closed-loop control still inherently neglects the wake dynamics. Closed-loop wind farm flow control using a dynamic wake model like FLORIDyn is still a poorly explored area of research. Gebraad et al. [28] employs the first version of the FLORIDyn model paired with a Kalman Filter to estimate the wind speed in a six-turbine wind farm. The same publication also uses FLORIDyn to perform a yaw angle optimization; however, not in closed-loop.

Control for energy maximization using dynamic models uses predominantly a receding horizon approach called Model-predictive control (MPC), for example, Goit and Meyers [29]; Munters and Meyers [30]; van den Broek et al. [12]; Becker et al. [31]. This means that the optimization is done over a pre-determined prediction horizon. The surrogate model is used to predict how the cost function will be impacted by the control

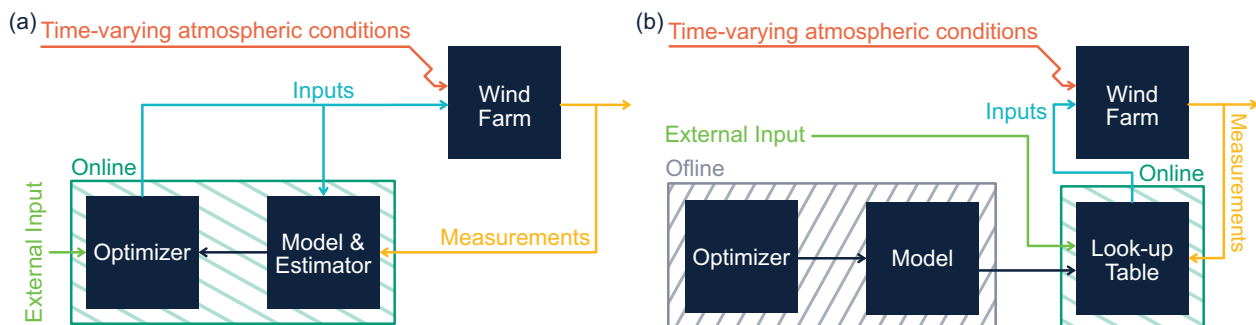


FIGURE 1 | (a) Closed-loop compared with (b) open-loop wind farm flow control. The figure is adapted from Meyers et al. [19].

actions taken. An optimizer is then used to determine the ideal control actions. The optimization is followed by an update time step, which is typically smaller than the prediction horizon. During this time, the previously optimized time series of control set points is applied. After the update time step has passed, a new optimization is done. MPC is typically applied to follow a reference value, for example, a reference wind farm power [32]. In contrast, this work aims to maximize a cost function, which leads to a different optimization problem. This is referred to as economic MPC (eMPC).

A re-occurring assumption in this context is the full knowledge of the flow field (e.g., Goit and Meyers [29]; Goit et al. [33]; Munters and Meyers [30]; Janssens and Meyers [34]), and full knowledge of the environmental changes ahead (e.g., van den Broek et al. [12, 35]). The latter is often referred to as preview and can lead to significant gains over preview-less control approaches, also for steady-state control approaches [35–37]. Yet, it is not clear how this preview information may be attained in a realistic WFFC scenario. The same holds true for the previously mentioned assumption of full flow field knowledge. State-estimation techniques can provide some of this knowledge. They use sensor data to correct the current state of the surrogate model. This is done by comparing predicted measurements to taken ones. Looking at FLORIDyn specifically, in the work by Gebraad et al. [28], the turbine power generated is used to correct the local wind speed estimate in FLORIDyn by employing a Kalman Filter. Similarly, Becker et al. [38] also uses the power generated, as well as wind direction measurements, to correct a heterogeneous flow field state in FLORIDyn using an Ensemble Kalman Filter (EnKF). This methodology has the advantage that it provides the state estimate as well as an uncertainty estimate. Additionally, it does not require a linear(-ization) of the model. In the later work of Braunbehrens et al. [39] and Di Cave et al. [40], the wake location is corrected by using an EnKF and a downstream turbine as a coarse sensor [41].

The literature review shows that there is a gap in closed-loop wind farm flow control approaches using dynamic surrogate models in time-varying conditions. Prior work has developed surrogate models like FLORIDyn to dynamically model the wakes in a farm. It has also proposed state-estimation frameworks to align the model's state with the true state. How to then use FLORIDyn to obtain the optimal control set points is still an unexplored topic.

This paper, therefore, proposes a novel closed-loop economic model predictive control framework for farm-wide energy maximization. To this end, the paper builds upon previous work with the following four main contributions: (i) a discussion around the nature of the cost functions for energy and power maximization, (ii) a dimensionality reduction of the optimization problem, which also leads to a realistic turbine operation, (iii) an implicit way to limit misalignment angles and to avoid nonlinear constraints, and (iv) the derivation of a closed-loop framework for dynamic model-based wind farm flow control. The proposed framework focuses on readily available sensor data, like turbine power and wind direction measurements at the turbine locations, and in contrast to previous work, does not assume knowledge of the full flow field nor a preview of future flow field changes. A block diagram of the closed-loop design is depicted in Figure 2. The derived optimization strategy is applied in a ten-turbine closed-loop case study. This is done with a wind farm simulated in LES under turbulent conditions and with realistic time-varying wind direction changes. In this context, the paper also addresses the impact of noisy and biased measurements on controller performance and yaw travel.

The remainder of the paper is structured as follows: Section 2 discusses the methodology of the closed-loop approach. This consists of a description of the model (Section 2.1), the state estimator (Section 2.2), and the novel discussion related to the controller (Section 2.3). The simulation methods are discussed in Section 3. This entails the description of the high-fidelity environment, as well as the tested wind conditions, sensor data, and controller settings. The proposed framework is tested in Section 4, where the results are presented. Lastly, Section 5 draws a conclusion of the work.

2 | Methodology

This section introduces the components of the closed-loop control approach presented in this paper: Section 2.1 describes the basics of the dynamic surrogate model, Section 2.2 of the state estimation. The main contribution of this paper, the controller, is discussed in Section 2.3. The list of parameters used for the components introduced in Section 2 can be found in Appendix A, alongside the tuning approach.

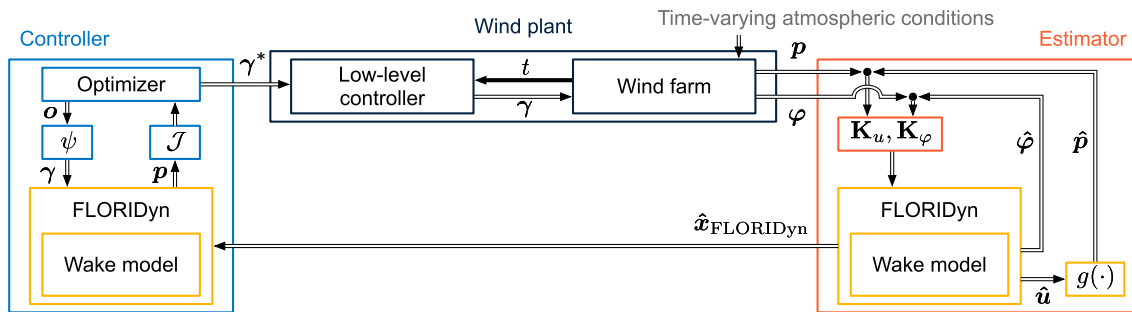


FIGURE 2 | Closed-loop design applied in this paper. The wind farm provides power and wind direction measurements at the turbine locations, which are used to correct the FLORIDyn model state. The identified state is then used to optimize the future yaw angle set points for all turbines. This is passed on to a low-level controller, which applies the set points.

2.1 | Surrogate Model

The Flow Redirection and Induction Dynamics Model (FLORIDyn) used in this study is based on the work presented in Becker et al. [16], with the extensions done in Becker et al. [38]. It simulates wake dynamics by employing so-called Observation Points (OPs), which carry states from the rotor plane downstream. The states one OP possesses are (i) turbine states (e.g., yaw angle), (ii) flow field states (e.g., wind direction), and (iii) its own positional states. The states are initialized at the rotor plane based on what the turbine measures. The OP position is advanced based on the wind speed and direction at its position.

While previously each OP would rely on its own state of the wind speed and direction to propagate, the wind speed and direction are now the result of a weighted average. This way, the OP takes neighboring OP states into account. The spatiotemporal Gaussian weighting function was first adapted by Lejeune et al. [17] for a FLORIDyn similar model and adapted in Becker et al. [38] to allow for better correlation between the OPs. This benefits the assumptions made with the state estimation; see Section 2.2. The FLORIDyn framework uses a Gaussian wake to model the wake deficit, wind speed reduction and wake deflection [13].

Unlike recently presented results (e.g., van den Broek et al. [35]) the simulations presented in this paper do not utilize any synthetic preview information of the wind direction. As previously mentioned, FLORIDyn does have flow field information stored in the OPs, which is propagated downstream. This provides other turbines with an estimate of the flow field changes ahead. During the prediction phase, each time step, the free stream turbines copy their previous state, and downstream turbines adapt the wind speed and direction by the weighted average of the data stored in the surrounding OPs.

2.2 | State Estimation

The employed state estimation is based on Becker et al. [38] and uses an Ensemble Kalman Filter (EnKF) [42, 43]. The EnKF was chosen for three main reasons: (i) it allows the state estimation of nonlinear systems without derivation of a linearized model, (ii) it provides uncertainties of the estimated states, which can be used for robust decision making, and (iii) it has previously been used for similar applications by Doekemeijer et al. [44] to estimate the flow field state of a wind farm simulation in the 2D solver WFSim [5]. Note that other state estimation formulations for FLORIDyn-like systems do exist. Gebraad et al. [28], for instance, employ an extended Kalman Filter. What distinguishes a dynamic system like FLORIDyn from other EnKF applications is that FLORIDyn is not a grid-based simulation but rather attaches its state to particles. Therefore, the flow is described by where the particles are, which may be different across the ensembles. The state estimation framework used addresses this issue by projecting the OPs of all Ensembles onto a common set, which is then corrected. To this end, two Kalman gain matrices (\mathbf{K}_u for wind speed and \mathbf{K}_φ for the wind direction) are derived from the common output matrices \mathbf{C}_u and \mathbf{C}_φ , as well as the

correlation of the ensembles. This is done differently to the state estimator proposed in Becker et al. [38], where only the wind speed would be corrected this way, while the wind direction would be corrected with ensemble individual Kalman Gain matrices. The correction loop is depicted in Figure 3. The matrices \mathbf{W} define the weighted projection onto the common states, \mathbf{W}^{-1} their inverse. As the inverse is generally not attainable, we set it to be equal to the identity matrix. Because \mathbf{W} does have a sparse and strongly diagonal shape, this approximation is valid. For more information on how the correction is derived we refer to the previously mentioned sources. There are EnKF-based designs for FLORIDyn-like models to correct the wake center, but these fall outside of the scope of this paper [39, 40]. Also outside of the scope of this paper fall online parameter estimation methods using an EnKF [23, 24].

2.3 | Controller

Section 2.3.1 compares the differences between the steady-state cost function formulation and two dynamic cost functions. In this section, as well as in Section 2.3.2, the receding horizon optimization problem is explained, where control actions taken within an action horizon are evaluated over a prediction horizon. Section 2.3.2 to 2.3.4 further discuss measures to simplify the optimization problem solved at runtime: Section 2.3.2 investigates a basis function approach to derive the yaw angle time series, followed by a methodology to incorporate constraints into the cost function (Section 2.3.3). Section 2.3.4 discusses how the wind farm is decomposed into smaller sections to reduce the number of turbines per optimization. Lastly, Section 2.4 discusses the reference controllers.

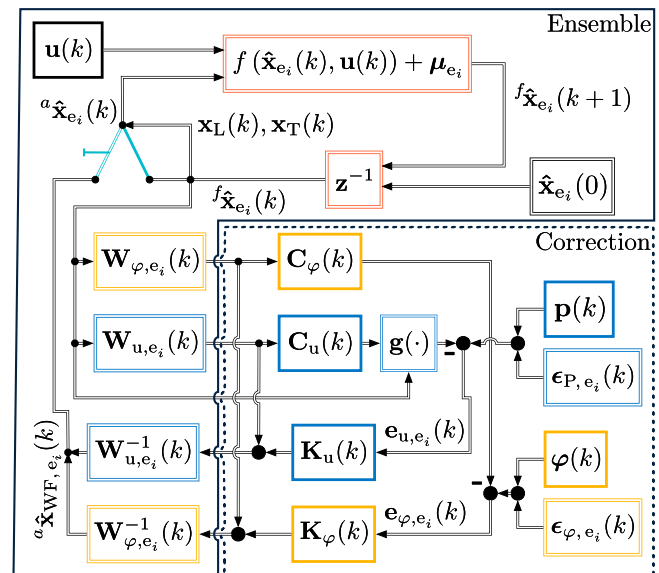


FIGURE 3 | State estimation cycle for an ensemble e_i . The upper section of the figure depicts the evolution of the wind farm simulation, here depicted by $f(\cdot)$, the lower left section the extraction of the measured values from the simulation. The lower right part depicts the correction based on the mismatch of the predicted and recorded measurements. Elements with a double outline are ensemble-specific, elements with a single outline are the same for all ensembles.

2.3.1 | Cost Functions

Three cost functions are compared in this work - (i) the steady state power, (ii) the accumulated energy, and (iii) the accumulated energy disregarding the time delays between the turbines. Each function is explained in further detail below.

The steady-state approach to formulate a cost function $\mathcal{J}_{\text{steady state}}$ to maximize the farm energy is to sum the power of all n_T turbines:

$$\mathcal{J}_{\text{steady state}}(\theta) = \sum_{i=1}^{n_T} p_i(\theta) \quad (1)$$

where p_i is the power generated by turbine i and θ are the optimization variables. Time is implicitly taken into account, as steady-state models do not assume any delay effects. If time is taken into account, the energy is calculated as the integral of the turbine power generated over time. Equation (1) is extended as follows to form the dynamic cost function:

$$\mathcal{J}_{\text{dynamic}}(\theta) = \Delta t \sum_{i=1}^{n_T} \sum_{k=1}^{\tau_{\text{ph}}} p_i(\theta, k) \quad (2)$$

where $p_i(\theta, k)$ is the power generated by turbine i at time step k , and Δt is the time step of the simulation. The cost function now approximates the energy generated by the farm across all time steps $k \in [1, \tau_{\text{ph}}]$, where τ_{ph} is the prediction horizon of the optimization. This additional time dependence inherently changes the result for θ that we get: The power loss due to yaw control actuation needs to be recouped within the prediction horizon for the control action to be profitable. This is in strong contrast to the steady-state cost function, which assumes to have eternity available to recoup the cost of control actions.

The following example further illustrates this difference based on a two-turbine wind farm: turbine T1 is located $5D$ downstream and $0.5D$ cross stream from the upwind turbine T0, D being the turbine diameter. This configuration favors positive yaw misalignment from the first turbine for power and energy maximization. In this experiment, the control set point of T0 is varied, and the power generated from both turbines is recorded in FLORIDyn. At $t = 0$ s, T0 starts fully aligned with the wind direction and then changes its orientation with 0.3 deg s^{-1} to the reference misalignment angle. The chosen yaw speed is within the commonly used range with up to 0.5 deg s^{-1} [45]. Figure 4 shows the efficiency of the yaw misalignment angle set point over time, (a) based on the power generated, and (b) based on the energy generated up to the given time step. Both figures show the initial loss connected to the misalignment of turbine T0. Turbine T1 starts to experience the benefits of the misalignment after ≈ 150 to 200 s. At this point, the steady state optimum is reached for $\gamma_{T0} \approx +16$ deg, which leads to a 7% increase in power generated compared with the baseline case. The energy generated, however, has integrated the losses leading up to the favorable wake redirection state. This means that this loss has to be recouped first before a yaw misalignment becomes a better choice than the baseline behavior. The more time is given, the more attractive larger yaw angles become: If the prediction horizon is 100 s long, $\gamma_{T0} = 0$ deg is the optimal solution




FIGURE 4 | Wind farm power (a), energy (b), and shifted energy (c) efficiency of a two turbine wind farm based on the yaw control set point of T0, normalized by the case of no misalignment. Turbine T0 starts at $t = 0$ s with $\gamma_{T0} = 0$ deg and continuously changes its angle with 0.3 deg s^{-1} to the set point, where it remains. Turbine T1 is fully aligned with the wind direction. The dotted line shows the break-even line, the continuous line the ideal value of γ_{T0} based on the time. As the shifted approach (c) moves the power signal by T1 in time, the data are shorter.

to Equation (2), for 300 s it is $\gamma_{T0} \approx 8$ deg and for 500 s $\gamma_{T0} \approx 12$ deg is optimal. With an infinite prediction horizon, the optimal solution of Equation (2) converges to the optimal solution of Equation (1), as the initial loss becomes increasingly negligible compared with the accumulated energy. This inherent characteristic of Equation (2) offers a conservative security, as it guarantees to recoup the initial investment in the given time frame under the assumption that conditions are steady. From this perspective wake steering becomes an investment-and-return problem, threatened by changing environmental conditions.

The described characteristics can also lead to unwanted behavior - if enough control degrees of freedom are given, T0 will initially engage in more aggressive control actions, followed by greedy behavior at the end of the time horizon. This is done as the downsides of the greedy control actions towards T1 fall outside of the time horizon. This is known as the “turnpike effect” and needs to be considered for dynamic wind farm flow control applications [12]. Ways to limit the turnpike effect is to not allow control actions at the end of the time horizon or to penalize them as part of the cost function.

An alternative is offered by Becker et al. [31]: The presented approach aims to synchronize the cause and effect of the control action with its impact on the actuated turbine as well as the downstream turbines. This formulation leads to a spatio-temporal split of the optimization problem into several smaller problems: (i) the spatial split ensures that wind turbines that

do not influence one another are optimized separately, while (ii) the temporal split connects the cost function of the yaw angles with their impact. The following example illustrates the working of the algorithm using the same two-turbine wind farm. The free wind speed is set to 8ms^{-1} , and the wake advection speed is assumed to be equal to the free stream velocity. The simulation time step is $\Delta t = 5\text{ s}$. It then takes $800\text{ m}/8\text{ ms}^{-1} = 100\text{ s}$ or $\Delta k = 20$ time steps for one particle of air to reach the downstream turbine. If the time window during which a control action can be taken, the action horizon τ_{ah} , is set to $\tau_{\text{ah}} = 10$ time steps, the cost function of the yaw trajectory γ_{T0} of the upstream turbine T0 is formulated as follows:

$$\min \mathcal{J}_{T0}(\gamma_{T0}) = -\Delta t \sum_{k=1}^{\tau_{\text{ah}}} p_0(\gamma_{T0}, k) + p_1(\gamma_{T0}, \gamma_{T1}, k + \Delta k).$$

This formulation sums the power of T0 during its action horizon and the power of T1 once the control actions arrive at the turbine. Note that for this formulation, the two turbines are simulated from $k = 1$ until $\tau_{\text{ph},T0} = \tau_{\text{ah}} + \Delta k$, but only sections of the prediction horizon are taken into account for the cost function. The power of T1 depends on both γ_{T0} and γ_{T1} . Because $\Delta k_{1 \leftarrow 0} > \tau_{\text{ah}}$, the yaw trajectory γ_{T1} can be optimized in a separate cost function:

$$\min \mathcal{J}_{T1}(\gamma_{T1}) = -\Delta t \sum_{k=1}^{\tau_{\text{ah}}} p_1(\gamma_{T1}, k).$$

For this cost function $\tau_{\text{ph},T1} = \tau_{\text{ah}}$. As a result we can formulate two optimization problems, $\mathcal{O}_1 = \min \mathcal{J}_{T0}$ and $\mathcal{O}_2 = \min \mathcal{J}_{T1}$. In this example, \mathcal{O}_1 depends on γ_{T1} , and thus on the result of \mathcal{O}_2 . Therefore, \mathcal{O}_1 cannot be solved independently. \mathcal{O}_2 is only evaluated over τ_{ah} time steps, and \mathcal{O}_1 over $\tau_{\text{ah}} + \Delta k_{1 \leftarrow 0}$ time steps, which is generally less than the otherwise necessary τ_{ph} to achieve similar yaw steering. Additionally, the smaller optimization problems can lead to better convergence within the given optimization budget. More details about the time-shifted approach can be found in Becker et al. [31] or in a similar application in Ciri et al. [46]. Figure 4c shows how the optimization landscape is affected by the change.

2.3.2 | Basis Function Approach

The optimized yaw time series must fulfill two constraints: (i) its rate of change needs to be lower or equal to the possible rate of change r_γ , limited by the actuators, and (ii) it should not exceed the prescribed maximum misalignment boundaries. This section focuses on the former aspect, while the latter is discussed in Section 2.3.3. In addition to the limited rate of change, it is desirable to derive yaw signals that can be realistically applied to existing yaw systems. Current yaw systems are comprised of a yaw drive, a bearing, and a brake [45]. To move the turbine, the brake is released and the yaw drive moves the nacelle with the maximum rate of change to the reference position. The brake is then engaged again to maintain the reference orientation. This behavior limits the set of feasible yaw time series, as, for example, a continuously changing yaw signal is undesirable. This can be exploited to simplify the optimization problem and to reduce

the number of inputs to the optimization. To this end, we rely on a basis function $\psi(\theta, t_n)$, which takes two optimization parameters, o_1 and o_2 as arguments, as well as $t_n \in [0, 1]$, which is the normalized time within the action horizon t_{AH} . The idea is to either in- or decrease the turbine's yaw angle with the maximum rate of change and to allow the change period to start at an early or late point in time within t_{AH} . The function ψ has been designed in such a way that $o_1, o_2 \in [0, 1]$, which provides a generic interface for an optimization algorithm and can improve its effectiveness of it. The basis function is further determined by $r_{\gamma,n} = r_\gamma t_{\text{AH}}$, the maximum rate of change within the normalized parameter space.

$$\begin{aligned} \psi(o, t_n) &= 2[o_1 - 0.5] \text{sat}_{[0,1]} \left(\frac{t_n - t_{s,n}(o)}{2[o_1 - 0.5]} \right) r_{\gamma,n}, \\ t_{s,n}(o) &= o_2 [1 - 2 \cdot |o_1 - 0.5|], \end{aligned} \quad (3)$$

where $\text{sat}_{[0,1]}(x)$ is a saturation function which is equal to 0 for $x \leq 0$, x for $x \in [0, 1]$, and 1 otherwise. The variable o_1 determines in which direction the orientation changes and for how long, e.g., $o_1 = 1$ is a constant increase across the entire action horizon, $o_1 = 0.5$ results in a steady yaw angle. The second optimization variable, o_2 , moves the starting point of the yaw angle change period. A possible result of ψ and the resulting $\Delta\gamma$ is shown in Figure 5a. The final yaw orientation trajectory is calculated as

$$\gamma(t) = \gamma(0) + \psi \left(o, \frac{t}{t_{\text{AH}}} \right) \quad \forall t \in [0, t_{\text{AH}}] \quad (4)$$

An example is given in Figure 5b. While the proposed function does limit the number of optimization parameters per turbine, it contains undesired regions of insensitivity. This namely affects the case when $o_1 = 0.5$, which results in no yaw change and

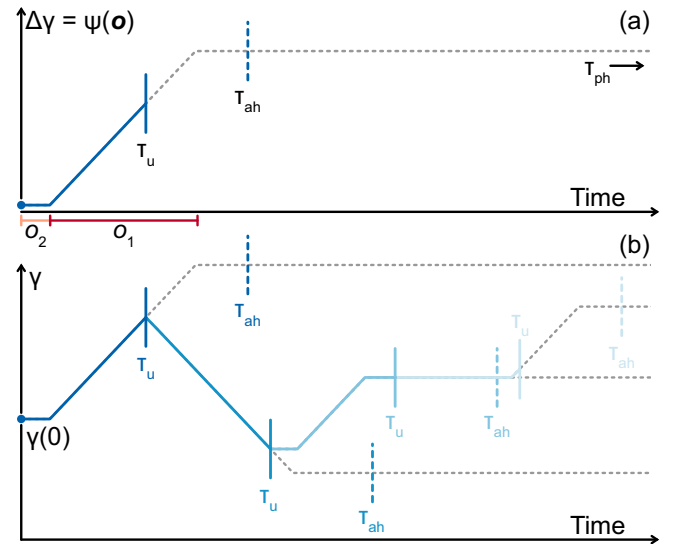


FIGURE 5 | Example of ψ and the resulting displacement $\Delta\gamma$ from the current orientation (a). The percentage of the action horizon τ_{ah} that is spent yawing is determined by o_1 , while o_2 determines the offset at which the angle change starts. At τ_u , the optimization is redone based on the updated model state. An example of how a turbine orientation might change over successive optimizations is given in (b).

removes any effect of ϕ_2 onto the resulting trajectory. Similarly, for $\phi_1 = 0$ or $= 1$, the entire duration of t_{AH} is spent yawing, which also nullifies the impact of ϕ_2 . This impacts the optimization landscape and introduces regions with no gradient sensitivity towards ϕ_2 . This affects the optimizer choice in Section 2.3.5.

2.3.3 | Implicit Misalignment Limitation

The methodology introduced in Section 2.3.2 limits the possible solution space of the yaw angle trajectories based on the yaw rate limit. However, it does not enforce a limit on what the resulting yaw angle might be. There are different ways of achieving a yaw misalignment limit, one would be to set it as a constraint. This is a linear constraint if the yaw misalignment is the input to the optimization problem. However, if wind direction changes are present, it is beneficial to optimize the turbine orientation instead of the misalignment, as the misalignment is a product of the uncontrollable wind direction and the turbine orientation. This switch, along with the basis function approach discussed in Section 2.3.2, leads to a nonlinear constraint of the yaw angle misalignment. This creates an additional layer to implement and can increase the complexity of the optimization problem. We instead chose to manipulate the power calculation in such a way that large yaw angles become much less attractive. This removes the constraint and makes an implicit part of the cost function. To this end, we formulate a weighting function for the power generated:

$$w_\gamma(\gamma, \gamma_{\max}, \gamma_{\min}) = \left[\frac{1}{2} \tanh(50[-\gamma + \gamma_{\max}]) + \frac{1}{2} \right] \cdot \dots \left[-\frac{1}{2} \tanh(50[-\gamma + \gamma_{\min}]) + \frac{1}{2} \right]. \quad (5)$$

The weighting function reduces the power generated by 50% at the limit yaw angles and smoothly transitions from the unmodified part of the power curve to the lowered part outside of the limits. The limits were chosen as $\gamma_{\max} = -\gamma_{\min} = 33$ deg. Note that Equation (5) reduces the power for yaw angles that are still within but close to the bounds. As a result, yaw angles > 33 deg are already unattractive to the optimizer.

2.3.4 | Wind Farm Decomposition

To reduce the computational cost of the optimization and to ensure a faster convergence, a decomposition of the wind farm in smaller subsets is useful [47]. This ensures that turbines that are unaffected by other turbines, nor affect any, can optimize for greedy behavior. Turbines that do affect one another are optimized together. In this work, we decompose the wind farm based on the current wind direction at each turbine's location. Each turbine calculates which other turbines are within ± 2 D crosswind distance and downstream based on the current wind direction, see Figure 6. These then become part of the optimization group related to this turbine. We then recursively determine which turbines affect one another. This creates a directed graph, similar to the model principles proposed by Starke et al. [48]. Each group is then optimized together, meaning that only those turbines in the group are simulated. This effectively reduces the ten-turbine wind farm to smaller farms with one to




FIGURE 6 | Demonstration of the wind farm decomposition into smaller subsets, based on the ± 2 D crosswind distance selection. Turbines that are grouped together share the same color.

six members. This wind farm decomposition is done with every optimization step given the current conditions and, therefore, changes over time.

2.3.5 | Optimization Algorithm

The presented framework uses a particle swarm optimization to solve the cost function, which is part of the evolutionary algorithms [49–51]. The code uses the implementation by MathWorks [52]. The optimizer was chosen for the main reason that the optimization landscape is not convex. This stems from the split nature of wake steering, as steering in both directions might yield an improvement while one is favorable. Another contribution is the insensitivity of yaw angle basis function for certain parameter combinations, see Section 2.3.2. A total of 100 particles is used for all optimizations. The stopping criterion is either a maximum number of four consecutive iterations with no cost function improvement or 20 iterations. These settings were chosen after also testing 2, 10 and 40 iterations. This design further leverages the code's capability to run multiple FLORIDyn instances in parallel, similar to the EnKF. Gradient-based methods have been tested during the development of the closed-loop controllers but have been outperformed by non-gradient-based methods. The development of a dedicated optimization strategy like the Serial Refine method [53] lies outside of the scope of this paper.

2.4 | Reference Controllers

We use two types of reference dead-band lookup-table (LuT) yaw steering controllers, one that aims for 0 deg yaw misalignment and one that does implement yaw steering. Their design is based on the work of Kanev [54] and has further been investigated in Becker et al. [55]. The control framework consists of two parts: (i) the dead-band filter for the wind direction and (ii) the yaw controller. The dead-band reads in the wind direction measurement $\phi(k)$ at time step k and updates its wind speed estimate $\hat{\phi}$ based on the difference between the two values:

$$\hat{\phi}(k) = \begin{cases} \phi(k) & \text{if } |\phi(k) - \hat{\phi}(k-1)| > \phi_{\lim}, \\ \phi(k) & \text{if } k_i \Delta t \left| \sum_{l=\tau}^{k-1} \phi(l) - \hat{\phi}(k-1) \right| > \phi_{\lim}, \\ \hat{\phi}(k-1) & \text{otherwise.} \end{cases} \quad (6)$$

Equation (6) also consists of a second update law based on the integrated difference between past measurements of φ and $\hat{\varphi}$ as the last time step τ at which $\hat{\varphi}$ was updated. The resulting $\hat{\varphi}(k)$ is then used in the LuT, denoted by f_{LuT} , which returns the yaw set points:

$$\gamma^*(k) = f_{\text{LuT}}(\hat{\varphi}(k)) \quad (7)$$

$$\gamma(k) = \gamma(k-1) + \text{sign}(\gamma^*(k) - \gamma(k-1)) \cdot \dots \min(|\gamma^*(k) - \gamma(k-1)|, \Delta t \delta_\gamma) \quad (8)$$

The turbine yaw angle $\gamma(k)$ is then updated based on Equation (8): It is either set to γ^* if the angle can be reached in Δt , or changes with the maximum rate of yawing δ_γ for Δt . For the baseline controller, we set $\Delta\varphi = 2$ deg and $k_i = 0.1$; for the LuT controller, we test $\Delta\varphi \in [2, 4]$ deg and $k_i = 0.1$. The LuT is calculated with FLORIDyn in steady-state conditions using the same particle swarm optimization as discussed in Section 2.3.5. This way all controllers use the same basis to make decisions. Figure 7 depicts the LuT for all turbines and wind directions. A characteristic of this LuT is that downstream turbines show very small misalignment angles to be optimal. This allows the downstream turbine to move slightly out of the way of a partial wake overlap. Toolboxes like FLORIS [56] avoid this behavior by setting downstream turbine misalignment angles to 0 deg by default. Because the CLC controllers tested in this work will also exhibit this behavior, it remains part of the LuT.

3 | Simulation Methods

Section 3.1 discusses the high-fidelity model used as surrogate for the real wind farm, along with the wind direction case. This is followed by Section 3.2, which further specifies the tested controllers and used measurements.

3.1 | “True Wind Farm” Setup

The experiments are performed in closed-loop with the LES code SOWFA as the true wind farm surrogate [2]. The wind farm is situated in a $5 \times 5 \times 1$ km domain, resolved in $300 \times 300 \times 100$

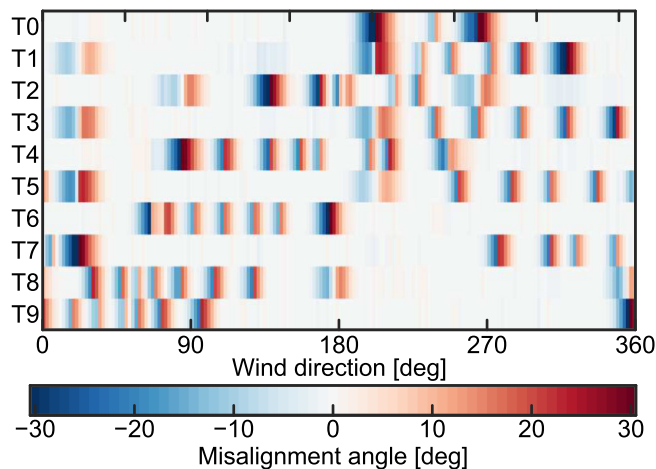


FIGURE 7 | Look-up table generated with the internal FLORIDyn wake model for the reference controllers.

cells with an average 8 ms^{-1} free wind speed at hub height and a 0.5 s time step. All simulations are done with a neutral turbulent precursor developed for $3 \cdot 10^4$ s with a surface roughness length of $2 \cdot 10^{-4}$ m. This leads to a low turbulence intensity case and more pronounced turbine wakes compared with a high turbulence flow field. Based on early flow estimates, the TI in FLORIDyn was set to 5.4%, the precursor turbulence intensity at hub height became $I_{0, u,v,w} \approx 4\%$. The shear resulting from the surface roughness is used to calculate the power-law shear coefficient α based on the mean wind speed magnitude. The ideal shear coefficient to describe the wind speed across the rotor plane lies between 0.07 and 0.083. Based on initial precursor values, α_s was set to 0.071. The mean precursor wind speed magnitude, as well as the FLORIDyn wind speed profile, are depicted in Figure 8b. The veer across the rotor plane is less than 2 deg in the LES and is neglected in FLORIDyn.

The closed-loop approach is tested in a 3-h long case, which is based on data collected by a vertical LiDAR at the Hollandse Kust Noord (HKN) site on the 28th of March 2023 [57]. During the measurement campaign (2019 - ongoing), wind speed and direction are recorded at different heights, with values available roughly every 20 s. First, the data were segmented into parts with sufficient data points. The values at 108 m and 133 m were used to resample and interpolate the wind direction at the turbine hub height of 119 m at a regular 20 s sampling rate. The data were then low-pass filtered with a Butterworth filter with a cut-off frequency of 1/600 Hz, equivalent to van den Broek et al. [35].




FIGURE 8 | Precursor profile of the wind speed in the precursor and in the FLORIDyn simulation.

The resulting sections were then investigated for completeness and interesting wind direction ramp events. For this work, one sub-section of 3 h is chosen. Figure 9 depicts the original data and how it was prepared for the closed-loop tests. The wind farm layout and dimensions are depicted in Figure 10, along with the rotated case domain to simulate varying inflow conditions. The same setup has been used to conduct the LuT controller study presented in Becker et al. [55].

We use OpenFOAM to implement spatial uniform wind direction changes following the predefined time series of wind directions. Within the LES domain, the southwest planes are used as inflow planes, while the northeast planes are outflow planes. The farm layout and initial wind direction are rotated to fit the 225 deg inflow of the turbulent precursor. The rotated layout is centered in the domain to balance the distance to the in- and outflow planes, see Figure 10. The ten DTU 10 MW turbines [58] are arranged to copy a subset of the HKN wind farm and are modeled as actuator discs (ADM). Turbines modeled as ADMs on a coarser grid tend to overestimate the power generated [59, 60], which is also an issue with this setup. Therefore, with the exception of Figure 14, we focus on normalized power and energy quantities.

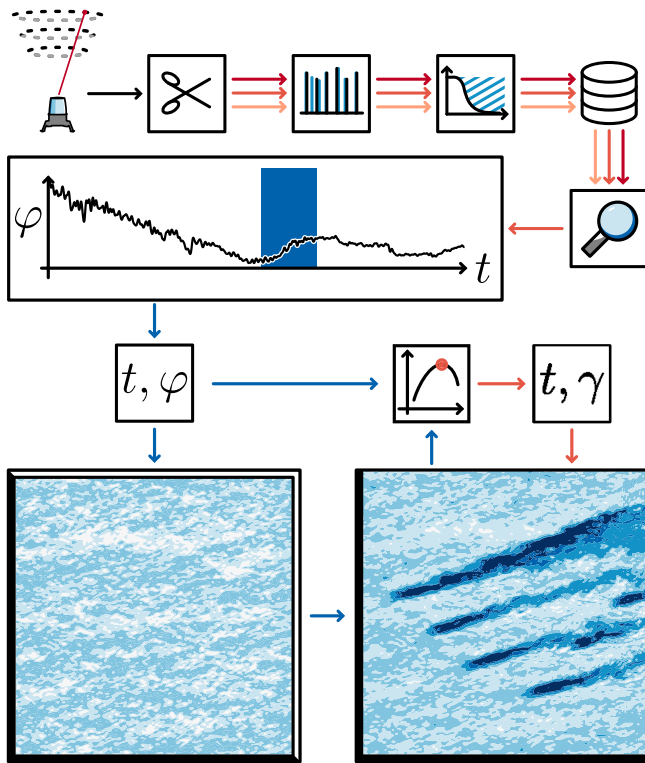


FIGURE 9 | Schematic of the data flow for the verification and test simulations: The initial LiDAR dataset was segmented into time series of sufficient completeness, followed by resampling, interpolation, and zero-phase low-pass filtering. The stored data has then been investigated for time series of long, uninterrupted wind direction trajectories with wake interaction. Out of a 24 h period, one 3 h segment has been chosen to be simulated in a closed loop. The simulation has a precursor with the wind direction changes from the LiDAR data, which is then used to drive the simulation with the wind farm. During the simulation, measurement data are sent from the LES to a controller, which then continuously updates a table with the yaw angle set points for all turbines.

The LES environment is extended by a wind farm-wide controller that receives turbine measurements and can provide set points during the simulation runtime. The measurements received from the LES include quantities like generator power and rotor speed. Wind direction measurements are provided in one of two ways: Either directly from the data that was used to create the precursor or based on probes at hub height at the turbine locations. The former results in a noise-free measurement of the underlying wind direction, while the latter is subject to ambient and turbine-added turbulence.

3.2 | Tested Controllers and Measurements

This section specifies the controller configurations tested in tandem with the LES and how they acquire their inputs.

3.2.1 | Model predictive controllers

All model predictive controllers use the Particle Swarm Optimizer coupled to the 2 degrees-of-freedom baseline function to derive the yaw trajectories; see Sections 2.3.1 and 2.3.2. Two controllers maximize the energy over 500 s and 1000 s. This is done by evaluating the cost function in Equation (2). Another controller uses the shifted turbine power signals in time to synchronize control actions with their effect on downstream turbines; see Section 2.3.1. This leads to a varying prediction horizon based on the turbines involved in the current optimization problem. Typical values are between 100 s and 600 s. The controllers are summarized in Table 1. All controllers update the optimal yaw set points every 60 s and use an action horizon of 100 s. This allows for a ± 30 deg orientation change based on the maximum yawing rate of $r_y = 0.3 \text{ deg s}^{-1}$.

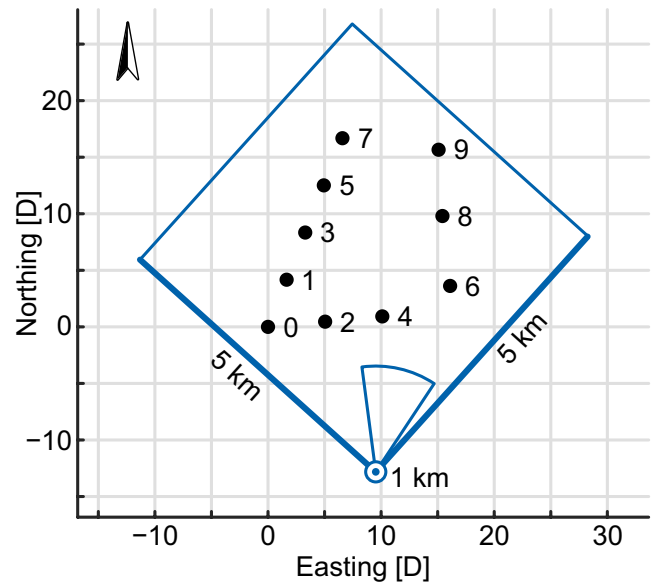


FIGURE 10 | Wind farm layout of the investigated case study. The wind farm is oriented as it is in the real world, while the $5 \times 5 \times 1 \text{ km}$ domain is rotated to fit the inflow direction. The inflow planes are marked by wider lines. The triangle marks the wind direction range of the case. The wind direction time series is chosen such that initially T6, T8, and T9 interact, and later T0, T1, T3, T5, and T7.

3.2.2 | Reference controllers

The reference yaw-steering and baseline controllers are based on the dead-band behavior described in Section 2.4. The used threshold and integral gain are given in Table 2. In the noise-free environment, the controllers are tested with a sampling time of 5 s to update the set points. With the disturbed wind direction measurements the LuT controllers are updated once every minute based on the past 1-min average.

3.2.2.1 | Measurements. The controllers depend on measurements of wind speed and direction to provide adequate yaw angles. Table 3 lists the different ways data are provided to the controllers. This concerns the background wind speed u_∞ , the wind direction φ , the sampling time step, and the averaging time of the φ measurement. The noise-free data relates to the filtered LiDAR data that was used to drive the precursor. It can, therefore, be considered as an ideal, noise-free signal of the background flow. The same holds for u_∞ , which is provided as a constant value to the reference controllers. The EnKF for the closed-loop controllers integrates new measurements every 15 s, but a new control decision is taken every 60 s. The LuT controllers do not have a state and rather act based on the current measurement, hence the lower sampling

time. Figure 11 showcases an example of the measurements. The black source data comes from the cleaned and zero-phase low-pass filtered field data; see Figure 9. The gray probe data are recorded in the LES at the rotor center of the turbine, which is, in this example, turbine T7. This was done to mimic a much-simplified version of the measurements a wind vane on a turbine might record. The plot also shows how the noise is reduced by the use of the past 1-min averaged data instead of the raw probe data. The probe data are characterized by higher noise levels and biases during waked conditions, which poses challenges for the state estimation and the dead-band controllers. Because the wake locations are unique to every simulation, also the LuT controllers have to run online and their control actions can not be pre-computed.

4 | Simulation Results

The simulation results first investigate the wake steering behavior in Section 4.1. Section 4.2 then compares the performance of the turbines throughout the case study, followed by the farm-wide performance in Section 4.3.

4.1 | Wake Steering

Figure 12a depicts the orientation of turbine T4 during the last half hour of the case study. During this time, the wind direction aligns T4 with T8, as indicated by the arrow. All controllers do engage in wake steering to avoid waking turbine T8 at 10.3 D distance. However, the magnitude of the misalignment differs. Between minutes 160 and 175, the wind direction varies only marginally and does not cross the line between T4 and T8. The LuT controllers engage in the largest misalignment angles, and the Shifted CLC controller acts similarly. The CLC 1000 s controller exhibits a smaller yaw angle, and lastly, the CLC 500 s controller shows little-to-no misalignment.

TABLE 1 | Selection of closed-loop model predictive controllers.

Controller	Cost function	τ_{ph}
PSO, MR	Energy	$0.5 \cdot 10^4$ s
PSO, MR	Energy	$1 \cdot 10^4$ s
PSO, MR	Shifted energy	Varying

Note: PSO refers to the Particle Swarm Optimization, MR to the maximum yawing rate basis function. All controllers have an action horizon of $\tau_{ph} = 100$ s and no preview.

TABLE 2 | List of reference controller settings.

Steady-state model	φ_{lim}	Optimization
FLORIDyn internal	2 deg	PSO offline
FLORIDyn internal	4 deg	PSO offline
Baseline	2 deg	—

Note: All reference controllers have a $k_t = 0.01$ s⁻¹.

TABLE 3 | Ways for the state estimator and controller to receive data from the simulation.

	u_∞	p	φ	Δt	average time φ
Mode 1	—	LES	noise free	15 s	0 s
Mode 2	—	LES	LES	15 s	60 s
Mode 3	given	—	noise free	5 s	0 s
Mode 4	given	—	LES	60 s	60 s

Note: Modes 1 and 2 are relevant to the ensemble Kalman filter, and Modes 3 and 4 are relevant to the reference LuT controllers.




FIGURE 11 | Wind direction data recorded at the location of turbine T7 during the simulation. The source data are based on low-pass filtered field measurements, which is used to drive the LES. The probe data are calculated from the u and v wind speed components at the turbine rotor center. Orange indicates the 1-min average values of the probe data. The probe data becomes more noisy and biased as the turbine is waked by other turbines, here from 1.5 to 2.5 h. Note that the wind direction is given with respect to the LES domain.




FIGURE 12 | The orientation of turbine T4 over time as a result of the noise-free wind direction measurements is given in (a). The arrow indicates the wind direction in which turbine T8 is located at 10.3 D distance. The resulting 10-min energy efficiency of T8 is given in (b), followed by the combined performance of T4 and T8 in (c).

Figure 12b then shows how the 10-min energy of T8 reacts to the yaw steering efforts of T4: The more aggressive yaw angles by the LuT controllers indeed leads to a better efficiency at T8, increasing its generation by +56%. The gain of T8, however, comes at the cost of misaligning T4 for a long time. Figure 12c shows the combined efficiency of T4 and T8. The data show that the period of outperforming the baseline is preceded by a period of underperformance. Based on the shallow misalignment angles, the closed-loop controllers overcome this period earlier than the LuT controllers, confirming the analysis done in Section 2.3.1 and that the CLC controllers are working as intended. They do engage in wake steering control and, based on their design, in a more or less aggressive manner. The presented example also shows that the controllers take turbines across a longer distance into account, something with which, for example, free vortex particle models of the wake can struggle with [35].

4.2 | Turbine Performance

Figure 13 shows the normalized energy gain separated for each turbine. Given the wind direction time series and




FIGURE 13 | Turbine individual difference between the energy generated in the controlled case and the baseline, normalized by the total baseline energy. The data in (a) relate to the controllers with noise-free wind direction measurements, and the data in (b) relate to the disturbed measurements. The data are normalized by the respective baseline.

the farm layout, T0, T2, T4, and T6 are upstream turbines, while T7, T8, and T9 are the most downstream turbines, see Figure 10.

The data show that for all controllers, the upstream turbines sacrifice energy by yawing, which is then recouped by the downstream turbines. There are differences in the magnitude: the CLC 500 s controller amplifies the energy generation of T1, T3, and T5 more than the generation of T7, T8, and T9. It also shows a decreased investment for turbines T0, T2, T4, and T6. This is consistent with the shorter prediction horizon length and with the analysis done in Section 4.1. The CLC 1000 s shows more committed control actions that further lower the energy generation of the upstream turbines but also result in larger returns for the downstream turbines. The same holds true for the shifted CLC controller, as well as for the LuT controllers. An advantage that the LuT controllers have over the CLC controllers in the noise-free environment is that they are able to engage T5 and T7 consistently.

With disturbed wind direction measurements, the performance of all controllers decreases. The LuT controllers especially sacrifice performance with T3, T5, T7, and T9. This can be a sign that the long-distance wake interactions fail as the wind direction measurement becomes more uncertain. The CLC controllers also sacrifice performance, mainly with T1, T3, and T5. We attribute this to the lowered yaw investment by T0 due to the more uncertain wind direction.

4.3 | Farm Level Performance

To investigate the performance of the controllers on a farm level, we compare the wind energy as the power generated over a sliding window integral of 10 min. Figure 14 depicts the absolute and relative energy generated by the farm throughout the case study. The absolute data shows the magnitude at which the energy is generated. We note that the largest reductions in energy appear around the 1.5 h mark, which relates to a brief back-and-forth shift in the wind direction. This crosses the 5-turbine line T0, T1, T3, T5, and T7. During this event, the controllers return their largest gains but also losses. The losses are a product of wake steering as a method: While the turbines are slightly misaligned, the wakes are redirected to one side of the downstream turbine. If the wind direction changes, the wake can only be further redirected up to a certain point at which it is worth redirecting the wake to the other side. During this switch, the baseline wakes have likely already arrived on the new side, which means that the baseline momentarily generates more than the control strategy.

Overall, all tested controllers return consistent gains in the noise-free case, but also in the disturbed case. Here, the performance of the controllers is generally weaker due to the worse sensor data. Qualitatively, the CLC 500 s controller is the closest to the BL performance and is followed by The CLC 1000 s controller. The remaining three controllers depict a similar performance.

To quantify the performance, the data of Figure 14c,d are summarized as a box plot in Figure 15, with the difference that all data are normalized with the noise-free baseline performance.

The two LuT controllers and the time-shifted controller still depict similar performance during noise-free conditions: The LuT 2 deg and 4 deg controllers have a median efficiency of +4.2% and +4.1%, respectively, the CLC shifted controller has a median of +4.0%. With disturbed wind direction measurements, the performance drops to a median of +1.7% and +1.4% for the LuT 2 deg and 4 deg controllers, while the CLC shifted controller reduces its median to 2.1%. We can conclude that the CLC is, therefore, more robust to sensor noise. This is mainly due to the EnKF, which by design assumes noise to be part of the sensor signal. With noise-free data, this leads to a disadvantage as the measurement is not considered to be fully “trustworthy”, but for disturbed measurements, this leads to a better estimate. This effect is also visible with the other two CLC controllers: The CLC 500 s controller drops from a median of +2.2% to +1.0%, the CLC 1000 s controller from +2.5% to +1.5%, which is a less significant decrease than the one of the LuT controllers.

A mostly neglected aspect of the cost function is the cost of actuation. It is only captured implicitly as a loss of the power generated by the actuated turbine. Actuation costs can include how much and how often the turbines yaw, how much time they spent in misalignment and how their inflow profile looks like. Recent work has suggested ways to create data-driven ways to estimate the loads on a turbine in a surrogate model [61, 62], but these are not yet included in the FLORIDyn model used in this work. We therefore resort to the yaw travel as quantity of interest.

Figure 16 shows the overall energy efficiency and normalized yaw travel. A common characteristic of the controllers under noise-free wind direction measurements is that all yaw controllers exhibit an approximately $2\times$ to $3\times$ larger amount of yaw

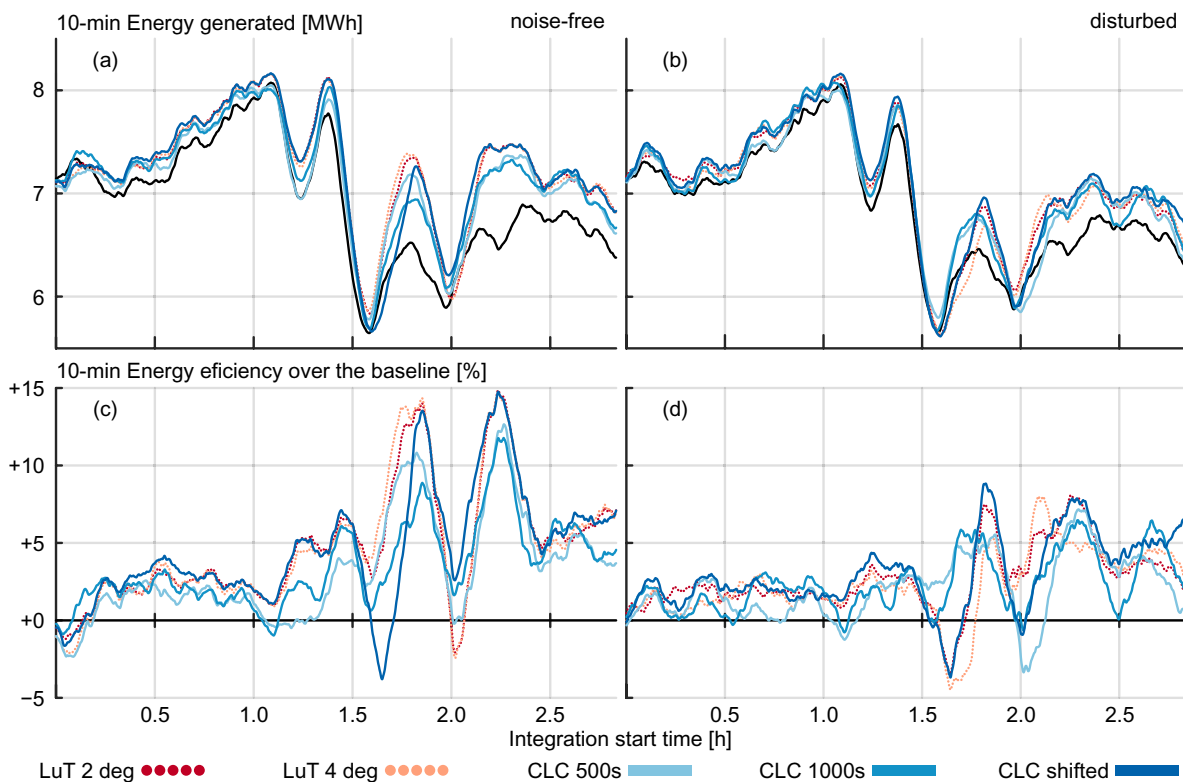


FIGURE 14 | Absolute (a, b) and relative (c, d) energy generated in a sliding time window of 10 min. The data in (a) and (c) are based on the noise-free wind direction measurements, and the data in (b) and (d) are based on the disturbed measurements.




FIGURE 15 | Box plot of the 10-min wind farm energy efficiency. The data are normalized with the noise-free baseline. Each pair depicts the noise-free performance on the left and the performance with disturbed wind direction measurements on the right. The median values from the left to the right are BL* 2 deg -0.5% , LuT 2 deg $+4.2\%$, LuT* 2 deg $+1.7\%$, LuT 4 deg $+4.1\%$, LuT* 4 deg $+1.4\%$, CLC 500 s $+2.2\%$, CLC* 500 s $+1.2\%$, CLC 1000 s $+2.5\%$, CLC* 1000 s $+1.5\%$, CLC shifted $+4.0\%$, CLC* shifted $+2.1\%$. The disturbed direction simulations are marked with *.

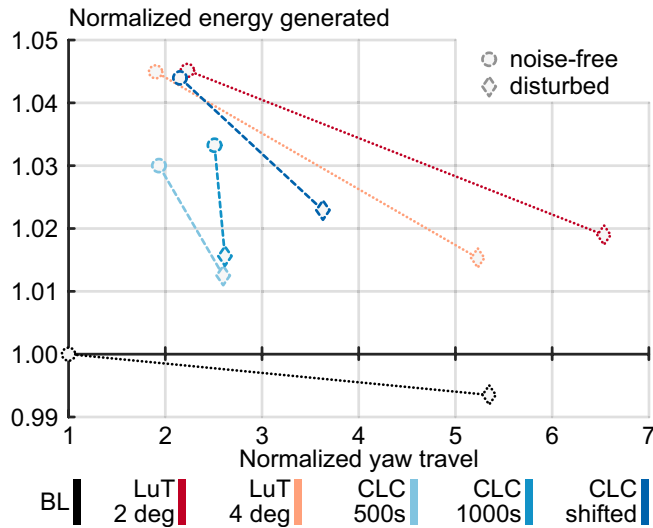


FIGURE 16 | Energy and yaw travel of the six controllers with both noise-free and disturbed wind direction measurements. The data are normalized with the noise-free baseline controller.

travel for a 3% to 4.5% gain in energy. This ratio changes as the wind direction measurements become disturbed. All controllers are affected by the disturbed measurements, however at a different scale. Similarly to the trend observed with Figure 13, the LuT are significantly more affected by poor data quality than the CLC controllers: All LuT controllers, including the baseline, exhibit a yaw travel amount that is ≈ 5 to $7 \times$ higher than the noise-free baseline, while the CLC controllers remain between ≈ 2.5 to $4 \times$ higher values.

4.4 | Computational Performance

The closed-loop algorithm has two steps: (i) the state estimation and (ii) the model predictive optimization. In the case study, the LES is paused for both the state estimation and the optimization. While unrealistic, this allows for more leeway and exploration of the methodology before imposing more challenging conditions.

The state estimation is real-time compatible and roughly requires 6 s of computational time for simulation, ensemble

combination, and state correction, applied every 15 s. The optimization, however, is not realtime applicable for larger numbers of turbines. Currently, groups of ≥ 3 turbines take longer than the allocated 60 s update time. These numbers are obtained in Matlab for 40 threads.

Future work should investigate the design of a dedicated optimization strategy similar to the serial refine approach [53]. And to reduce the cost per simulation. In its current form, also the FLORIDyn model used in the optimization uses the spatiotemporal weighting of the OPs. This adds a considerable cost to the model, and while it is necessary for the EnKF, it might not be for the optimization.

5 | Conclusions

This paper introduces a novel closed-loop control framework to maximize the energy generated by a wind farm under time-varying inflow conditions based on the dynamic wake model FLORIDyn. The observed case-study results show that the framework can lead up to a median energy gain over 10 min of $+4.0\%$ using a shifted cost function and $+2.2\%$ to $+2.5\%$ using an energy maximizing cost function. This, however, falls short behind the tested LuT controllers with gains of $+4.1\%$ and $+4.2\%$. These results are obtained with noise-free wind direction measurements. If disturbed measurements from the LES are used, the performance of all controllers decreases while the yaw travel increases. However, the closed-loop controllers are less affected by this change, highlighting their robustness due to the Ensemble Kalman Filter and cost-function design. Their median 10-min energy gain reduces to $+2.1\%$, $+1.5\%$, and $+1.2\%$, while the LuT performance decreases to $+1.7\%$ and $+1.4\%$. Future research should explore the scalability of this approach to larger wind farms, as well as in heterogeneous flow conditions.

Author Contributions

M.B.: conceptualization, methodology, software, validation, investigation, writing – original draft, visualization. **M.J.v.d.B.:** conceptualization, writing – review. **D.A.:** conceptualization, supervision. **J.W.v.W.:** writing – review and editing, conceptualization, supervision, resources, funding acquisition.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

The FLORIDyn framework is available with the DOI 10.4121/038777f7-f497-494f-9f61-85d90a00074a. The repository includes the model, the state estimation, and the closed-loop control files. The simulation data to recreate plots and accompanying information given with the DOI 10.4121/63fa3d6c-6bd2-4efe-9548-ddcaf1bfbeb9 The version of the LES solver SOWFA [2] used in this work is based on the GitHub repository <https://github.com/TUdelft-DataDrivenControl/SOWFA>.

Peer Review

The peer review history for this article is available at <https://publons.com/publon/10.1002/we.70044>.

References

1. UNECE, *Carbon Neutrality in the UNECE Region: Integrated Life-Cycle Assessment of Electricity Sources* (Geneva: UNITED NATIONS ECONOMIC COMMISSION FOR EUROPE, 2022).
2. M. J. Churchfield, S. Lee, J. Michalakos, and P. J. Moriarty, "A Numerical Study of the Effects of Atmospheric and Wake Turbulence on Wind Turbine Dynamics," *Journal of Turbulence* 13 (2012): N14, <https://doi.org/10.1080/14685248.2012.668191>.
3. P. Chatelain, S. Backaert, G. Winckelmans, and S. Kern, "Large Eddy Simulation of Wind Turbine Wakes," *Flow, Turbulence and Combustion* 91, no. 3 (2013): 587–605, <https://doi.org/10.1007/s10494-013-9474-8>.
4. A. Rott, S. Boersma, J. W. van Wingerden, and M. Kühn, "Dynamic Flow Model for Real-Time Application in Wind Farm Control," *Journal of Physics: Conference Series* 854 (2017): 012039, <https://doi.org/10.1088/1742-6596/854/1/012039>.
5. S. Boersma, B. Doekemeijer, M. Vali, J. Meyers, and J. W. van Wingerden, "A Control-Oriented Dynamic Wind Farm Model: WFSim," *Wind Energy Science* 3, no. 1 (2018): 75–95, <https://doi.org/10.5194/wes-3-75-2018>.
6. M. J. van den Broek and J. W. van Wingerden, "Dynamic Flow Modelling for Model-Predictive Wind Farm Control," *Journal of Physics: Conference Series* 1618 (2020): 022023.
7. M. J. van den Broek, B. Sanderse, and J. W. van Wingerden, "Flow Modelling for Wind Farm Control: 2D vs. 3D," *Journal of Physics: Conference Series* 2265, no. 3 (2022a): 032086, <https://doi.org/10.1088/1742-6596/2265/3/032086>.
8. Larsen, G.C. (Ed.) 2007. "Dynamic Wake Meandering Modeling". No. 1607 in Risø R, Report. Roskilde: Risø National Laboratory.
9. Jonkman, J.M., Annoni, J., Hayman, G., Jonkman, B. & Purkayastha, A 2017. Development of FAST.Farm: A New Multi-Physics Engineering Tool for Wind-Farm Design and Analysis. In: *35th Wind Energy Symposium*, Grapevine, Texas: American Institute of Aeronautics and Astronautics.
10. Liew, J. 2022. "Dynamic Modelling of Wind Farms for Closed-Loop Control". Ph.D. thesis, Technical University of Denmark.
11. Marten, D. 2020. "QBlade: A Modern Tool for the Aeroelastic Simulation of Wind Turbines".
12. M. J. van den Broek, D. De Tavernier, B. Sanderse, and J. W. van Wingerden, "Adjoint Optimisation for Wind Farm Flow Control With a Free-Vortex Wake Model," *Renewable Energy* 201 (2022b): 752–765, <https://doi.org/10.1016/j.renene.2022.10.120>.
13. M. Bastankhah and F. Porté-Agel, "Experimental and Theoretical Study of Wind Turbine Wakes in Yawed Conditions," *Journal of Fluid Mechanics* 806 (2016): 506–541, <https://doi.org/10.1017/jfm.2016.595>.
14. P. M. O. Gebraad and J. W. van Wingerden, "A Control-Oriented Dynamic Model for Wakes in Wind Plants," *Journal of Physics: Conference Series* 524 (2014): 012186, <https://doi.org/10.1088/1742-6596/524/1/012186>.
15. M. Becker, B. Ritter, B. Doekemeijer, et al., "The Revised FLORIDyn Model: Implementation of Heterogeneous Flow and the Gaussian Wake," *Wind Energy Science* 7, no. 6 (2022): 2163–2179, <https://doi.org/10.5194/wes-7-2163-2022>.
16. M. Becker, D. Allaerts, and J. W. van Wingerden, "FLORIDyn - A Dynamic and Flexible Framework for Real-Time Wind Farm Control," *Journal of Physics: Conference Series* 2265, no. 3 (2022b): 032103, <https://doi.org/10.1088/1742-6596/2265/3/032103>.
17. M. Lejeune, M. Moens, and P. Chatelain, "A Meandering-Capturing Wake Model Coupled to Rotor-Based Flow-Sensing for Operational Wind Farm Flow Prediction," *Frontiers in Energy Research* 10 (2022): 884068, <https://doi.org/10.3389/fenrg.2022.884068>.
18. R. Braunbehrens, J. Schreiber, and C. L. Bottasso, "Application of an Open-Loop Dynamic Wake Model With High-Frequency SCADA Data," *Journal of Physics: Conference Series* 2265, no. 2 (2022): 022031, <https://doi.org/10.1088/1742-6596/2265/2/022031>.
19. J. Meyers, C. Bottasso, K. Dykes, et al., "Wind Farm Flow Control: Prospects and Challenges," *Wind Energy Science* 7, no. 6 (2022): 2271–2306, <https://doi.org/10.5194/wes-7-2271-2022>.
20. P. Fleming, J. Annoni, J. J. Shah, et al., "Field Test of Wake Steering at an Offshore Wind Farm," *Wind Energy Science Discussions* 2, no. 1 (2017): 229–239, <https://doi.org/10.5194/wes-2017-4>.
21. T. Duc, O. Coupiac, N. Girard, G. Giebel, and T. Göçmen, "Local Turbulence Parameterization Improves the Jensen Wake Model and Its Implementation for Power Optimization of an Operating Wind Farm," *Wind Energy Science* 4, no. 2 (2019): 287–302, <https://doi.org/10.5194/wes-4-287-2019>.
22. B. M. Doekemeijer, S. Kern, S. Maturu, et al., "Field Experiment for Open-Loop Yaw-Based Wake Steering at a Commercial Onshore Wind Farm in Italy," *Wind Energy Science* 6, no. 1 (2021): 159–176, <https://doi.org/10.5194/wes-6-159-2021>.
23. B. M. Doekemeijer, D. van der Hoek, and J. W. van Wingerden, "Closed-Loop Model-Based Wind Farm Control Using FLORIS Under Time-Varying Inflow Conditions," *Renewable Energy* 156 (2020): 719–730, <https://doi.org/10.1016/j.renene.2020.04.007>.
24. M. F. Howland, A. S. Ghate, S. K. Lele, and J. O. Dabiri, "Optimal Closed-Loop Wake Steering – Part 1: Conventionally Neutral Atmospheric Boundary Layer Conditions," *Wind Energy Science* 5, no. 4 (2020): 1315–1338, <https://doi.org/10.5194/wes-5-1315-2020>.
25. M. F. Howland, A. S. Ghate, J. B. Quesada, et al., "Optimal Closed-Loop Wake Steering – Part 2: Diurnal Cycle Atmospheric Boundary Layer Conditions," *Wind Energy Science* 7, no. 1 (2022): 345–365, <https://doi.org/10.5194/wes-7-345-2022>.
26. Sood, I. & J. Meyers 2023. Development and Validation of a Large Eddy Simulation Based Virtual Environment for Optimal Wind Farm Control. Ph.D. thesis, KU Leuven.
27. P. Bachant, P. Ireland, B. Burrows, et al., "Development and Validation of a Hybrid Data-Driven Model-Based Wake Steering Controller and Its Application at a Utility-Scale Wind Plant," *Wind Energy Science* 11 (2024): 2235–2259.
28. Gebraad, P.M.O., P. A. Fleming & J. W. van Wingerden 2015. "Wind Turbine Wake Estimation and Control Using FLORIDyn, a Control-Oriented Dynamic Wind Plant Model", <https://doi.org/10.1109/ACC.2015.7170978>.
29. J. P. Goit and J. Meyers, "Optimal Control of Energy Extraction in Wind-Farm Boundary Layers," *Journal of Fluid Mechanics* 768 (2015): 5–50, <https://doi.org/10.1017/jfm.2015.70>.
30. W. Munters and J. Meyers, "Dynamic Strategies for Yaw and Induction Control of Wind Farms Based on Large-Eddy Simulation and

- Optimization,” *Energies* 11, no. 1 (2018): 177, <https://doi.org/10.3390/en11010177>.
31. M. Becker, D. Allaerts, and J. W. Van Wingerden, “Time-Shifted Cost Function Design for More Efficient Dynamic Wind Farm Flow Control*,” in *2024 IEEE Conference on Control Technology and Applications (CCTA)* (Newcastle upon Tyne, United Kingdom: IEEE, 2024a): 440–445.
32. A. Sterle, C. A. Hans, and J. Raisch, “Model Predictive Control of Wakes for Wind Farm Power Tracking,” *Journal of Physics: Conference Series* 2767, no. 3 (2024): 032005, <https://doi.org/10.1088/1742-6596/2767/3/032005>.
33. J. Goit, W. Munters, and J. Meyers, “Optimal Coordinated Control of Power Extraction in LES of a Wind Farm With Entrance Effects,” *Energies* 9, no. 1 (2016): 29, <https://doi.org/10.3390/en9010029>.
34. N. Janssens and J. Meyers, “Towards Real-Time Optimal Control of Wind Farms Using Large-Eddy Simulations,” *Wind Energy Science* 9, no. 1 (2024): 65–95, <https://doi.org/10.5194/wes-9-65-2024>.
35. M. J. van den Broek, M. Becker, B. Sanderse, and J. W. van Wingerden, “Dynamic Wind Farm Flow Control Using Free-Vortex Wake Models,” *Wind Energy Science* 9, no. 3 (2024): 721–740, <https://doi.org/10.5194/wes-9-721-2024>.
36. B. A. M. Sengers, A. Rott, E. Simley, M. Sinner, G. Steinfeld, and M. Kühn, “Increased Power Gains From Wake Steering Control Using Preview Wind Direction Information,” *Wind Energy Science* 8, no. 11 (2023): 1693–1710, <https://doi.org/10.5194/wes-8-1693-2023>.
37. Simley, E., P. Fleming, J. King & M. W. Sinner 2021. “Steering Wind Farm Control With Preview Wind Direction Information.” In: *2021 American Control Conference (ACC)*, New Orleans, LA, USA: IEEE, pp. 1783–1789.
38. M. Becker, D. Allaerts, and J. W. van Wingerden, “Ensemble-Based Flow Field Estimation Using the Dynamic Wind Farm Model FLORIDyn,” *Energies* 15, no. 22 (2022a): 8589, <https://doi.org/10.3390/en15228589>.
39. R. Braunbehrens, S. Tamaro, and C. L. Bottasso, “Towards the Multi-Scale Kalman Filtering of Dynamic Wake Models: Observing Turbulent Fluctuations and Wake Meandering,” *Journal of Physics: Conference Series* 2505, no. 1 (2023): 012044, <https://doi.org/10.1088/1742-6596/2505/1/012044>.
40. J. Di Cave, R. Braunbehrens, J. Krause, A. Guilloiré, and C. L. Bottasso, “Closed-Loop Coupling of a Dynamic Wake Model With a Wind Inflow Estimator,” *Journal of Physics: Conference Series* 2767, no. 3 (2024): 032034, <https://doi.org/10.1088/1742-6596/2767/3/032034>.
41. J. Schreiber, C. L. Bottasso, and M. Bertelè, “Field Testing of a Local Wind Inflow Estimator and Wake Detector,” *Wind Energy Science* 5, no. 3 (2020): 867–884, <https://doi.org/10.5194/wes-5-867-2020>.
42. G. Evensen, “Sequential Data Assimilation With a Nonlinear Quasi-Geostrophic Model Using Monte Carlo Methods to Forecast Error Statistics,” *Journal of Geophysical Research: Oceans* 99, no. C5 (1994): 10143–10162, <https://doi.org/10.1029/94JC00572>.
43. G. Evensen, *Data Assimilation: The Ensemble Kalman Filter*, 2nd ed. (Dordrecht: Springer, 2009).
44. Doekemeijer, B.M., Boersma, S., Pao, L.Y. & van Wingerden, J.W. 2017. Ensemble Kalman Filtering for Wind Field Estimation in Wind Farms. In: *2017 American Control Conference (ACC)*, pp. 19–24.
45. M. G. Kim and P. H. Dalhoff, “Yaw Systems for Wind Turbines – Overview of Concepts, Current Challenges and Design Methods,” *Journal of Physics: Conference Series* 524 (2014): 012086, <https://doi.org/10.1088/1742-6596/524/1/012086>.
46. U. Ciri, M. A. Rotea, and S. Leonardi, “Model-Free Control of Wind Farms: A Comparative Study Between Individual and Coordinated Extremum Seeking,” *Renewable Energy* 113 (2017): 1033–1045, <https://doi.org/10.1016/j.renene.2017.06.065>.
47. F. Bernardoni, U. Ciri, M. A. Rotea, and S. Leonardi, “Identification of Wind Turbine Clusters for Effective Real Time Yaw Control Optimization,” *Journal of Renewable and Sustainable Energy* 13, no. 4 (2021): 043301, <https://doi.org/10.1063/5.0036640>.
48. G. M. Starke, C. Meneveau, J. R. King, and D. F. Gayme, “A Dynamic Model of Wind Turbine Yaw for Active Farm Control,” *Wind Energy* 27, no. 11 (2023): 1302–1318, <https://doi.org/10.1002/we.2884>.
49. IEEE Neural Networks Council, ed., *Proceedings/1995 IEEE International Conference on Neural Networks, the University of Western Australia, Perth, Western Australia, 27 November - 1 December 1995* (Piscataway, NJ: IEEE Service Center, 1995).
50. Pedersen, M.E.H. 2010. “Good Parameters for Particle Swarm Optimization”. Hvass Laboratories, Luxembourg. Technical Report no. HL1001.
51. E. Mezura-Montes and C. A. Coello Coello, “Constraint-Handling in Nature-Inspired Numerical Optimization: Past, Present and Future,” *Swarm and Evolutionary Computation* 1, no. 4 (2011): 173–194, <https://doi.org/10.1016/j.swevo.2011.10.001>.
52. MathWorks, *MATLAB - Optimization Toolbox Version 9.5* (MathWorks Inc., 2023).
53. P. A. Fleming, A. P. J. Stanley, C. J. Bay, et al., “Serial-Refine Method for Fast Wake-Steering Yaw Optimization,” *Journal of Physics: Conference Series* 2265, no. 3 (2022): 032109, <https://doi.org/10.1088/1742-6596/2265/3/032109>.
54. S. Kanev, “Dynamic Wake Steering and Its Impact on Wind Farm Power Production and Yaw Actuator Duty,” *Renewable Energy* 146 (2020): 9–15, <https://doi.org/10.1016/j.renene.2019.06.122>.
55. M. Becker, M. Lejeune, P. Chatelain, D. Allaerts, R. Mudafort, and J. W. Van Wingerden, “A Dynamic Open-Source Model to Investigate Wake Dynamics in Response to Wind Farm Flow Control Strategies,” *Wind Energy Science* 10, no. 6 (2024b): 1055–1075.
56. NREL 2023. “FLORIS. Version 3.4. *GitHub repository*”.
57. Knoop, S. 2019. “Wind - Lidar Wind Profiles Measured at North Sea Wind Farm TenneT Platforms 1 Second Raw Data”.
58. Bak, C., Zahle, F., Bitsche, R., Kim, T., Yde, A., Henriksen, L.C. et al. *The DTU 10-MW Reference Wind Turbine* (DTU Orbit, 2013).
59. Martinez, L., Leonardi, S., Churchfield, M. & Moriarty, P. 2012. A Comparison of Actuator Disk and Actuator Line Wind Turbine Models and Best Practices for Their Use. In: *50th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, Nashville, Tennessee: American Institute of Aeronautics and Astronautics.
60. C. R. Shapiro, D. F. Gayme, and C. Meneveau, “Filtered Actuator Disks: Theory and Application to Wind Turbine Models in Large Eddy Simulation,” *Wind Energy* 22, no. 10 (2019): 1414–1420, <https://doi.org/10.1002/we.2376>.
61. A. Guilloiré, F. Campagnolo, and C. L. Bottasso, “A Control-Oriented Load Surrogate Model Based on Sector-Averaged Inflow Quantities: Capturing Damage for Unwaked, Waked, Wake-Steering and Curtailed Wind Turbines,” *Journal of Physics: Conference Series* 2767, no. 3 (2024): 032019, <https://doi.org/10.1088/1742-6596/2767/3/032019>.
62. J. Liew, R. Riva, M. Friis-Møller, and T. Göçmen, “Wind Farm Control Optimisation Under Load Constraints Via Surrogate Modelling,” *Journal of Physics: Conference Series* 2767, no. 9 (2024): 092039, <https://doi.org/10.1088/1742-6596/2767/9/092039>.
63. V. V. Dighe, M. Becker, T. Göçmen, B. Sanderse, and J. W. van Wingerden, “Sensitivity Analysis and Bayesian Calibration of a Dynamic Wind Farm Control Model: FLORIDyn,” *Journal of Physics: Conference Series* 2265, no. 2 (2022): 022062, <https://doi.org/10.1088/1742-6596/2265/2/022062>.

Appendix A

Parameter Tuning

The total number of parameters is high in the proposed framework. Table A1 lists all parameters and constants: Six parameters to parametrize the OP weighting, one to limit the advection speed, eleven for the wake and turbine model, eight for the state estimation, and six for the optimization. Ideally, one would tune all parameters to fit a generalized scenario. The number of parameters to tune is high, which makes it difficult to replicate the related work [23, 24, 63]. We, therefore, resort to Latin Hypercube sampling of the remaining parameters and test the different combinations in a 20-min sub-set simulation of the final setup. For each simulation, a set of farm-wide error quantities is calculated:

- Mean bias in turbine power
- Mean absolute turbine power error
- Mean squared turbine power error
- Mean squared farm power error
- Farm power bias
- Mean squared turbine power error weighted by the predicted spread of the EnKF

This is complemented by turbine individual error quantities:

- Best possible power correlation based on a variable signal time-shift
- Time-shift at which the best correlation is achieved
- Turbine power bias
- Absolute power error
- EnKF weighted power error

Out of 600 simulations, the parameter combinations that perform the best are compared for similarities. The remaining parameters are tuned manually, informed by parameters that were deemed ideal in previous studies. During this process, $\eta = 1$ was locked to further reduce the number of parameters. However, the turbine model used in the simulations overestimates the power, see Section 3.1. As the power was used to correct the wind speed, the state estimation returns a $\approx 8\%$ too high wind speed. This was captured by the error quantities and corrected by adjusting the free parameters. The wake advection factor d is smaller, and the remaining parameters lead to a faster wake recovery. The resulting parameter set may work for this setup but may not be ideal for generalization.

TABLE A1 | Collection of all parameters and constants, their use, the investigated range, and selected value, as well as the component they belong to.

Parameter	Use	Range	Selected	Unit	Framework component
$\sigma_{w,dw,\varphi}$	weighting OP, Dir.	—	2.87	D	FLORIDyn
$\sigma_{w,cw,\varphi}$	weighting OP, Dir.	—	2.87	D	FLORIDyn
$\sigma_{w,t,\varphi}$	weighting OP, Dir.	—	50	s	FLORIDyn
$\sigma_{w,dw,u}$	weighting OP, Vel.	[0.3, 3.5]	0.6966	D	FLORIDyn
$\sigma_{w,cw,u}$	weighting OP, Vel.	[0.3, 2]	0.3570	D	FLORIDyn
$\sigma_{w,t,u}$	weighting OP, Vel.	[100, 300]	206.2331	s	FLORIDyn
d	Advection factor	[0.5, 1.0]	0.7396	—	FLORIDyn
η	Turbine efficiency	—	1.0	—	Turbine model
p_p	Yaw exponent (Power)	[1.7, 2.7]	2.2	—	Turbine model
α	Near wake length	—	2.32	—	Wake model
β	Near wake length	[0.07, 0.39]	0.154	—	Wake model
k_a	Wake expansion	[0.17, 0.92]	0.38371	—	Wake model
k_b	Wake expansion	—	0.003678	—	Wake model
k_{fa}	Added turbulence	—	0.73	—	Wake model
k_{fb}	Added turbulence	[0, 8]	0.8325	—	Wake model
k_{fc}	Added turbulence	[0, 0.5]	0.0325	—	Wake model
k_{fd}	Added turbulence	—	−0.32	—	Wake model
k_{TI}	TI spread	[1, 4]	3	—	Wake model
$l_{loc,\varphi}$	Localization, Dir.	—	2.8	D	EnKF
$\sigma_{\mu,\varphi}$	Process noise, Dir.	—	0*	deg	EnKF
$\sigma_{v,\varphi}$	Measurement noise, Dir.	—	3	deg	EnKF
$l_{loc,u}$	Localization, Vel.	[3.5, 8]	6.8011	D	EnKF
$\sigma_{\mu,u}$	Process noise, Vel.	[0.1, 0.5]	0.1991	m/s	EnKF
$\sigma_{v,p}$	Measurement noise, Pow.	[0.01, 0.3]	0.08	MW	EnKF
n_e	Number of ensembles	—	50	—	EnKF
k_{enkf}	Estimation sample time	—	3	Δt	EnKF
τ_{ah}	Action horizon	—	20	Δt	eMPC
r_γ	Yaw rate limit	—	0.3	deg s^{-1}	eMPC
$n_{\text{iter, max}}$	Max. optimization iterations	—	20	—	eMPC
k_{mpc}	Optimization sample time	—	12	Δt	eMPC
γ_{max}	Yaw limitation	—	33	deg	eMPC
γ_{min}	Yaw limitation	—	−33	deg	eMPC

Note: The turbine diameter $D = 178.4$ m and time step $\Delta t = 5$ s are used to normalize some parameters. Parameters without range have not been tuned.

*The process noise was unintentionally set to $\sigma_{\mu,\varphi} = 0$ deg, and should, for future experiments, be set to a higher value.