

A robust active power control algorithm to maximize wind farm power tracking margins in waked conditions

Simone Tamaro, Filippo Campagnolo, and Carlo L. Bottasso

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

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

Abstract. We present an active wind farm power control (APC) algorithm that operates wind turbines to maximize their power availability in order to robustly track a reference power signal in the presence of turbulent wind lulls. The operational setpoints of the wind turbines are optimized using an engineering flow model by combining induction control with wake steering. The latter has the goal of deflecting low-momentum wakes and increasing power margins. The algorithm also features a 5 proportional-integral closed loop inspired by the literature to correct potential errors deriving from the offline computation of the setpoints.

First, we demonstrate the new approach in steady-state conditions, showing how the availability of power is increased by mitigating wake interactions. We observe that our proposed method is particularly effective in conditions of strong wake impingement, occurring in scenarios of high power demand. Next, considering two wind farm layouts, we compare the performance 10 of the algorithm to three state-of-the-art reference APC formulations in unsteady scenarios using large-eddy simulations coupled with the actuator line method (LES-ALM). We show that the occurrence and treatment of local, temporary instances of power unavailability (*saturations*) dramatically affects power tracking accuracy. The proposed method yields superior power tracking due to the increased power margins that limit the occurrence of saturation events. Additionally, we show that this performance is achieved with reduced structural fatigue.

15 1 Introduction

As renewable energy sources occupy a larger portion of the electricity mix, they must become also capable of providing extra functionalities to the grid, beyond the pure generation of power (Aho et al., 2012; Ela et al., 2014). Among these extra operating modes, active power control (APC) is a strategy where generating assets are intentionally operated below their maximum output to satisfy operational constraints imposed by the transmission system operator (TSO).

20 In the context of wind energy, the APC problem is particularly challenging due to the dynamics of mesoscale weather phenomena, the two-way interaction of the atmospheric boundary layer with a wind farm, the complex development of the flow within the plant, and its local interaction with the aeroelastic behavior of wind turbines. The maximum available power that can be generated by a wind farm at any given time is strongly influenced by site and local turbine-specific ambient conditions, which change over time in uncertain and difficult-to-predict ways (van Kuik et al., 2016). As a result, sudden drops in wind

25 speed or inaccuracies in forecasts can result in inadequate power reserves, making it difficult to follow a given reference signal (Fleming et al., 2016).

In a wind farm, the situation is further complicated by the presence of low-momentum turbulent wakes, which lead to power losses and contribute to fatigue loading on frequently waked turbines (Vermeer et al., 2003; Lee et al., 2013; Guilloré et al., 2024). A variety of approaches have been suggested to reduce the impact of wake effects, such as induction, yaw control, and 30 mixing (Meyers et al., 2022). Yaw control, in particular, involves “steering” the wakes away from downstream turbines. Its effectiveness in boosting power has been demonstrated in numerical simulations (Jiménez et al., 2010; Fleming et al., 2014; Vollmer et al., 2016), wind tunnel experiments (Medici and Alfredsson, 2006; Campagnolo et al., 2016; Schottler et al., 2017), and real-world field trials (Fleming et al., 2019; Doekemeijer et al., 2021). In recent years, machine learning techniques have 35 been increasingly integrated into wake mitigation strategies, further enhancing their effectiveness (Meyers et al., 2022; Ally et al., 2025).

Different APC approaches have been described and tested in the literature. The most straightforward method uses an open-loop strategy, where each turbine is given a predefined power-share setpoint (Fleming et al., 2016). However, unsurprisingly, the absence of feedback reduces the power tracking accuracy, particularly under strong wake impingement conditions. Additionally, uniformly distributing power among turbines may result in suboptimal performance, as the local power availability 40 varies among different turbines due to wake effects.

Various authors have used model predictive control (MPC) for APC (Shapiro et al., 2017; Boersma et al., 2018). The main drawback of such methods lies in the need for a dynamic farm flow model, which increases complexity and computational cost. While such methods undoubtedly have their own merits, here we are interested in solutions that are closer to practical applicability and a more rapid uptake from industry.

45 Classical PID (proportional-integral-derivative) controllers are widely used in an extremely broad range of different industrial applications. As such, they have also been widely studied for wind farm APC (van Wingerden et al., 2017). While these methods lack the sophistication of MPC, they do not require a dynamic wind farm flow model and can offer quick response times with straightforward implementations.

The APC PI (proportional-integral) controller proposed by van Wingerden et al. (2017) operates on the tracking error, 50 adjusting the power demands to match a reference, and distributes power among the turbines in a predefined static manner. The approach incorporates gain scheduling, which is based on the proportion of wind turbines in saturation, defined as those whose available power is less than the demanded power.

Multiple authors have formulated APC methods that – beyond tracking a power signal – also try to control dynamic loads in order to reduce fatigue (Kanev et al., 2018; Vali et al., 2019; Silva et al., 2022). In particular, Vali et al. (2019) introduced a 55 nested PI loop to dynamically adjust the setpoints of the wind turbines, with the goal of equalizing their loading. So far, these PI-based methods have been applied only to induction control.

Although these methods perform well in many conditions, their performance may be significantly impacted by saturation events, caused by a temporary local lack of power (i.e., when the local *reserves* are exhausted). Saturations harm tracking

accuracy, and therefore their occurrence should be limited as much as possible. Saturations, or more in general power reserves, 60 are not explicitly accounted for nor monitored in the existing APC PI implementations, which is a gap in the existing literature.

The effects of wind variability can be mitigated by hybridizing wind farms with storage solutions (Sinner et al., 2023). While storage has a crucial role to play in the transition towards a large penetration of renewables, here we are interested in mitigating the effects of wind variability on APC performance without the addition of extra hardware, but simply through a better, more robust way of controlling wind farms.

65 The success of wake steering in power boosting control makes this technology a primary candidate to increase power reserves, thereby improving APC tracking robustness. In fact, some initial attempts in this direction have been recently presented by Starke et al. (2023); Oudich et al. (2023). These studies showed that the time scales required by wake redirection are compatible with secondary grid frequency regulation. However, these same studies lack a comprehensive modeling of misaligned conditions, which are significantly affected by curtailed operation (Cossu, 2021; Campagnolo et al., 2023; Heck et al., 2023).

70 This paper introduces a novel wind farm control algorithm designed to enhance power tracking accuracy under conditions of strong, persistent wakes, especially when the power demand approaches the maximum available power of the wind farm. The algorithm improves tracking performance by explicitly maximizing power reserves with the goal of mitigating the impact of wind lulls. This innovative approach integrates wake steering with induction control. The associated power losses in misaligned conditions are accounted for with a recent model by Tamaro et al. (2024a), which takes into account the way a turbine is 75 controlled as it yaws out of the wind. Wake steering is achieved using an open-loop, model-based optimal setpoint scheduler. This approach is based on the offline optimization of the control setpoints that, once stored in look-up tables, are interpolated at runtime. The resulting relatively simple online implementation, which is still based on a sophisticated offline optimization, has recently gained popularity in power-boosting wind farm control (Meyers et al., 2022), and it is being deployed in a growing number of industrial applications. Induction control is implemented through a rapid closed-loop corrector to enhance tracking 80 accuracy (Tamaro and Bottasso, 2023).

The methodology described in this paper has been preliminarily tested in a simulation environment by Tamaro and Bottasso (2023), and also demonstrated through experiments performed with scaled models in a boundary layer wind tunnel by Tamaro et al. (2024b). While both these already published articles reported promising results, they did not describe the formulation in detail, which was also still not completely mature; additionally, these papers only considered a limited number of cases, 85 and a comprehensive detailed comparison to state-of-the-art APC methods was missing. To address and correct these limits of our previous studies, here we more thoroughly describe and improve our APC method, and expand its testing across a wider range of operating conditions through new dedicated simulations. To understand whether and to what extent this new approach improves on the state of the art, we perform a comparative analysis with respect to three alternative recently described APC methods, considering power tracking accuracy and fatigue loading. Particular attention is devoted to saturation events, 90 as they are strong drivers of performance for both power and load metrics. The study is conducted on a small cluster of wind turbines with different layouts, first in steady-state with the FLOW Redirection and Induction in Steady-state (FLORIS v3) code (Gebraad et al., 2016; NREL, 2023), and then in unsteady conditions using a TUM-modified version of NREL's large-

eddy simulation with actuator line model (LES-ALM) Simulator fOr Wind Farm Applications (SOWFA) (Fleming et al., 2014; Wang et al., 2019).

95 The paper is organized as follows: Sect. 2 presents the new formulation, the tools developed to support it, and the reference APC algorithms. Section 3 describes the simulation set-up and reports and discusses results from the steady and unsteady analyses. Finally, Sect. 4 concludes and offers an outlook towards future work.

2 Active power control formulation

100 Here we present the proposed active power control formulation, which will be referred to in the following as CL+MR, which stands for closed loop with maximum reserves. The presentation covers all the APC-relevant aspects at the farm, turbine, and flow levels. The overall CL+MR APC method is described in §2.1, together with its setpoint scheduler and closed-loop corrector. Next, in §2.2 we present the reference APC methods, which will be used later on to perform a comparative analysis of the performance of our new approach. The next two sections discuss turbine-level aspects of the problem. The identification and treatment of saturation events – which play a central role in the behavior of APC methods – are discussed in §2.3, while 105 §2.4 presents the wind turbine controller. The final two sections are devoted to the flow-related aspects of the problem. A steady-state model, described in §2.5, is used to synthesize the open-loop setpoint scheduler; the same model is used later on for a preliminary performance assessment in steady conditions. Next, §2.6 describes a dynamic flow model used for gain tuning.

2.1 Robust APC: closed loop with maximum reserve (CL+MR)

110 The central component of the wind farm control system is an open-loop, model-based optimal setpoint scheduler. This scheduler determines the yaw misalignment of each turbine and calculates its contribution (i.e., power share) to meet the power demand set by the TSO, based on the current ambient conditions. Real-time wind speed and turbulence intensity can be obtained from SCADA data. Wind sensing techniques (Bertelè et al., 2021, 2024) can be employed to account for wind shear, for a finer adjustment of the controller to the ambient conditions. Additionally, a feedback loop is used to correct any tracking 115 errors that may occur from the open loop in real-time. Such errors are always present in practice, and are due to two inevitable limiting factors: first, open-loop setpoints are computed using models, which are clearly always of only a limited fidelity to reality; second, setpoints are chosen based on the knowledge of the current ambient conditions, which are always obviously known only with a limited accuracy. A block diagram of the control system is provided in Fig. 1.

120 The closed and open loops operate at different time rates due to the varying time scales of the physical phenomena they control. Specifically, the open loop, which adjusts the yaw setpoints γ_i and power shares α_i , updates at a slow rate, because of the time required by wakes to propagate downstream. On the other hand, the closed loop operates at a much faster rate, as it is charged with locally correcting small tracking errors. In this study, the open loop is updated three orders of magnitude faster than the closed loop.

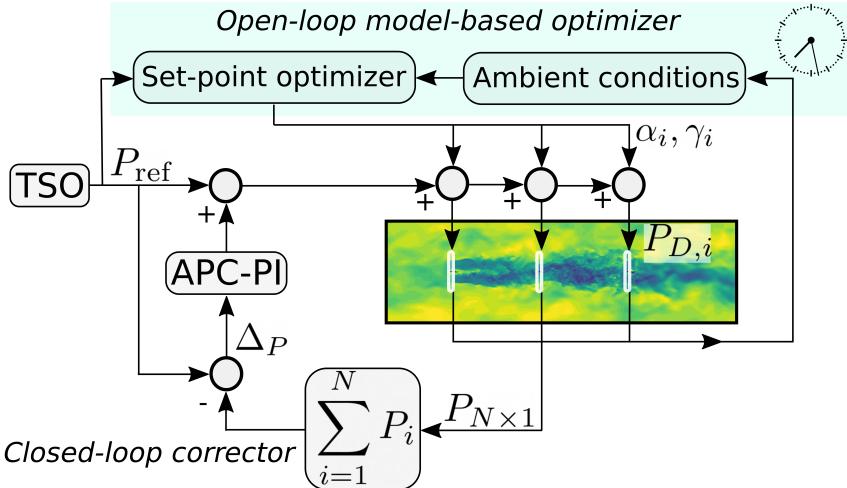


Figure 1. Schematic representation of the APC controller, featuring an open-loop model-based optimizer and a closed-loop corrector.

2.1.1 Open-loop setpoint optimal scheduler

125 The open-loop part of the algorithm calculates the optimal setpoints for yaw misalignment and power share. These are determined by solving an optimization problem that maximizes the minimum power reserve across the farm, while meeting a specified overall power demand. All necessary flow-related quantities are computed from a fast engineering model, which in this work is based on the FLORIS v3 implementation (NREL, 2023). The optimization is based on a gradient-based approach (Brayton et al., 1979). In fact, this class of methods typically outperforms derivative-free approaches (Thomas et al., 2023)
130 in the presence of active constraints at convergence (the power demand, in this case) and when the solution space is smooth (Nocedal and Wright, 2006).

The power of the i -th turbine is noted $P_i = P_i(\mathbf{A}_i, \mathbf{u}_i)$, where $\mathbf{A}_i = (U, \psi, k, I)$ indicates the local ambient conditions, here assumed to include the rotor-equivalent wind speed U , wind direction ψ , vertical shear k , and turbulence intensity (TI) I . The symbol $\mathbf{u}_i = (\gamma_i, \alpha_i)$ indicates the control inputs, which are represented by the yaw misalignment γ_i and power share $135 \alpha_i = P_{D,i} / \sum_{i=1}^N P_{D,i}$, where $P_{D,i}$ is the power demand. Power is computed using the wind farm flow model described in §2.5. The power of misaligned and curtailed wind turbines is computed based on Tamaro et al. (2024a).

The maximum power that turbine i can capture by modifying its control setpoints \mathbf{u}_i (with the setpoints of the other turbines held constant) is calculated as

$$P_{a,i} = \max_{\mathbf{u}_i} P_i(\mathbf{A}_i, \mathbf{u}_i) = \frac{1}{2} \rho \pi R^2 U^3 C_P(\lambda, \theta, \gamma), \quad (1)$$

140 where ρ is the air density, R is the rotor radius, $C_P(\lambda, \theta, \gamma)$ is the power coefficient, $\lambda = \Omega R/U$ is the tip speed ratio, where Ω is the rotor speed, and θ is the blade pitch angle.

Following Tamaro et al. (2024a) (Sect. 3.2), the power coefficient is calculated as

$$C_P(\lambda, \theta, \gamma) = C_P(\lambda, \theta, 0) \eta_P(\lambda, \theta, \gamma), \quad (2)$$

where $C_P(\lambda, \theta, 0)$ is the power coefficient in wind-aligned conditions, and η_P indicates the scaling factor that accounts for power losses due to a misalignment γ . This model can also take into account the effects resulting from a rotor tilt δ (Tamaro et al., 2024a).

As explained in more detail in Tamaro et al. (2024a), the power model accounts for the turbine-level control strategy. In fact, given the desired power setpoint α_i , misalignment angle γ_i , and present wind speed U , the tip speed ratio λ and pitch setting θ appearing in the power model Eq. (2) are computed based on the turbine control strategy. The turbine controller used here is described in §2.5, although the approach is agnostic to these specific details and other control approaches could be readily used as well.

The algorithm seeks the setpoint combination that minimizes the maximum power ratio $P_i/P_{a,i}$ across all N turbines in the farm, while ensuring the power demand of the TSO, noted P_{ref} , is met. This can be expressed as

$$\min_{\mathbf{u}} \max_{i \in [1, N]} \frac{P_i}{P_{a,i}} \text{ such that } \sum_{i=1}^N P_i = P_{\text{ref}}. \quad (3)$$

In fact, a smaller power ratio $P_i/P_{a,i}$ results in a larger reserve $r_i = 1 - P_i/P_{a,i}$, which can be used to mitigate local drops in wind.

Equation (3) represents a constrained optimization problem. However, this optimization does not need to be carried out in real-time during operation. Instead, it is performed offline for a range of ambient conditions and wind farm curtailment levels. The results are stored in a look-up table, which is then interpolated at runtime, similarly to the approach used in power-boosting wind farm control (Meyers et al., 2022). Although executing the optimization online may be feasible, performing it offline offers the advantage of validating the resulting setpoints in advance, which is beneficial from an operational safety perspective. When the maximum reserve of the farm is zero, the optimization provides the traditional maximum power solution (Meyers et al., 2022), making power boosting a limiting case of the proposed APC formulation. This feature ensures a seamless and smooth transition between the power-boosting and APC modes of operation.

By using a load surrogate model (Guilloré et al., 2024), damage could be readily introduced as a cost term or a constraint in the optimization, although this option was not considered further in this work.

2.1.2 Closed-loop corrector

The closed-loop corrector is directly adopted from the work of van Wingerden et al. (2017). It features a simple PI feedback loop that operates based on the power tracking error, which results from the open-loop part of the control system. The tracking error Δ_P is defined in this work as

$$\Delta_P = P_{\text{ref}} - \sum_{i=1}^N P_i. \quad (4)$$

The PI gains are obtained with a tuning procedure based on a simple Simulink (The MathWorks Inc., 2022) model described in Sect. 2.6. van Wingerden et al. (2017) proposed a gain scheduling based on the number of saturated turbines. This approach is not used here because it was found that – when used with a limited number of wind turbines – it may cause abrupt variations 175 in the gains that can lead to instabilities. Instead, when a turbine saturates, its local power tracking error is redistributed equally among the non-saturated ones, as proposed by Vali et al. (2019). This is explained in more detail later in Sect. 2.3.

The controller features an anti-windup term on the integrator when all turbines are saturated, and the integrator is reset when no saturation occurs, as proposed by Silva et al. (2022).

2.2 Reference APC formulations for performance comparison

180 Three reference wind farm APCs are considered with the goal of comparing the performance of the proposed power reserve boosting method, namely:

- **Open loop (OL):** this simple approach assigns predetermined setpoints α_i to each turbine as fractions of the power demanded by the TSO, so that $\sum_{i=1}^N \alpha_i = 1$. These setpoints are scheduled with the current wind direction to account for different local power availabilities, but they are independent of P_{ref} .
- 185 – **Closed loop (CL):** this method is the same as OL with the addition of the PI feedback loop described in Sect. 2.1.2.
- **Closed loop with load balance (CL+LB):** this method consists of CL without the fixed scheduling of the power share setpoints α_i . Instead, an additional PI loop is nested to distribute the α_i ’s with the purpose of balancing loads within the wind farm (Vali et al., 2019). In this work, the tower-base fore-aft bending moment is chosen as the target load. As proposed by Silva et al. (2022), the mean load is computed considering only non-saturated turbines, and an anti-windup 190 term is added to the integrator of a turbine when it saturates.

All open-loop setpoints are computed based on the engineering model of the plant (NREL, 2023). CL, CL+LB and CL+MR feature the same PI control block described in Sect. 2.1.2, so the closed-loop part is exactly the same. A visual comparison of all APC strategies considered in this work is presented in Fig. 2.

2.3 Identification and treatment of saturation conditions

195 Local turbine saturations are detected when the blade pitch lies at its optimal value, and the tracking error exceeds a given threshold, set to 1% of rated power. Both conditions need to be verified to activate saturation. The threshold for the pitch angle is set to 0.4° . When a wind turbine enters saturation, its power demand is fixed to the last value recorded, while its local tracking error is equally redistributed to non-saturated turbines in the form of an additional power demand in order to ensure that $\sum_{i=1}^N \alpha_i = 1$. This way locally isolated saturation events – even if persistent – do not introduce significant tracking errors 200 as long as other turbines with enough power reserve can compensate.

On the other hand, conditions in which all wind turbines are close to saturation are particularly harmful to the tracking accuracy, as a cascade effect can be triggered that may lead to all turbines being saturated. In fact, in this case, all wind turbines

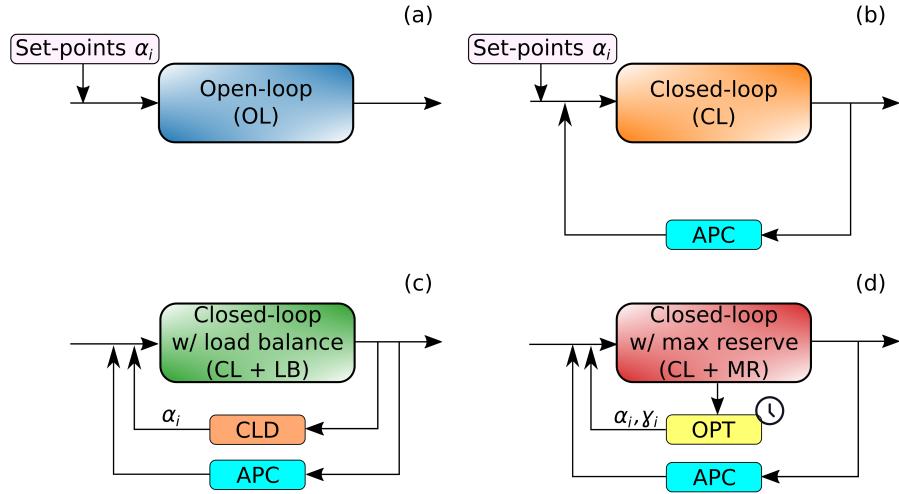


Figure 2. Visual comparison of the three APCs considered in this work, open loop (a), closed loop (b), closed loop with load balancing setpoint distributor (c), closed loop with max-reserve setpoint distributor (d).

operate in greedy mode and, if the TSO demand drops, a significant (negative with respect to Eq. 4) tracking error will arise. To avoid this situation, when the wind farm produces more than the instantaneous demand (with a threshold set to 3% of the 205 turbine rated power), every saturation condition is forcibly reset.

2.4 Wind turbine controller

The wind turbine controller is characterized by two distinct regimes, namely:

- Below rated: the blades lie at the optimal pitch angle, while the generator torque Q is related to the rotor speed Ω as $Q(\Omega) = \kappa\Omega^2$, where κ is a constant (Bossanyi, 2000).
- Above rated: each turbine yields its demanded power by collectively pitching the blades based on a standard PI loop. Q and Ω are fixed and equal to the rated values, i.e. Q_R and Ω_R , respectively.

An intermediate regime is also often present in typical controllers for noise or load reduction (e.g., thrust clipping) (Abbas et al., 2022); however, such an intermediate control regime is not considered in this work for simplicity.

The transition between the below and above rated regimes occurs when Ω exceeds Ω_R . Turbines can track a given power 215 demand $P_D = \alpha_i P_{\text{ref}}$ by adjusting Ω_R as

$$\Omega_R = \sqrt[3]{\frac{P_D}{\kappa}}, \quad (5)$$

and by setting $Q_R = P_D/\Omega_R$. In case a gearbox is present, the rotor angular velocity should be corrected by the gearbox ratio to yield the high-speed shaft velocity.

With this control approach, the blade pitch angle θ can be used to measure the margin of a curtailed wind turbine (Tamaro and Bottasso, 2023). Generally, high values of θ indicate a high margin since the turbine operates at a sub-optimal C_P . The lowest limit for θ is the optimal pitch angle θ_{opt} , which yields the maximum C_P .

This type of controller was chosen because it explicitly receives a power demand as an input, which is accurately tracked using the blade pitch angle. Other control methods could also be used, possibly significantly affecting the APC performance.

2.5 Steady-state model for control synthesis

The engineering farm flow model FLORIS v3 (NREL, 2023) is used both to synthesize the open-loop part of the controller and to perform steady-state analyses. The results reported here are based on the Gauss Curl Hybrid wake model with default wake parameters, due to its fast computation time and accurate modeling of lateral wake deflection (King et al., 2021). To model off-rated operation, the lookup tables of power coefficient C_P and thrust coefficient C_T are modified to consider curtailed conditions. The sum of squares freestream superposition model by Katic et al. (1987); Annoni et al. (2018) is used to model the interaction of the wakes.

Since off-rated operation spans a wide range of C_T – and misaligned operation is strongly dependent on C_T (Cossu, 2021; Heck et al., 2023) – the standard FLORIS model is coupled with an analytical model to predict rotor performance (Tamaro et al., 2024a), rather than relying on the \cos^{pp} method (Liew et al., 2020). In fact, the latter may yield inaccurate power estimates in the scenarios considered here, as it neglects to consider how the rotor is controlled as it is yawed out of the wind. Conversely, the misalignment model by Tamaro et al. (2024a) determines how much power is lost for a given misalignment angle γ based on the tip speed ratio λ and pitch setting θ used by the turbine. Since λ and θ are the results of the control strategy, which is itself reacting to γ , the problem is implicit. Hence, to compute $C_P = C_P(\lambda, \theta, \gamma)$ and $C_T = C_T(\lambda, \theta, \gamma)$, the balance equation between aerodynamic power P and demanded power P_D is solved iteratively. Following Tamaro et al. (2024a), the equation

$$P = \frac{1}{2} \rho \pi R^2 U^3 C_P(\lambda, \theta, \gamma) = Q(P_D, \Omega) \Omega = P_D \quad (6)$$

is solved for λ with $\theta = \theta_{\text{opt}}$. If the resultant Ω is smaller than the rated value Ω_R – which is computed from P_D via Eq. (5) – the turbine operates in below-rated conditions, whereas if $\Omega \geq \Omega_R$ it is in rated conditions. In this case, Eq. (6) is instead solved for θ , with $\lambda = \Omega_R R / U$.

In the optimization phase, we normalize P_D using the rated power P_R , defining a curtailment parameter $\epsilon = P_D / P_R$. Figure 3 presents the setpoints of λ and θ for a range of ϵ values for $U = 8 \text{ m s}^{-1}$, plotted as functions of the misalignment γ .

Figure 3a shows that when ϵ is decreased, λ is reduced due to the smaller Ω_R (see Eq. 5). Accordingly, the blade pitch angle in Fig. 3b increases to reduce the power output.

The quantities ϵ and γ are used as optimization variables. The curtailment ϵ is limited to $0 \leq \epsilon_i \leq P_{a,i} / P_R$ to ensure that the optimizer never asks for a power demand that exceeds the locally available one, while the yaw angle is bounded to $\gamma_{\min} \leq \gamma \leq \gamma_{\max}$, these turbine-dependent limits typically being imposed by ultimate and/or fatigue loads.

A different optimization is solved for each ambient condition (wind speed, wind direction, optionally turbulence intensity), to generate the associated setpoints. When the optimization converges, the resulting power share setpoints are computed as

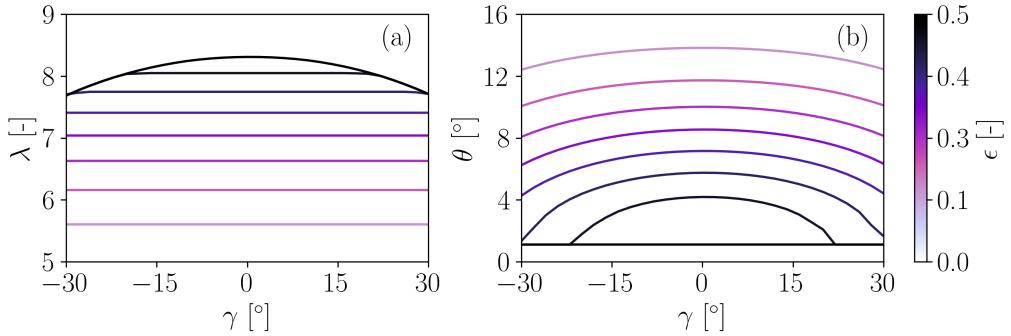


Figure 3. Control setpoints λ **(a)** and θ **(b)** plotted as functions of the misalignment angle γ for $U = 8 \text{ m s}^{-1}$ and for different values of the normalized power demand ϵ , as indicated by the different colors of the lines.

$\alpha_i = P_i / \sum_{i=1}^N P_i$, and they are stored together with the yaw setpoints in a lookup table. The optimization problem is solved with the gradient-based Sequential Quadratic Programming (SQP) method (Brayton et al., 1979). A total of 102 optimizations are performed, i.e., 51 per wind direction, each corresponding to a wind farm power request from 50% to 100% of the greedy one. During operation, α_i and γ_i are linearly interpolated from the lookup tables based on the average power demand computed over the last 30 seconds. Clipping to the last available value of the lookup table is applied to avoid extrapolation.

Figure 4 reports the yaw setpoints and power share percentages that maximize the minimum local available power for the aligned turbines of Sect. 3.1.

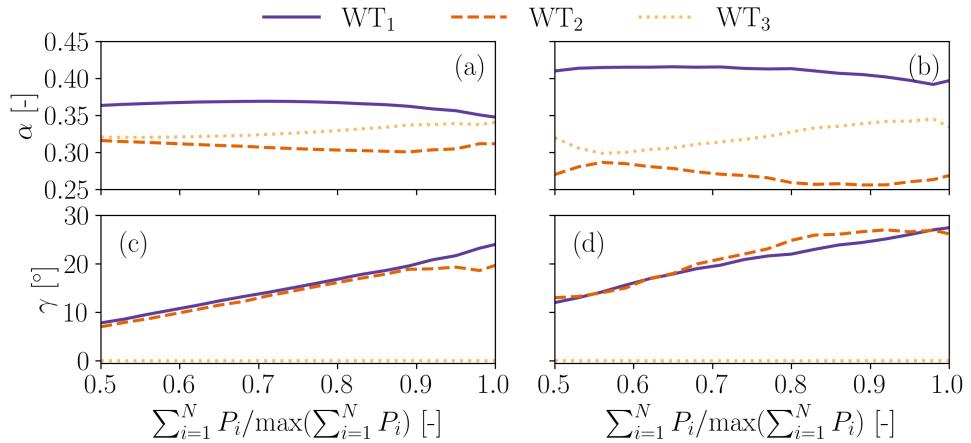


Figure 4. Optimal setpoints that maximize the minimum power reserve for the three aligned turbines of Sect. 3.1. Power setpoints **(a,b)**, yaw setpoints **(c,d)**. Wind direction $\psi = 7.1^\circ$ (50% wake overlap) **(a,c)**, wind direction $\psi = 3.6^\circ$ (75% wake overlap) **(b,d)**. The power setpoints are plotted as fractions of the available wind farm power.

The figure shows that the most upstream turbines are misaligned with respect to the wind, with the goal of increasing the power reserves of the downstream ones. The wind turbines are misaligned more when the rotor overlap is larger. The power shares result from different local inflow conditions, yaw misalignment, and wake effects.

2.6 Digital twin for gain tuning

The gains of the APC controllers are synthesized with a reduced wind turbine model coupled with a simple dynamic wake model. The Simulink (The MathWorks Inc., 2022) implementation features three wind turbines with their own controllers. The Jensen wake model (Jensen, 1983) is combined with the instantaneous thrust coefficient C_T to estimate the wake deficit for downstream wind turbines. Wake effects are delayed based on the local wind speed and wind turbine separation distance to simulate the time needed for wake effects to propagate downstream. Figure 5 presents a sketch of the digital twin. This rather crude wake model is adopted here purely because, although very simple, it still provides for a realistic estimation of wake deficit in this aligned setup; although probably not necessary for the goal of gain tuning, it is however clear that other more sophisticated models could be used.

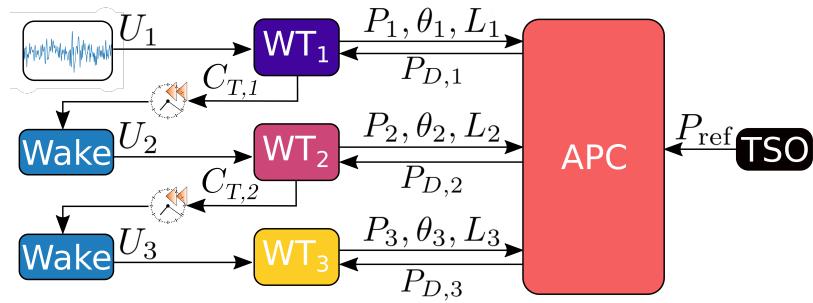


Figure 5. Block diagram of the structure of the Simulink model used for optimizing the APC gains.

The wind turbines are assumed to be fully aligned, and the inflow is taken from wind tunnel measurements in a turbulent boundary layer. The CL and CL+LB APC supercontrollers are implemented in the digital twin. Their gains are optimized with the interior-point gradient-based algorithm, where the cost function is the Root Mean Square (RMS) of the power tracking error. To improve the robustness of the gains, zero-mean white noise is added to the input measured power.

In the case of CL+LB, the cost function is defined as $J = 0.75\bar{\Delta}_P + 0.25\sum_{i=1}^N \bar{\Delta}_{L,i}$, and it represents the weighted sum of the non-dimensional tracking error $\bar{\Delta}_P$ and non-dimensional overall load-balancing error $\bar{\Delta}_{L,i}$, defined as $\bar{\Delta}_{L,i} = L_i - \sum_{i=1}^N L_i$. Both quantities are non-dimensionalized to lie in the interval [0,1].

3 Results

3.1 Numerical set-up

280 Tests are performed on an array of three wind turbines located at a streamwise distance of 4 diameters ($4D$). The use of a small cluster of turbines is motivated by the computational cost of the LES-ALM simulations. However, the small distance between the turbines implies a short propagation time that requires a very robust APC. At the same time, this layout – while still capturing the strong aerodynamic couplings caused by single and double wake interactions on downstream rotors – is sufficiently simple to clearly isolate and understand the underlying effects within the wind farm.

285 Two scenarios are considered, with wind directions ψ equal to 3.6° and 7.1° , corresponding to 75% and 50% rotor overlaps, respectively. Rotor overlap is defined as the lateral offset between the centers of two wind turbine rotors, expressed as a percentage of the rotor diameter. An overlap of 100% means that the rotors are perfectly aligned (center-to-center), while 50% overlap means that the downstream rotor is laterally displaced by half a rotor diameter. These scenarios were chosen because they involve partial wake impingement, which is particularly relevant for fatigue considerations (Guilloré et al., 2024).

290 The most upstream turbine is labeled WT_1 , the most downstream one WT_3 , and the one in the center is referred to as WT_2 . Simulations are conducted using the IEA 3.4 MW reference wind turbine, a typical onshore machine with a contemporary design (Bortolotti et al., 2019). Here, we only note that the turbine has a rotor diameter D of 130 m, 5° up tilt angle (i.e. $\delta = -5^\circ$) (Tamaro et al., 2024a), the optimal pitch is $\theta_{\text{opt}} = 1.09^\circ$, and the constant κ for the generator torque in below-rated regime is $\kappa = 1,804 \text{ kN m s}^2 \text{ rad}^{-2}$. The maximum yaw rate is equal to $0.8 \text{ }^\circ \text{s}^{-1}$. The PI gains for the CL+MR controller 295 are $K_{APC}^P = 1.3127$ [-] and $K_{APC}^I = 0.2063 \text{ s}^{-1}$, while the ones for the coordinated load distribution (CLD) of CL+LB are $K_{CLD}^P = 0.03427 \text{ Nm}^{-1}$ and $K_{CLD}^I = 0.07959 \text{ Nm}^{-1} \text{ s}^{-1}$. The APC open loop is updated every 30 seconds, and the closed loop every 0.01 seconds. The closed loop frequency of 100 Hz is faster than the 1-10 Hz reported by other studies (Liu et al., 2019; Vali et al., 2019), and it is due to the small time step used by the LES-ALM solver. In CL+LB, loads are filtered by applying exponential smoothing with a time constant of 0.1 s.

300 LES-ALM simulations are used for testing the performance of the new APC formulation, because they can represent the complex dynamics typical of wind turbine wakes and their interactions (Wang et al., 2019). We use an in-house version of the LES-ALM code SOWFA (Troldborg et al., 2007; Wang, 2021), which includes the smearing correction to blade tip forces proposed by Meyer Forsting et al. (2019). The incompressible solver is based on a finite-volume formulation and uses a standard Smagorinsky model to treat subgrid scales, with a constant of 0.13 (Sagaut, 2006).

305 The LES Cartesian mesh comprises approximately 14.3 million cells and includes six refinement levels. The smallest cells measure 1 m and are located in correspondence with the three rotors. Two tilted hexahedral regions are used to refine the wind farm array. The computational domain, grids, and turbine layout are shown in Fig. 6.

310 Unsteady tests are run with a turbulent wind obtained from a precursor generated in stable atmospheric conditions with periodic inlet-outlet boundary conditions. The inflow is characterized by a turbulence intensity $TI = 5.7\%$, a hub height wind speed $U_\infty = 9.54 \text{ ms}^{-1}$, a power-law shear coefficient $k = 0.21$, and an integral length scale of turbulence at hub height equal

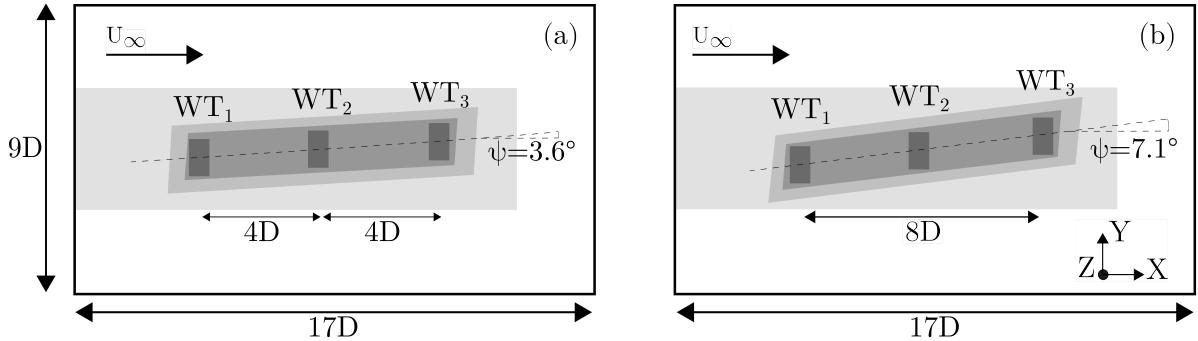


Figure 6. Wind farm layout and simulation scenarios for 75% **(a)**, and 50% **(b)** rotor overlap. The shaded areas indicate the mesh refinement levels.

to $0.79D$. The normalized mean streamwise velocity and turbulence intensity fields at the inlet – $4D$ upstream of WT_1 – are shown in Fig. 7.

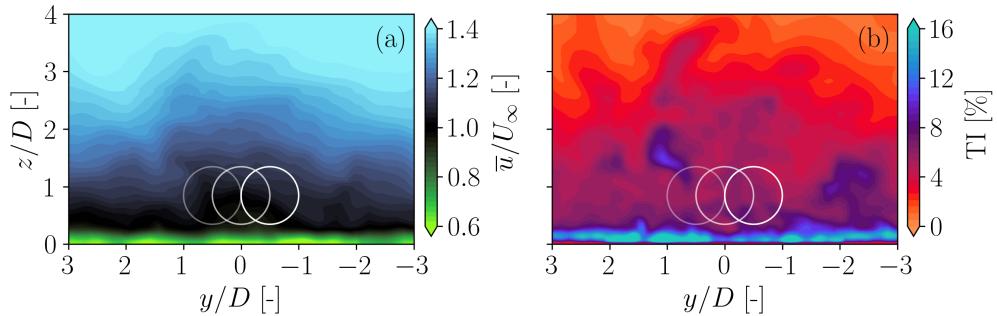


Figure 7. Slices of inflow fields. Time-averaged wind speed normalized by the free-stream wind speed **(a)**; turbulence intensity field **(b)**. The locations of the three wind turbines for $\psi = 7.1^\circ$ (50% wake overlap) are marked with a line transparency that increases along with the distance from the inflow slice.

Each LES-ALM simulation is run for 1,200 seconds. The first 200 seconds are considered as the initial transient, and hence they are discarded. The total time of 1,000 seconds is approximately equivalent to 1.6 standard 10-minute seeds, which is 315 less than the minimum recommended value (Liew and Larsen, 2022), but in line with the numerical study of van Wingerden et al. (2017). The figure shows that the three rotors are immersed in the boundary layer and that the flow is not perfectly uniform, which is expected given the integral length scale of $0.79D$ and the averaging time of only 1,000 seconds. Given the short duration, results are likely dependent on the specific inflow realization, especially at high power demands due to the multiple simultaneous saturations. To mitigate this effect, LES-ALM simulations are also performed with all turbines operating 320 in greedy mode. The results of the different APC formulations are then compared to the corresponding greedy results.

3.2 Reference power demand signal

A dynamic reference power signal typical of automatic generation control (AGC) is used as a reference power signal. AGC is the secondary response regime of grid frequency control, and it consists of the modification of the power output of a plant depending on dynamically changing requests by the TSO (Aho et al., 2012). A similar signal has also been considered by other authors (Fleming et al., 2016; van Wingerden et al., 2017; Shapiro et al., 2017; Boersma et al., 2018; Vali et al., 2019). The signal is defined as

$$P_{\text{ref}}(t) = P_{\text{greedy}}(\psi) \left(b + cn_k^{\text{AGC}}(t) \right), \quad (7)$$

where $n_k^{\text{AGC}}(t)$ is a normalized perturbation from a standard test signal, P_{greedy} is the time-averaged available power of the wind farm in greedy conditions, and b and c are parameters that shift and change the amplitude of n_k^{AGC} , respectively. A “ground-truth” greedy power output P_{greedy} is computed from a LES-ALM simulation in which the turbines are aligned with the inflow and operate in greedy power mode (i.e., $P_D = P_R$, see Sect. 2.4). Figure 8 shows the time history of the normalized perturbation n_k^{AGC} .

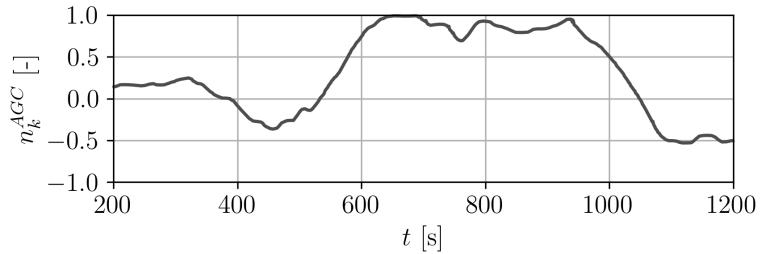


Figure 8. Time history of the normalized perturbation signal n_k^{AGC} of Eq. (7).

P_{greedy} is computed on a preliminary LES-ALM simulation under the same ambient conditions of the APC runs, where the turbines operate aligned and in greedy mode. The value $c = 0.1$ is used in all cases, while three different values of b are considered, i.e., $b = [0.7, 0.8, 0.9]$. In general, a higher value of b makes the TSO signal harder to track, due to a closer proximity to the greedy power available.

3.3 Steady-state analysis

First, the proposed method CL+MR is compared to OL using the steady state engineering flow model FLORIS. This analysis is performed on the cluster of three IEA 3.4 MW turbines at a streamwise distance of 4D and three inflow angles $\psi = [0^\circ, 3.6^\circ, 7.1^\circ]$, corresponding to rotor overlaps of 100%, 75%, and 50%, respectively.

In addition to this basic comparison to the simplest possible (open-loop) approach, here we also would like to understand the role of wake steering, and whether it alone would be sufficient to achieve a satisfactory APC performance. To this end, we

consider the situation where the OL power setpoints α_i are superimposed with the yaw misalignment setpoints that optimize power output, i.e., the classical wake-steering-based power boosting. Such a method is – in principle – attractive because it is simpler than CL+MR, but, similarly to it, includes wake steering. This approach is referred to in the following as OL+power boosting.

Results are shown in Fig. 9, where we report the difference between the rotor power and the locally available one, normalized by the rated power. We use this quantity instead of the power reserve r in order to compare the APCs in absolute terms.

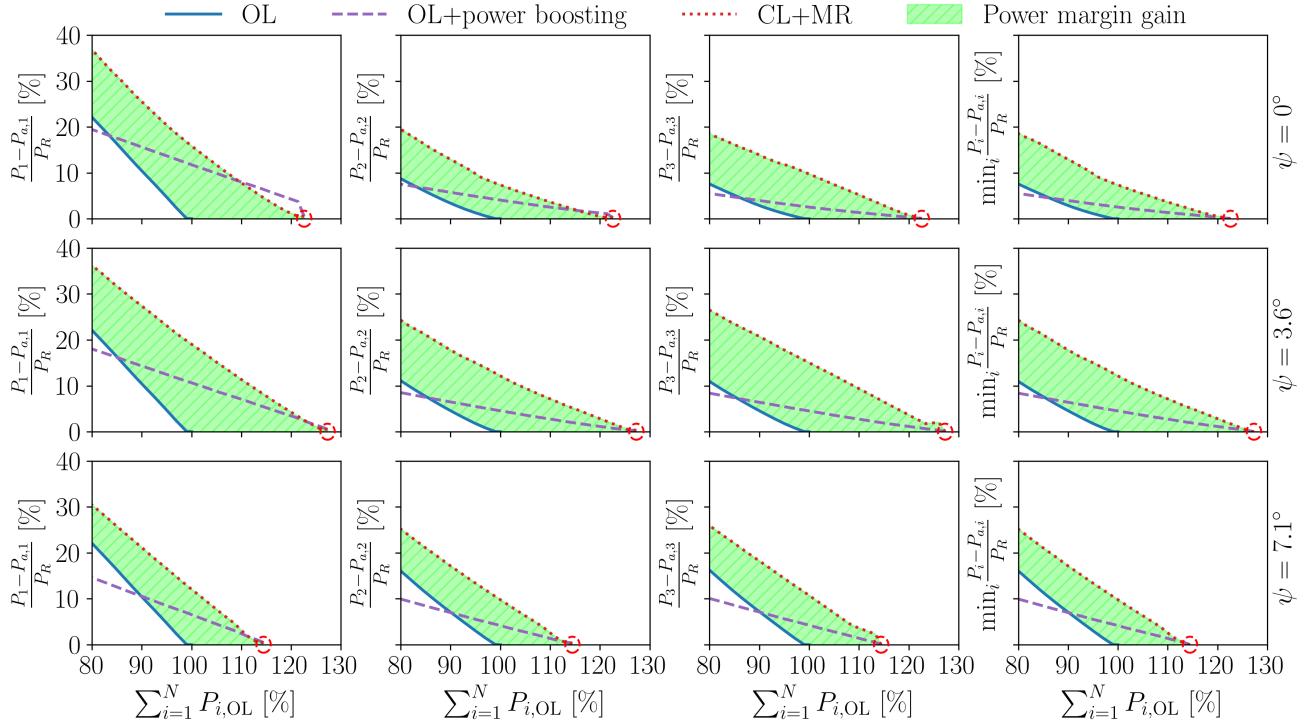


Figure 9. Local power margin in percentage obtained with a purely inductive wind farm control (labeled OL), with an inductive APC that operates with statically misaligned rotors according to the power boosting solution (labeled OL+power boosting), and with the proposed controller that dispatches power and yaw set-points to optimize power reserves (labeled CL+MR). The first three columns show the power margin for WT₁, WT₂, and WT₃, respectively, while the fourth column indicates the minimum power margin in the wind farm. The first row is obtained at $\psi = 0^\circ$, i.e., 100% rotor overlap; the second row at $\psi = 3.6^\circ$, i.e., 75% rotor overlap; and the third row at $\psi = 7.1^\circ$, i.e., 50% rotor overlap. On the x -axis, wind farm power is normalized by the maximum value for OL. The red circles indicate the classical maximum power solution, and the colored area represents the power reserve improvement of CL+MR with respect to OL.

The figure shows that the margin drops to zero in correspondence with the maximum power of the plant, whereas it increases as the power demand is lowered and the wind turbines are curtailed. Wake steering effectively extends the power available to the wind farm. In agreement with the literature, the effectiveness of wake steering depends on the direction of the wind with

respect to the alignment of the turbines, an effect that is here captured by the ψ angle. In all cases, the proposed CL+MR strategy increases the power margin for all wind farm power demands. The highest power margin improvement is observed for $\psi = 3.6^\circ$ (75% wake overlap), and it is in excess of +20%. As the power demand $\sum_{i=1}^N P_{i,OL}$ is reduced, the added margin of CL+MR diminishes since wake effects get weaker and the effectiveness of wake steering is reduced, in both cases because of the lower C_T . In the figure, the classical maximum power solutions are shown using red circle symbols.

The plots show that OL+power boosting achieves the same power reserve improvement of CL+MR only at the red-circled points, where both the α_i and γ_i setpoints are identical. However, as the wind farm power demand decreases, the power margin of OL+power boosting recovers at a much slower rate than for CL+MR, resulting in a smaller improvement compared to the standard OL approach. At lower power demands – specifically when $\sum_{i=1}^N P_{i,OL} \approx 90\%$ – OL+power boosting actually yields a smaller minimum reserve than OL. This can be attributed to an imbalance in the distribution of power reserves across the wind farm. Specifically, upwind rotors that are yawed for steering their wakes have reduced margins compared to downstream, aligned rotors, which benefit from wake steering. As a result, in CL+MR, the optimizer tends to realign the yawed wake-steering rotors when the power demand drops. Conversely, in OL+power boosting, these rotors remain misaligned: the effectiveness of wake steering diminishes with reducing C_T , which in turn decreases the locally available power. This effect is more pronounced at $\psi = 7.1^\circ$ (50% wake overlap), where wake steering is inherently less effective.

Overall, WT₁ presents the largest power margin for both OL and CL+MR, since it is driven by a clean inflow, while for downstream rotors the margin recovers at a lower rate, due to the non-linear behavior of wake recovery.

3.4 Unsteady conditions

Next, results from the LES-ALM simulations are presented. Figure 10 reports slices of the time-averaged free-stream velocity fields $\bar{u}(x,y)$ at hub height, normalized by the free-stream value U_∞ . Figure 11 shows the turbulence intensity TI. The plots are shown for OL and the proposed CL+MR for $\psi = 3.6^\circ$ (75% wake overlap) and $\psi = 7.1^\circ$ (50% wake overlap), for a power demand level $b = 0.8$. CL and CL+LB are not shown for the sake of simplicity, since they are qualitatively similar to OL.

These figures highlight the different extent of wake impingement that occurs on the downstream rotors at $\psi = 3.6^\circ$ (75% wake overlap) and $\psi = 7.1^\circ$ (50% wake overlap). Similarly, Figs. 10b and 10d allow one to appreciate the effect of yaw misalignment on the wakes, which are significantly deflected to one side compared to fully aligned conditions. The effect is remarkable when considering that the turbines are curtailed, and hence the smaller C_T reduces the deflection compared to a greedy scenario.

The plots of turbulence intensity in Fig. 11 highlight the fact that waked rotors operate in regions of significant turbulence, as expected. The turbulence levels reach rather large values in correspondence with the wake of WT₃, especially in OL with 75% rotor overlap (see Fig. 11a). Overall, the rather high turbulence intensity levels are likely to be also the result of the unsteady operation of the rotors due to power tracking. For the proposed CL+MR method, the turbulence intensity levels are generally significantly lower than OL for two reasons. First, wake steering reduces wake overlaps and, secondly, the rotors operate generally at lower C_T values, which in turn lead to reduced amplitudes of wake meandering (Foti et al., 2018) and to smaller speed gradients over the rotor.

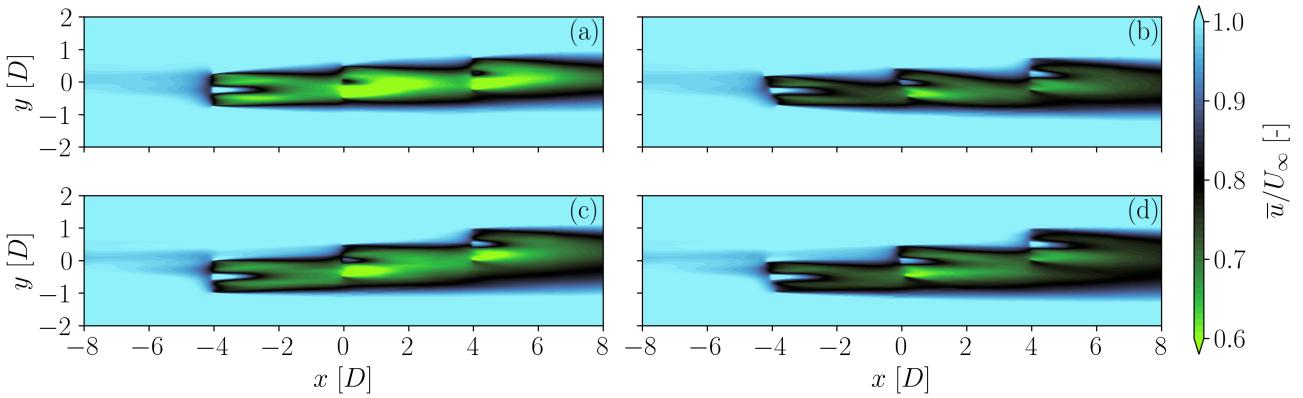


Figure 10. Time-averaged streamwise velocity fields for a TSO power request $b = 0.8$, for two APC control strategies: open loop (OL) (a,c); closed loop with optimal power reserve (CL+MR) (b,d). $\psi = 3.6^\circ$ (75% wake overlap) (a,b); $\psi = 7.1^\circ$ (50% wake overlap) (c,d). The slices are extracted at hub height.

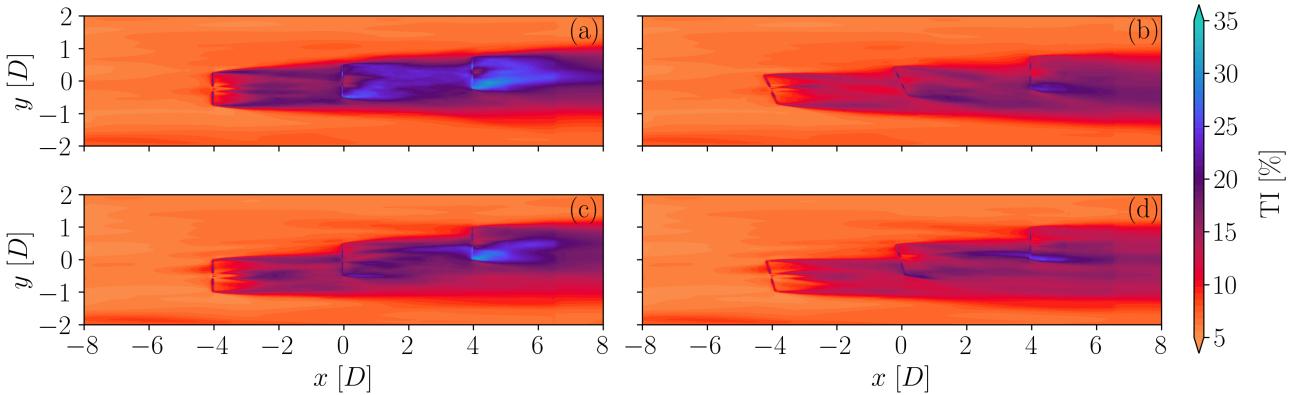


Figure 11. Turbulence intensity fields for a TSO power request $b = 0.8$, for two APC control strategies: open loop (OL) (a,c); closed loop with optimal power reserve (CL+MR) (b,d). $\psi = 3.6^\circ$ (75% wake overlap) (a,b); $\psi = 7.1^\circ$ (50% wake overlap) (c,d). The slices are extracted at hub height.

Next, the time series of produced and demanded power are shown in Figs. 12 and 13 for $\psi = 3.6^\circ$ (75% wake overlap) and $\psi = 7.1^\circ$ (50% wake overlap), respectively. The grey background shows the “ground truth” wind farm power from a preliminary LES-ALM simulation where the turbines operate aligned in greedy mode. This greedy power can differ and even be lower than the one produced in the APC cases, because of the slightly different wake effects due to curtailment. Nevertheless, 390 it is reported here to provide a proxy for the instantaneous power available in the wind farm.

Both figures show that at the low demand value $b = 0.7$ all methods track power somewhat accurately. In fact, in this case, the power demanded by the TSO is always smaller than the greedy power, and in general, all turbines have enough power reserve to avoid saturations. As the power demand increases, it gets closer to the available one, leading to tracking inaccuracies driven by local saturation events. Still, for $b = 0.8$, the effectiveness of the closed-loop methods is evident. This is mostly due to the 395 setpoint redistribution logic that makes unsaturated turbines compensate for the tracking error of the saturated ones. However, when $b > 0.8$, the power available to the wind farm is often not enough to track the reference signal, and this is especially clear in the time interval $650 \text{ s} < t < 1000 \text{ s}$. In such cases, the higher power availability made possible by wake steering allows only CL+MR to track the signal, while other methods clearly lack the power reserve to do so.

In the zoomed regions of the plots, the saturation events can be clearly observed. In the open-loop case, the power demand 400 remains unaffected since no countermeasure is implemented. Conversely, in the closed-loop cases, the power demand remains constant for saturated rotors since the other ones are called to compensate.

3.4.1 Saturation events

Figures 14 and 15 provide a visual representation of the occurrences of local saturation conditions for each APC strategy and different TSO request scenarios for $\psi = 3.6^\circ$ (75% wake overlap) and $\psi = 7.1^\circ$ (50% wake overlap).

405 These plots indicate that waked wind turbines – i.e., WT_2 and WT_3 – are often saturated because of their generally small power reserve. This effect is exacerbated in the scenario with stronger wake impingement when $\psi = 3.6^\circ$ (75% wake overlap). CL does not reduce the extent of saturation events compared to OL, but it rather increases them since other turbines are requested to compensate.

The strategy CL+LB presents the highest number of saturation events of waked turbines. This is because, in order to balance 410 the tower-base fore-aft bending moment, waked turbines are requested to yield a power that is not actually available. At high power demands (i.e., $b > 0.8$), waked wind turbines operate in greedy mode, while WT_1 is responsible for following the TSO signal. In this case, the difference between 50% and 75% rotor overlap is remarkable, since in the latter case WT_1 is often saturated, while in the former it is not. This is due to the different wake impingement that modifies the load distribution.

To further understand the reasons for the observed saturation events, the collective blade pitch angle θ is considered. In fact, 415 this parameter can be seen as a proxy of the local power reserve, as discussed in Sect. 2.4. Figures 16 and 17 show the mean collective blade pitch angle $\bar{\theta}$, binned according to the instantaneous required power, normalized by the wind farm available power, as determined from a greedy LES-ALM simulation in the same conditions. This type of binning is chosen to decouple the aerodynamic effects of the APC strategy from sporadic events introduced by the characteristics of the inflow seed. All simulations with different power demand (b) levels are included in the average for a given ψ , and only bins with a minimum

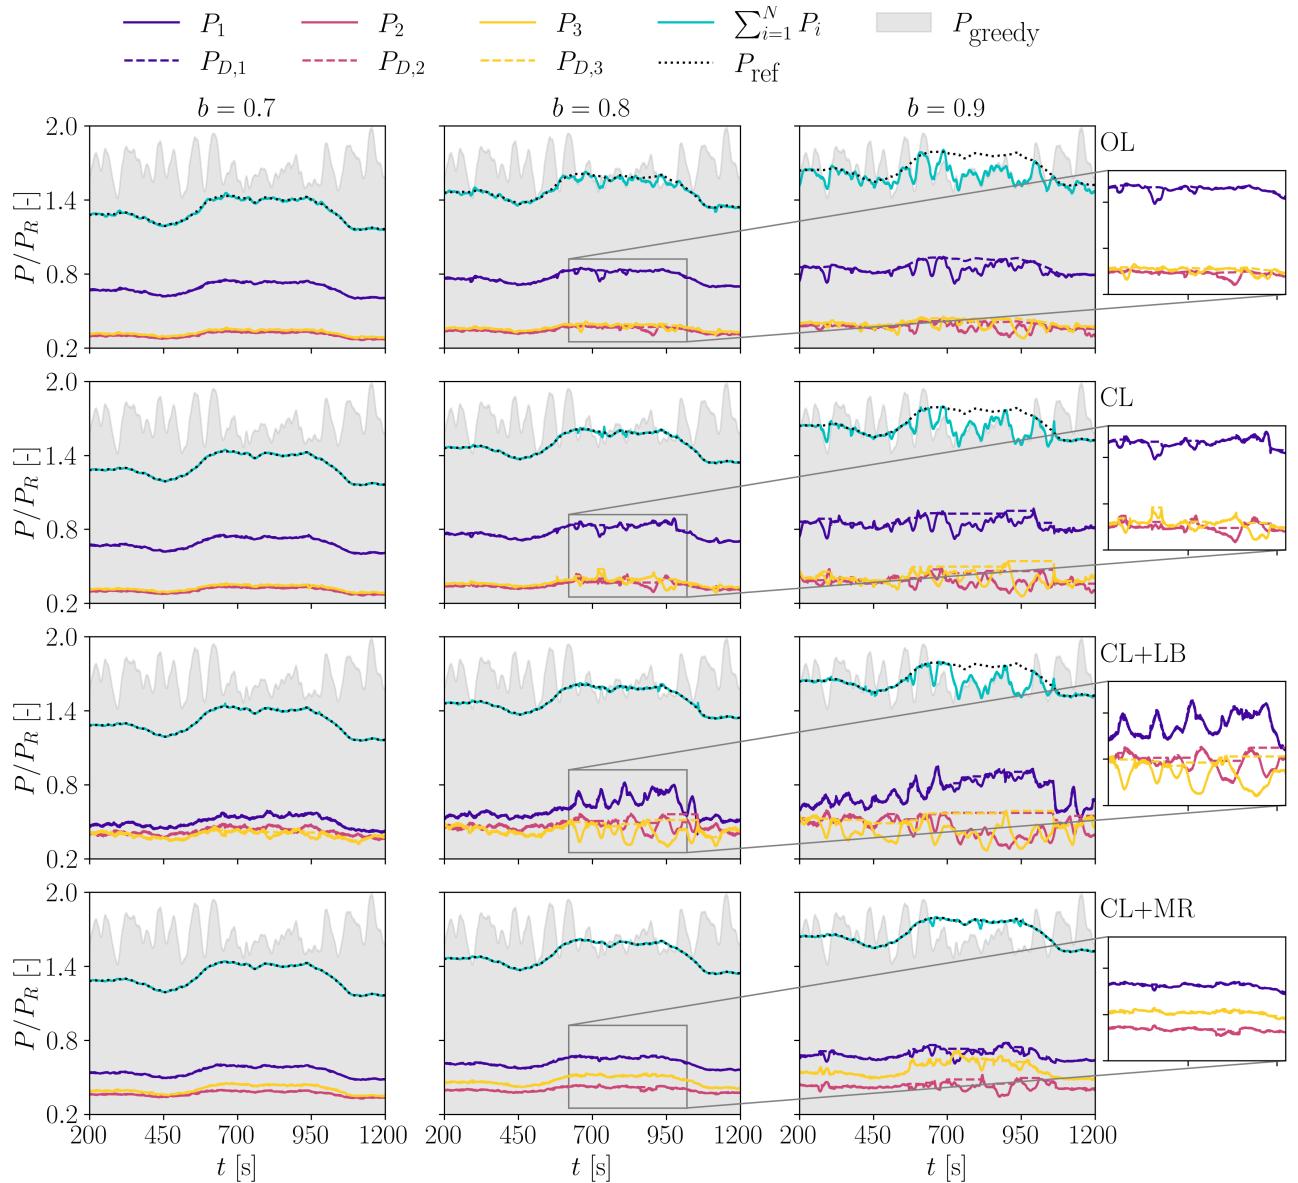


Figure 12. Time series of power demand $P_{D,i}$ and power generated P_i for $\psi = 3.6^\circ$ (75% wake overlap). The y axes are non-dimensionalized by the rated power P_R of the IEA 3.4 MW reference wind turbine. The wind farm power obtained in greedy operation (without wake steering) is displayed in the background in grey.

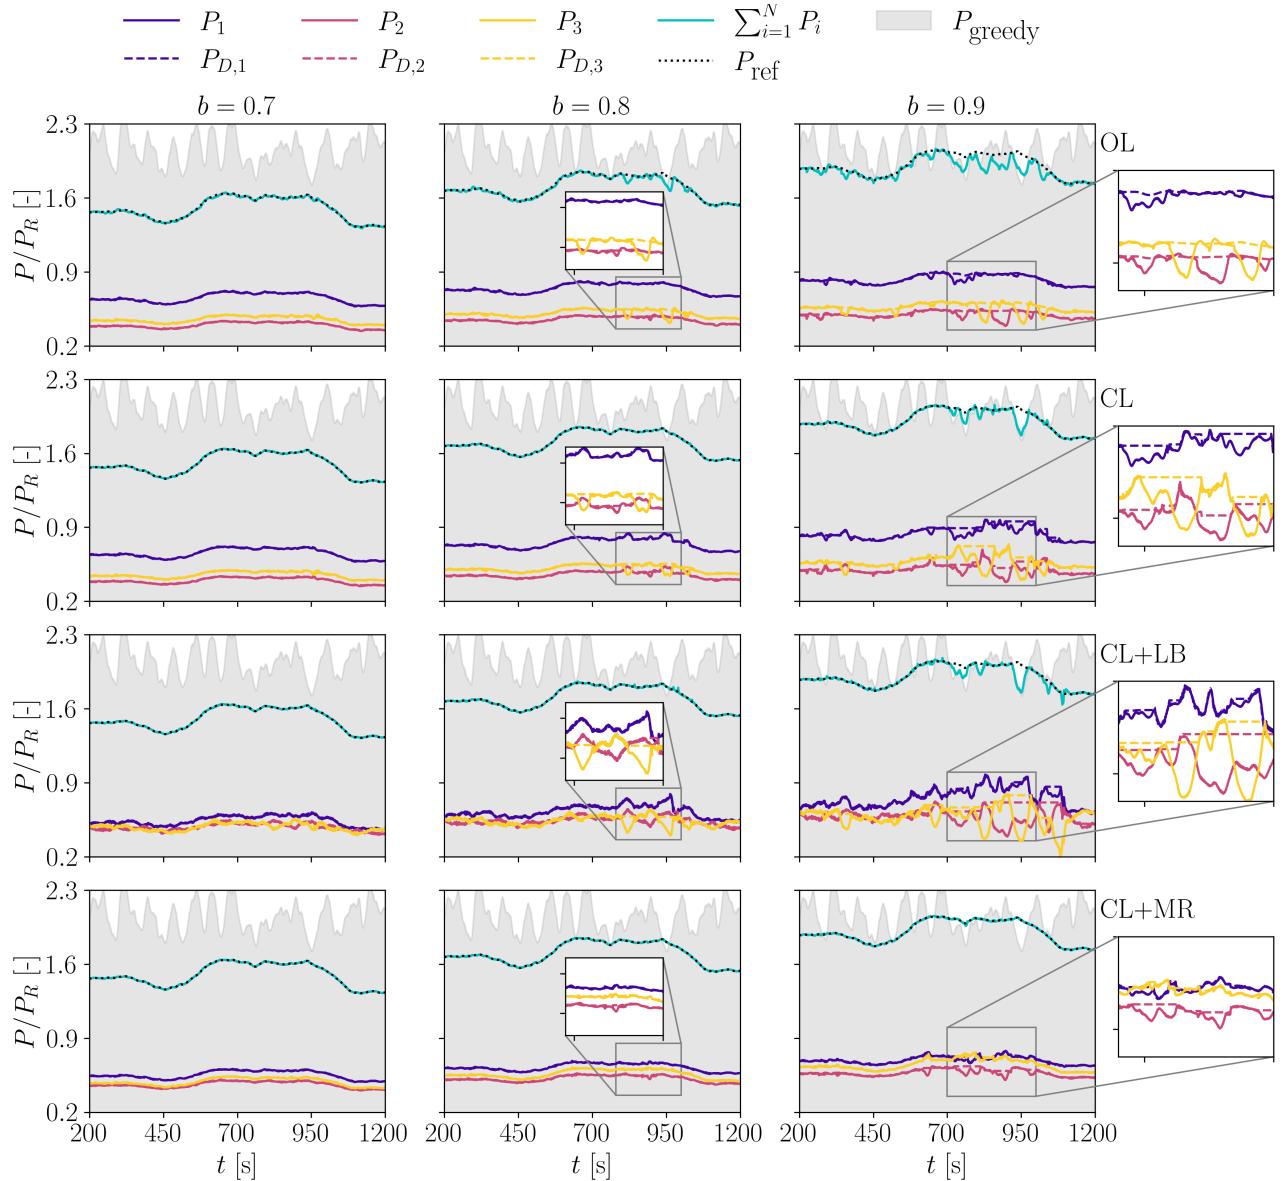


Figure 13. Time series of power demand $P_{D,i}$ and power generated P_i for $\psi = 7.1^\circ$ (50% wake overlap). The y axes are non-dimensionalized by the rated power P_R of the IEA 3.4 MW reference wind turbine. The wind farm power obtained in greedy operation (without wake steering) is displayed in the background in grey.

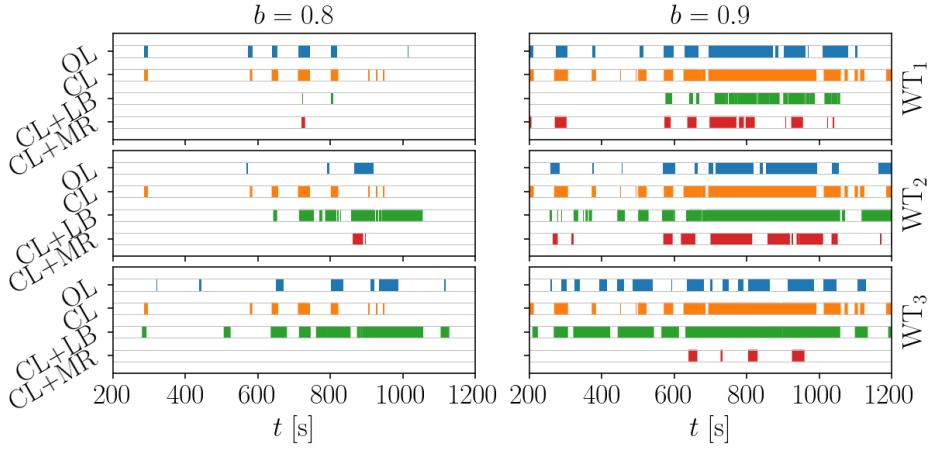


Figure 14. Occurrence of local saturation events for the three turbines in the array and the different control strategies for $\psi = 3.6^\circ$ (75% wake overlap). Only TSO requests with higher demand ($b \geq 0.8$) are considered, as few saturations were observed for low demand values ($b = 0.7$).

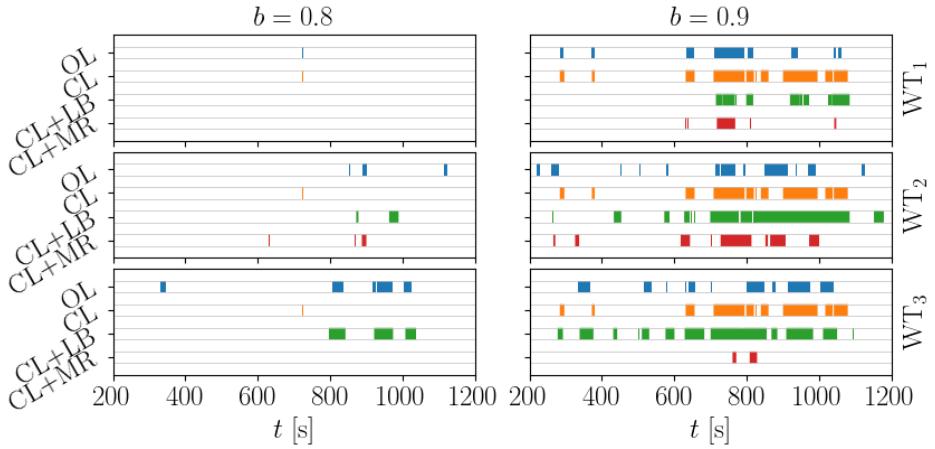


Figure 15. Occurrence of local saturation events for the three turbines in the array and the different control strategies for $\psi = 7.1^\circ$ (50% wake overlap). Only TSO requests with higher demand ($b \geq 0.8$) are considered, as few saturations were observed for low demand values ($b = 0.7$).

420 total length of 30 seconds are included. Figure 16 refers to the scenario $\psi = 3.6^\circ$ (75% wake overlap), while Fig. 17 reports the case for $\psi = 7.1^\circ$ (50% wake overlap).

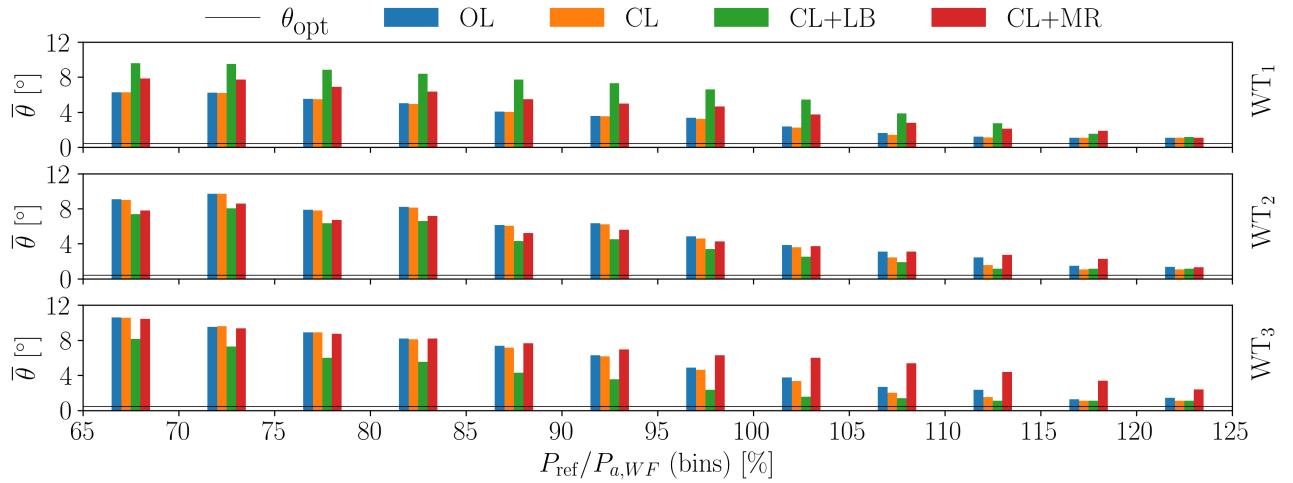


Figure 16. Mean collective blade pitch angle binned by the instantaneous available power computed from the time series of greedy power and TSO power request for $\psi = 3.6^\circ$ (75% wake overlap). The values represent the means over all tested demands (b values). Only bins with a minimum total length equivalent to 30 seconds are considered. The optimal blade pitch angle value is indicated by a black horizontal line.

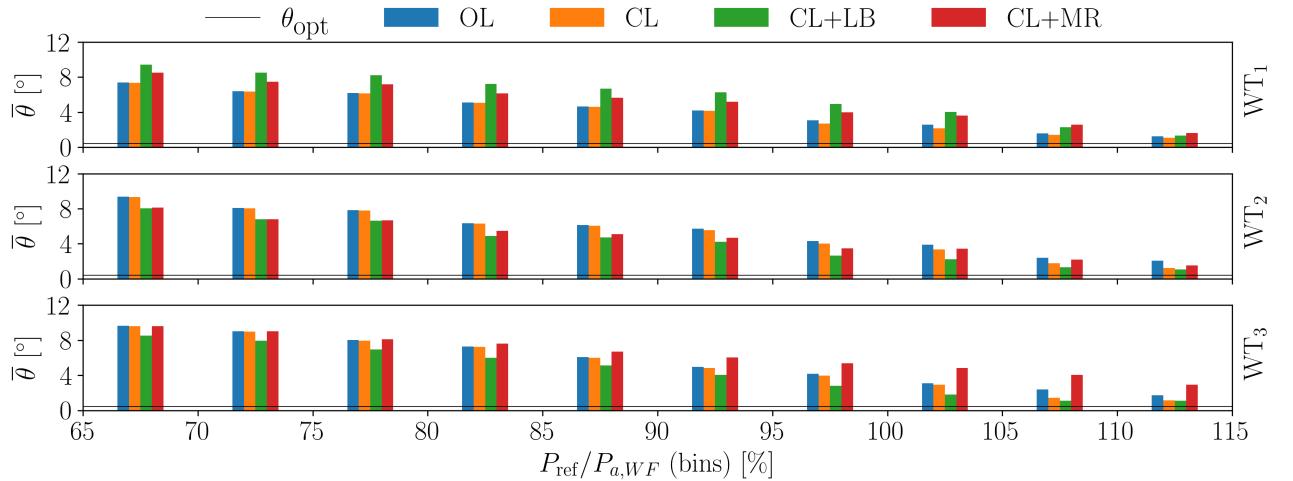


Figure 17. Mean collective blade pitch angle binned by the instantaneous available power computed from the time series of greedy power and TSO power request for $\psi = 7.1^\circ$ (50% wake overlap). The values represent the means over all tested demands (b values). Only bins with a minimum total length equivalent to 30 seconds are considered. The optimal blade pitch angle value is indicated by a black horizontal line.

These results show that OL and CL share the same power reserve distribution since, in both cases, the same power setpoints α_i are used. As expected, when the power demand increases, all CL-based methods present a lower $\bar{\theta}$ than in the OL case; this is due to the closed-loop correction and to the setpoint redistribution that occurs with saturations.

425 In line with the results shown previously, CL+LB presents a distribution of power reserves that is strongly unbalanced towards the upstream turbine. In fact, the goal of balancing loads implies that waked turbines receive relatively high power demand setpoints, which push them closer to saturation as the power demand increases. This is an important undesired and yet unreported side effect of the load balancing strategy; as shown later on, in turn, this effect undermines the very ability of the strategy to balance loads in certain situations. The proximity to saturation is visible in Figs. 16 and 17, as the mean blade 430 pitch angles approach the optimum value θ_{opt} . Conversely, the proposed CL+MR approach presents a distribution of reserves that is balanced within the wind farm, resulting in a generally higher power reserve for waked turbines and a lower one for the upstream turbine.

435 Additionally, CL+LB appears to be sensitive to the wind direction ψ , as a strong wake impingement at $\psi = 3.6^\circ$ (75% wake overlap) yields a more unbalanced power reserve distribution than for $\psi = 7.1^\circ$ (50% wake overlap). Conversely, the effect of the amount of wake overlap is less evident for the proposed CL+MR. This is expected since, in this case, the setpoints specifically depend on the wind direction.

3.4.2 Power tracking accuracy

440 Figure 18 presents the RMS of the power tracking error Δ_P for both wind direction (i.e., wake overlap) scenarios. Results indicate that closed-loop methods can significantly reduce the tracking error compared to the open-loop approach. The effectiveness is higher for low demand values $b \leq 0.8$ and gradually reduces as the power demand is increased, as it appears already in the time series reported in Figs. 12 and 13. In all cases, the proposed maximum reserve method CL+MR presents the lowest power tracking error. This is to be expected since, for $b > 0.7$, there are short events where the power demand exceeds the available greedy power. Such events degrade the tracking accuracy for methods that do not include wake steering. This highlights the importance of including wake steering in APC, as an effective tool for increasing power reserves.

445 To perform a more comprehensive comparison, the same analysis is repeated. This time, however, the tracking error time series are binned according to the instantaneous required power normalized by the wind farm available power, as done earlier for $\bar{\theta}$. Only bins with a minimum length of 30 seconds are considered, and the values of Δ_P in each bin from each power demand scenario (b request) are averaged. Only positive values of Δ_P are considered here, since they represent a lack of power, i.e., the inability of the plant to deliver what was requested. Results are shown in Figs. 19 and 20 for $\psi = 3.6^\circ$ (75% wake overlap) and $\psi = 7.1^\circ$ (50% wake overlap), respectively.

450 This analysis indicates that all closed-loop methods perform similarly at moderate TSO demands, i.e., $b < 0.8$, with a remarkable performance improvement over OL. The improvement of closed-loop methods is due to the faster response of the wind farm and to the treatment of saturation conditions.

455 However, as the power demand approaches $P_{\text{ref}}/P_{a,WF} \approx 100\%$, the proposed CL+MR exhibits a much improved tracking accuracy than all other methods. If, on the one hand, this is to be expected when $P_{\text{ref}}/P_{a,WF} > 100\%$ due to the overall higher

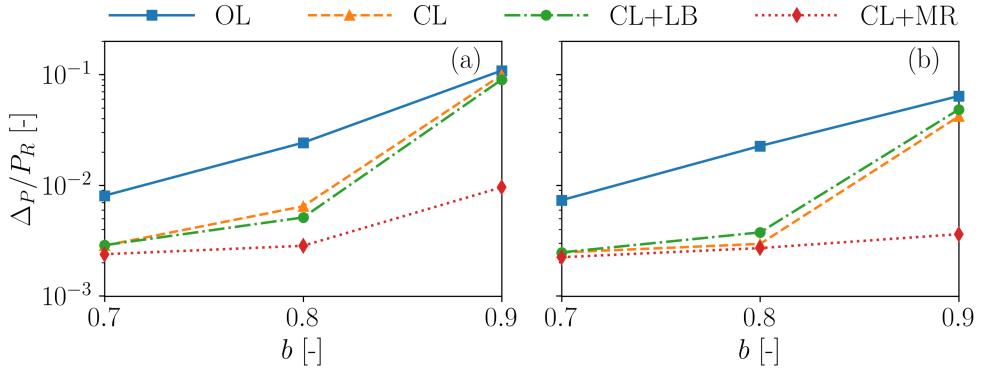


Figure 18. RMS of tracking error normalized by the turbine rated power, as functions of the TSO demand level (b values). $\psi = 3.6^\circ$ (75% wake overlap) (a); $\psi = 7.1^\circ$ (50% wake overlap) (b).

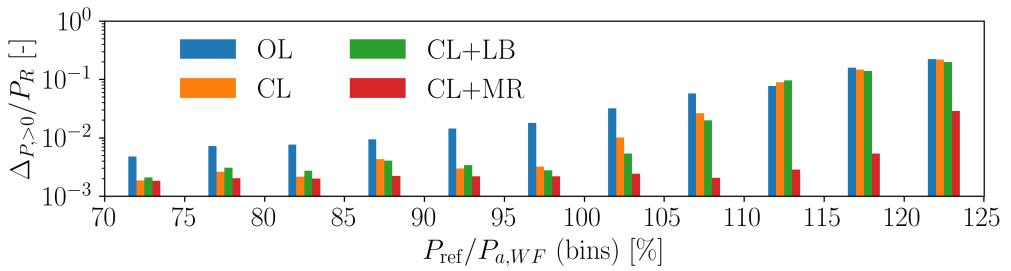


Figure 19. Positive tracking error binned by the instantaneous available power computed from the time series of greedy power and TSO power request for $\psi = 3.6^\circ$ (75% wake overlap). The values on the y axis are normalized by the wind turbine rated power. The values represent the means over all tested demands (b values). Only bins with a minimum total length equivalent to 30 seconds are considered.

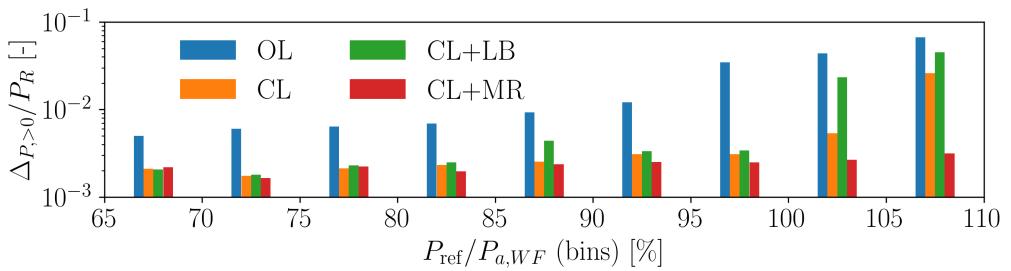


Figure 20. Positive tracking error binned by the instantaneous available power computed from the time series of greedy power and TSO power request for $\psi = 7.1^\circ$ (50% wake overlap). The values on the y axis are normalized by the wind turbine rated power. The values represent the means over all tested demands (b values). Only bins with a minimum total length equivalent to 30 seconds are considered.

wind farm power, the improvements observed for $80\% < P_{\text{ref}}/P_{a,WF} < 100\%$ are less obvious. In fact, since CL, CL+LB, and CL+MR share the exact same closed-loop part of the APC controller, the better performance of the new CL+MR is to be attributed to the more balanced power reserve distribution, consistently with the mean blade pitch angles shown in Figs. 16 and 17. This is also in agreement with the results of the steady-state analysis shown in Sect. 3.3.

460 3.4.3 Load and fatigue analysis

Here we characterize the performance of the proposed CL+MR strategy on the loading of the turbines, and we compare it to the other reference APC methods.

First, the average loading of the turbines is analyzed by considering the mean tower-base fore-aft bending moment L . Since the turbines are yawing out of the wind when performing wake steering, the moment is computed by always considering the 465 component orthogonal to the rotor orientation. Figures 21 and 22 report the results for the three wind turbines in the array, for the cases $\psi = 3.6^\circ$ (75% wake overlap) and $\psi = 7.1^\circ$ (50% wake overlap), respectively.

The figures show that L increases with the TSO power demand level b . In all cases except CL+LB, WT_1 presents the largest load, while L is rather comparable for WT_2 and WT_3 . CL+LB successfully balances L in the wind farm. It only fails to do so when $b = 0.9$ for $\psi = 3.6^\circ$ (75% wake overlap), due to the persistent saturation of waked rotors. In this case, in fact, the 470 controller is designed to prioritize power tracking over load balancing.

Interestingly, the proposed CL+MR also presents a rather balanced load distribution for $\psi = 7.1^\circ$ (50% wake overlap), but with slightly lower values than CL+LB. Conversely, for $\psi = 3.6^\circ$ (75% wake overlap), the loading on WT_3 grows along with b , due to the larger power share of this rotor compared to the other strategies.

Next, the effects on fatigue are analyzed by computing the damage equivalent loads (DELs) for the tower-base fore-aft 475 bending moment using the rainflow counting algorithm (Downing and Socie, 1982). Simulations were performed with rigid wind turbine models, and hence some dynamic effects are neglected. Nevertheless, this analysis still captures the effects generated by the different APC strategies on aerodynamically induced loads.

Figures 23 and 24 present the DELs of the three wind turbines in the array for $\psi = 3.6^\circ$ (75% wake overlap) and $\psi = 7.1^\circ$ (50% wake overlap), respectively.

480 Similarly to the average loading discussed earlier, DELs increase with TSO power demand b because of the higher loads that are involved. In general – and as expected –, it appears that DELs grow as the turbines are more impinged by wakes. In fact, WT_3 is very often the most loaded turbine, possibly because of the relatively high turbulence associated with partial wake impingements.

There is however a notable exception for CL+LB (which is explicitly designed to ensure an equal balancing of the loads), 485 where sometimes WT_1 is the most highly loaded machine. This is explained by the previous analyses (see Fig. 14), which showed that WT_1 is often the only turbine responsible for tracking the TSO signal, while the waked WT_2 and WT_3 operate in greedy mode. This points to the fact that load balancing by itself may not always be able to achieve the desired effect of an equal distribution of damage, defying its very design goal.

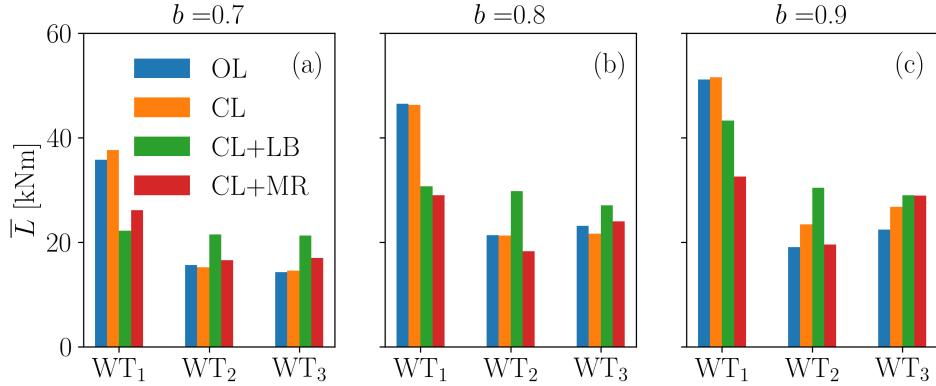


Figure 21. Mean tower-base fore-aft bending moment for the TSO demand levels $b = 0.7$ (a), $b = 0.8$ (b), $b = 0.9$ (c), for $\psi = 3.6^\circ$ (75% wake overlap).

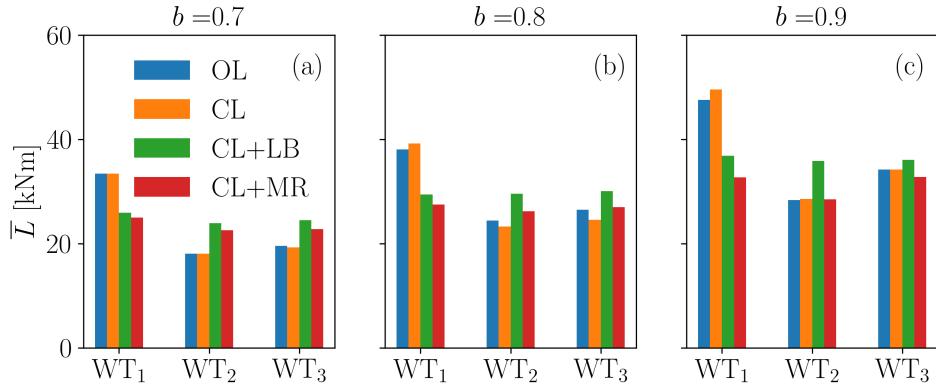


Figure 22. Mean tower-base fore-aft bending moment for the TSO demand levels $b = 0.7$ (a), $b = 0.8$ (b), $b = 0.9$ (c), for $\psi = 7.1^\circ$ (50% wake overlap).

It is clear that the DELs in Figs. 23 and 24 could be biased by some unique events. These conditions can arise due to the 490 high-amplitude cycles that occur, especially in correspondence with simultaneous saturation events. For this reason, no further conclusions are drawn from these plots. Rather, the fatigue analysis is repeated by binning the load time series according to the instantaneously required power, normalized by the wind farm available power, as done already in Sect. 3.4.1. In this case, multiple DELs are computed on continuous time segments of at least 45 seconds that belong to the same seed, and they are later summed together. DELs are computed in this way for each TSO demand scenario (b values), and they are then averaged. 495 The results are shown in Figs. 25 and 26 for $\psi = 3.6^\circ$ (75% wake overlap) and $\psi = 7.1^\circ$ (50% wake overlap), respectively.

These results allow for some further insight into the behavior of the various controllers. For all APCs, damage increases along with $P_{\text{ref}}/P_{a,\text{WF}}$. OL and CL do not present significant differences, as the power share distribution is similar. In these cases, the turbines operating in waked conditions are clearly more damaged than the upstream one. CL+LB works especially

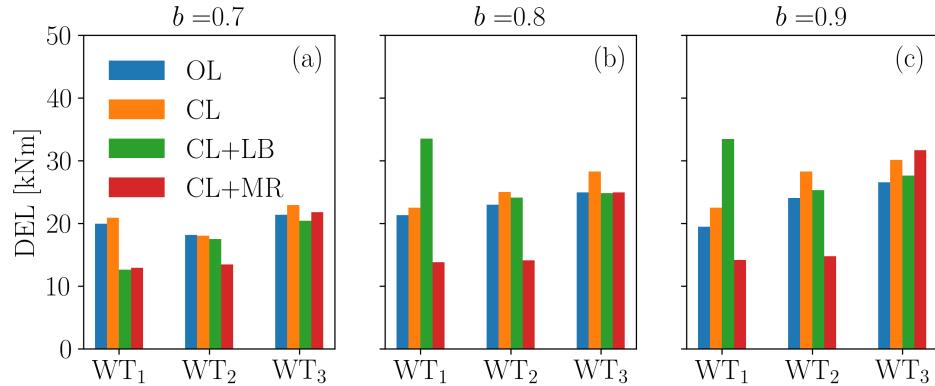


Figure 23. Tower-base fore-aft bending DELs for the cases $b = 0.7$ (a), $b = 0.8$ (b), $b = 0.9$ (c), for $\psi = 3.6^\circ$ (75% wake overlap).

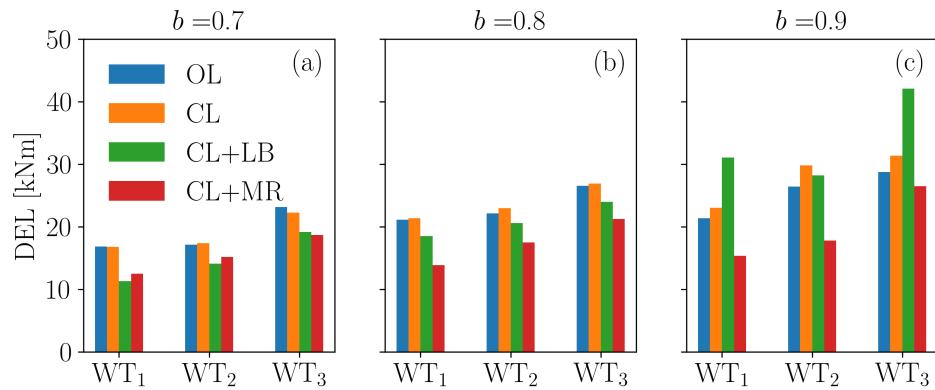


Figure 24. Tower-base fore-aft bending DELs for the cases $b = 0.7$ (a), $b = 0.8$ (b), $b = 0.9$ (c), for $\psi = 7.1^\circ$ (50% wake overlap).

well for the simpler case of strong curtailments $P_{\text{ref}}/P_{a,WF} < 80\%$, with rather low DELs in accordance with Vali et al. (2019).

500 Damage is also well distributed among the wind turbines, but only in the less difficult cases.

Notably, the proposed CL+MR presents a relatively constant damage distribution, with improved performance especially in the more difficult cases for $P_{\text{ref}}/P_{a,WF} > 80\%$. Results also indicate a strong damage reduction on WT₁ and WT₂ for the larger rotor overlap condition (75%, $\psi = 3.6^\circ$).

The previous analysis quantified fatigue damage but did not attribute it to specific dynamic sources. To distinguish the 505 dominant sources of load unsteadiness, we categorize damage d based on the load cycle duration T obtained by rainflow counting, according to the following three bins:

- High frequency: $0 \text{ s} \leq T \leq 5 \text{ s}$, generated by aerodynamic loads due to small turbulent eddies, rotational effects (1P, 3P harmonics), and the associated fast responses of the turbine and farm-level controllers.

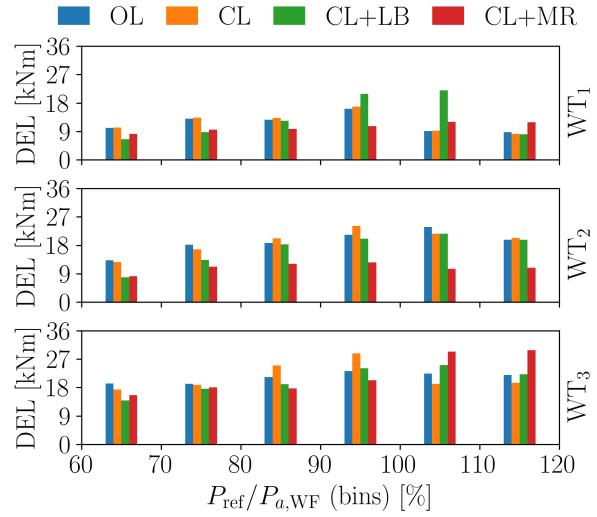


Figure 25. Tower-base fore-aft bending DELs for $\psi = 3.6^\circ$ (75% wake overlap). DELs are binned with the instantaneous available power computed with the time series of greedy power and TSO power request. The values represent the mean over all the tested demands (b values). Only continuous time segments with a minimum duration of 45 seconds are considered.

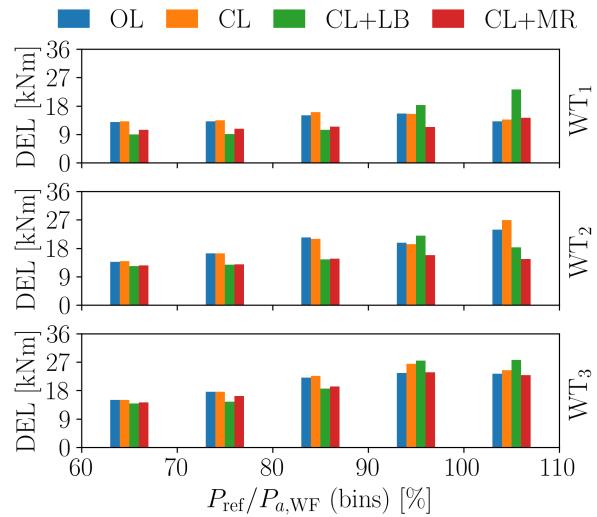


Figure 26. Tower-base fore-aft bending DELs for $\psi = 7.1^\circ$ (50% wake overlap). DELs are binned with the instantaneous available power computed with the time series of greedy power and TSO power request. The values represent the mean over all the tested demands (b values). Only continuous time segments with a minimum duration of 45 seconds are considered.

– Medium frequency: $5 \text{ s} \leq T \leq 30 \text{ s}$, associated with larger-scale flow eddies (the integral time scale of the inflow is approximately 11 s), including wake meandering, and the resulting control actions.

– Low frequency: $30 \text{ s} \leq T \leq 20 \text{ min}$, caused by slower phenomena such as secondary grid frequency control (power tracking), long-period gusts, large-scale wake effects, and the resulting variations in power share (and hence loading).

Clearly this definition of the bins is somewhat arbitrary, but still enables a rough initial categorization of the driving phenomena. Since the simulations assume rigid wind turbine structures, aeroelastic dynamics are excluded from this analysis.

Figures 27 and 28 present the results of this binning, where we have considered the total damage deriving from the simulations with all the different TSO demand levels (b values). The damage that occurred in a specific cycle has been normalized with the damage of the baseline greedy control case (grey area in Figs. 12 and 13), so that the bars labeled “Greedy” sum up to one.

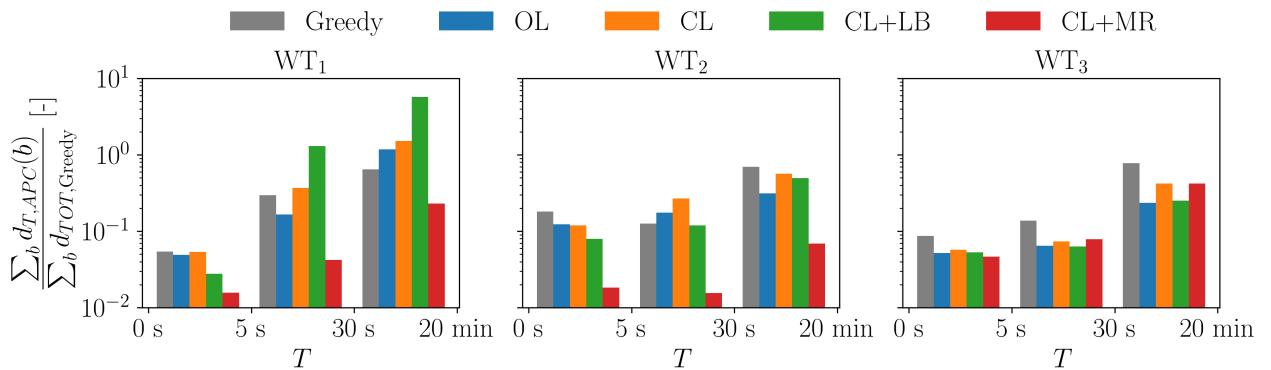


Figure 27. DELs of the tower-base fore-aft bending moment for $\psi = 3.6^\circ$ (75% wake overlap) binned by cycle duration T , and normalized by the total damage of a baseline greedy control case in that bin.

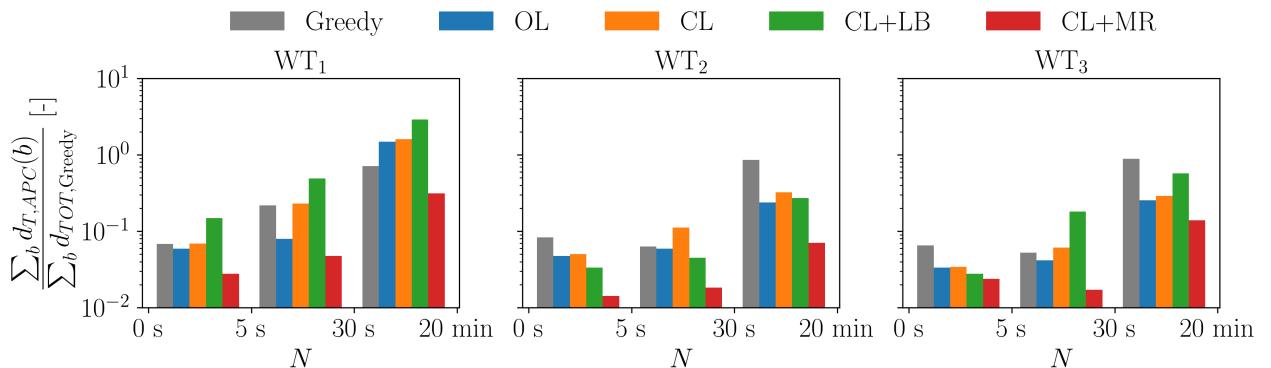


Figure 28. DELs of the tower-base fore-aft bending moment for $\psi = 7.1^\circ$ (50% wake overlap) binned by cycle duration T , and normalized by the total damage of a baseline greedy control case in that bin.

Results indicate that the waked turbines WT_2 and WT_3 present more damage in the high-frequency range than WT_1 , as a
520 result of the small eddies that characterize wake turbulence.

The greedy case presents the highest damage in the low-frequency range for waked rotors (WT_2 and WT_3), likely due to strongly fluctuating inflows caused by wake impingement. All APC methods perform better than the greedy case, because the low-frequency rotor dynamics are driven by the smooth reference power demand signal, so that large load fluctuations that would derive from wake turbulence are somehow mitigated.

525 As expected, closing the loop (CL) results in larger damage compared to operating in open loop (OL), especially in the medium frequency range, due to the extra turbine control activity that derives from the closed-loop correction. This is clearly visible in Fig. 13, where the local power oscillations of the various turbines are much larger when using CL than for OL, especially for the higher TSO demand levels ($b \geq 0.8$). However, as already noted, these larger oscillations at the level of the turbines result into a much improved tracking of the power signal by the whole wind farm (cf. Fig. 13 and 18).

530 As already observed earlier, CL+LB in general performs rather poorly for WT_1 when compared to all other methods. This is due to the extra thrust-loading that results from the need to compensate the saturations of WT_2 and WT_3 . The effect of the added loading on damage is dramatic, due to the persistent saturations resulting from the CLD loop, as shown in Figs. 14 and 15.

535 The proposed CL+MR significantly outperforms all methods in the entire frequency spectrum for WT_1 and WT_2 , due to the mitigation of wake impingement by wake steering, a generally reduced blade pitch actuation, and a particularly smooth power output made possible by the larger reserves, consistently with the time series shown in Figs. 12 and 13. This is true also for WT_3 in the milder waking scenario ($\psi = 7.1^\circ$, 50% wake overlap), thanks to wake steering that is able to clean its inflow. For the stronger waking case ($\psi = 3.6^\circ$, 75% wake overlap), while the benefit on WT_2 is dramatic, the performance for WT_3 is similar to the other methods, probably due to the combined effects of two impinging wakes.

540 Overall, results indicate that low-frequency load cycles are responsible for the largest portion of damage. As shown in the time series of Figs. 12 and 13, low-frequency cycles of power (and therefore load) often coincide with saturations, especially when a compensation mechanism is in place, like in CL, CL+LB, and CL+MR. In fact, as a result of saturation, the compensating rotor has to track an additional power signal, which amplifies its control activity and increases fatigue. Furthermore, saturated rotors operate in greedy mode, which means that their loads are subjected to low-frequency variations deriving from
545 the local inflow, which is often a waked one. This highlights once again the importance of the treatment and reduction of saturation events.

4 Conclusions and outlook

We have presented a new wind farm APC method to robustly track a reference power signal in turbulent wind conditions. The controller collectively operates the wind turbines to maximize the minimum local available power. This reserve can then be
550 exploited for accurate power tracking, ensuring minimal insurgence of saturation events.

The algorithm combines one open loop and one closed loop, which operate at different time rates. A modified version of FLORIS was utilized to compute the wind turbine setpoints for the open-loop branch in a gradient-based optimization. The PI gains of the closed-loop branch were tuned with a digital twin that mimics the wind turbine dynamics and the propagation of wake effects.

555 The new methodology was tested on a small cluster of three wind turbines with persistent wakes. We quantified the power reserve with a steady-state analysis and compared the new algorithm with a standard open-loop APC method. We observed that the new methodology is particularly effective at increasing power reserves when it can mitigate strong wake impingements.

In addition, the new approach was tested with unsteady LES-ALM simulations, showing an accurate power tracking performance, which was, in most cases, largely superior to the one provided by reference controllers representing the state of the art 560 in APC. We have shown that this better accuracy is explained by the strong reduction of saturation events and the evenly spread power reserve, as measured by the mean collective blade pitch angle. We have binned the data according to the instantaneous available wind farm power to exclude biasing by isolated events.

Overall, the following observations should be highlighted:

- The power reserve of the wind turbines in a wind farm is significantly affected by the extent of wake impingement.
- 565 – The new methodology is very effective at creating additional reserves when the wind farm operates close to its maximum power capacity, because wake deficits are particularly strong in those conditions, and the effectiveness of wake steering is maximum.
- Combining wake steering with induction control can improve wind farm performance by leveraging the strengths of both 570 strategies. This can be particularly effective for optimizing objectives beyond power maximization, such as reducing structural loads, mitigating fatigue, and managing dynamic responses in general.
- Power tracking accuracy dramatically depends on the occurrence of saturation events. In this regard, it should be noted that:
 - When one wind turbine saturates, it is extremely beneficial to redistribute its power tracking error in the form of an additional power demand to other turbines that are not saturated.
 - 575 – The main hindrance to tracking accuracy is represented by conditions in which simultaneous saturations occur. These can trigger cascading effects of power redistribution, or – in a worst case scenario – can push all wind turbines to operate in greedy mode even if the wind farm power demand is exceeded.
 - When applying PI methods, it is extremely important to implement anti-windup procedures to hedge against saturations.
 - 580 – In the presence of saturations, non-saturated turbines are called to compensate for the saturated ones, resulting in high-amplitude, low-frequency load cycles, which have a large, negative effect on fatigue.

- Saturations not only harm power tracking, but also have significant effects on loads. In fact, load balancing (Vali et al., 2019) by itself fails to effectively balance loads in the presence of saturations. We have observed and reported here several instances where the proposed method, although lacking in the present implementation a dedicated load-balancer, results in more uniform load distributions than CL+LB. This suggests that an improved version of the proposed CL+MR method might exhibit an even better performance if it included a load balancing criterion. Beyond balancing, the proposed approach generally produced much reduced loadings on the turbines when compared to the alternative methods.

The main limitation of this work is the rather small duration of the LES-ALM simulations. Although in line with similar studies in the literature, this limited duration could have biased some results due to particular events occurring within the inflow time histories. To account for this, we performed greedy simulations to quantify the actual wind farm available power at every time instant. Another limitation of the unsteady results is the use of rigid wind turbine models in the simulations. If, on the one hand, this should not play a major role in the behavior of the far wakes (Salavatidezfouli et al., 2025), on the other, it somehow limits the conclusions that can be drawn from the analysis of fatigue.

Appendix A: Nomenclature

595	A	Ambient conditions
	<i>b</i>	Shift of normalized power signal (power demand level)
	C_P	Power coefficient
	C_T	Thrust coefficient
	<i>c</i>	Amplitude of normalized power signal
600	<i>D</i>	Rotor diameter
	<i>d</i>	Damage
	K^I	Control gain (integral)
	K^P	Control gain (proportional)
	<i>L</i>	Reference load signal (tower-base fore-aft bending moment)
605	n_k^{AGC}	Normalized power demand perturbation
	<i>N</i>	Number of turbines in a farm
	<i>P</i>	Wind turbine power
	P_D	Wind turbine power demand
	P_R	Rated power
610	P_{ref}	Reference power signal
	p_p	Cosine law exponent (yaw misalignment)
	<i>R</i>	Rotor radius
	<i>r</i>	Power reserve
	<i>s</i>	Saturation
615	<i>t</i>	Time
	<i>T</i>	Load cycle duration
	<i>U</i>	Rotor-equivalent wind speed
	U_∞	Free-stream wind speed at hub height
	u	Control inputs
620	<i>x</i>	Cartesian coordinate
	<i>y</i>	Cartesian coordinate
	<i>z</i>	Cartesian coordinate
	α	Power-share setpoint
625	Δ_L	Load balancing error
	Δ_P	Power tracking error
	ϵ	Curtailment

η_P	Power loss factor
γ	Rotor yaw angle
630 λ	Tip speed ratio
Ω	Rotor angular velocity
Ω_R	Rated rotor angular velocity
ψ	Wind direction
ρ	Air density
635 θ	Blade pitch angle
AGC	Automatic generation control
APC	Active power control
CL	Closed loop
640 CLD	Coordinated load distribution
CL+LB	Closed loop with load balancing
CL+MR	Closed loop with maximum reserve
DEL	Damage equivalent load
FA	Fore-aft
645 FLORIS	FLOW Redirection and Induction in Steady State
OL	Open loop
PI	Proportional integral
RMS	Root mean square
SOWFA	Simulator fOr Wind Farm Applications
650 TI	Turbulence intensity
TSO	Transmission system operator
WF	Wind farm
WT	Wind turbine

655 *Code and data availability.* The supercontroller codes in C++ for all APC methods discussed here are available on Zenodo at <https://doi.org/10.5281/zenodo.4500333>.
 Videos of one of the simulations are available at https://youtu.be/dS_FrPhw3EM.

Author contributions. CLB developed the formulation of the new APC method and supervised the overall research. ST implemented the model, performed the experiments, and conducted the steady-state analyses with FLORIS. FC supported the implementation of the methods.

All authors contributed to the interpretation of the results. CLB and ST wrote the manuscript, with contributions by FC. All authors provided
660 important input to this research work through discussions and feedback and improved the manuscript.

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

Financial support. This work has been supported in part by the PowerTracker project, which receives funding from the German Federal
Ministry for Economic Affairs and Climate Action (FKZ: 03EE2036A). This work has also been partially supported by the MERIDIONAL,
665 SUDOCO and TWAIN projects, which receive funding from the European Union's Horizon Europe Programme under grant agreements
101084216, 101122256, and 101122194 respectively.

Acknowledgements. The authors express their gratitude to the Leibniz Supercomputing Centre (LRZ) for providing access and computing
time on the SuperMUC Petascale System under Projekt-ID pr84be “Large-Eddy Simulation for Wind Farm Control”. The authors would
like to thank Abhinav Anand and Anik H. Shah for their support in the implementation of the steady-state model and in the load and fatigue
670 analysis, respectively.

References

Abbas, N. J., Zalkind, D. S., Pao, L., and Wright, A.: A reference open-source controller for fixed and floating offshore wind turbines, *Wind Energy Science*, 7, 53–73, <https://doi.org/10.5194/wes-7-53-2022>, 2022.

Aho, J., Buckspan, A., Laks, J., Fleming, P., Jeong, Y., Dunne, F., Churchfield, M., Pao, L., and Johnson, K.: A tutorial of wind turbine control for supporting grid frequency through active power control, in: 2012 American Control Conference (ACC), pp. 3120–3131, <https://doi.org/10.1109/ACC.2012.6315180>, 2012.

Ally, S., Verstraeten, T., Daems, P.-J., Nowé, A., and Helsen, J.: Modular deep learning approach for wind farm power forecasting and wake loss prediction, *Wind Energy Science*, 10, 779–812, <https://doi.org/10.5194/wes-10-779-2025>, 2025.

Annoni, J., Fleming, P., Scholbrock, A., Roadman, J., Dana, S., Adcock, C., Porte-Agel, F., Raach, S., Haizmann, F., and Schlipf, D.: Analysis of control-oriented wake modeling tools using lidar field results, *Wind Energy Science*, 3, 819–831, <https://doi.org/10.5194/wes-3-819-2018>, 2018.

Bertelè, M., Bottasso, C. L., and Schreiber, J.: Wind inflow observation from load harmonics: initial steps towards a field validation, *Wind Energy Science*, 6, 759–775, <https://doi.org/10.5194/wes-6-759-2021>, 2021.

Bertelè, M., Meyer, P. J., Sucameli, C. R., Fricke, J., Wegner, A., Gottschall, J., and Bottasso, C. L.: The rotor as a sensor – observing shear and veer from the operational data of a large wind turbine, *Wind Energy Science*, 9, 1419–1429, <https://doi.org/10.5194/wes-9-1419-2024>, 2024.

Boersma, S., Rostampour, V., Doekemeijer, B., van Geest, W., and van Wingerden, J.-W.: A constrained model predictive wind farm controller providing active power control: an LES study, *Journal of Physics: Conference Series*, 1037, 032 023, <https://doi.org/10.1088/1742-6596/1037/3/032023>, 2018.

Bortolotti, P., Tarres, H. C., Dykes, K. L., Merz, K., Sethuraman, L., Verelst, D., and Zahle, F.: IEA Wind TCP Task 37: Systems Engineering in Wind Energy - WP2.1 Reference Wind Turbines, Tech. rep., National Renewable Energy Lab. (NREL), <https://doi.org/10.2172/1529216>, 2019.

Bossanyi, E. A.: The Design of closed loop controllers for wind turbines, *Wind Energy*, 3, 149–163, <https://doi.org/10.1002/we.34>, 2000.

Brayton, R., Director, S., Hachtel, G., and Vidigal, L.: A new algorithm for statistical circuit design based on quasi-Newton methods and function splitting, *IEEE Transactions on Circuits and Systems*, 26, 784–794, <https://doi.org/10.1109/TCS.1979.1084701>, 1979.

Campagnolo, F., Petrović, V., Schreiber, J., Nanos, E. M., Croce, A., and Bottasso, C. L.: Wind tunnel testing of a closed-loop wake deflection controller for wind farm power maximization, *Journal of Physics: Conference Series*, 753, 032 006, <https://doi.org/10.1088/1742-6596/753/3/032006>, 2016.

Campagnolo, F., Tamaro, S., Mühle, F., and Bottasso, C. L.: Wind Tunnel Testing of Combined Derating and Wake Steering, *IFAC-PapersOnLine*, 56, 8400–8405, <https://doi.org/10.1016/j.ifacol.2023.10.1034>, 22nd IFAC World Congress, 2023.

Cossu, C.: Wake redirection at higher axial induction, *Wind Energy Science*, 6, 377–388, <https://doi.org/10.5194/wes-6-377-2021>, 2021.

Doekemeijer, B. M., Kern, S., Maturu, S., Kanev, S., Salbert, B., Schreiber, J., Campagnolo, F., Bottasso, C. L., Schuler, S., Wilts, F., Neumann, T., Potenza, G., Calabretta, F., Fioretti, F., and van Wingerden, J.-W.: Field experiment for open-loop yaw-based wake steering at a commercial onshore wind farm in Italy, *Wind Energy Science*, 6, 159–176, <https://doi.org/10.5194/wes-6-159-2021>, 2021.

Downing, S. D. and Socie, D. F.: Simple Rainflow Counting Algorithms, *International Journal of Fatigue*, 4, 31–40, [https://doi.org/10.1016/0142-1123\(82\)90001-4](https://doi.org/10.1016/0142-1123(82)90001-4), 1982.

Ela, E., Gevorgian, V., Fleming, P., Zhang, Y. C., Singh, M., Muljadi, E., Scholbrook, A., Aho, J., Buckspan, A., Pao, L., Singhvi, V., Tuohy, A., Pourbeik, P., Brooks, D., and Bhatt, N.: Active Power Controls from Wind Power: Bridging the Gaps, Tech. Rep. NREL/TP-5D00-60574, National Renewable Energy Lab. (NREL), <https://doi.org/10.2172/1117060>, 2014.

710 710 Fleming, P., Aho, J., Gebraad, P., Pao, L., and Zhang, Y.: Computational fluid dynamics simulation study of active power control in wind plants, in: 2016 American Control Conference (ACC), pp. 1413–1420, <https://doi.org/10.1109/ACC.2016.7525115>, 2016.

Fleming, P., King, J., Dykes, K., Simley, E., Roadman, J., Scholbrook, A., Murphy, P., Lundquist, J. K., Moriarty, P., Fleming, K., van Dam, J., Bay, C., Mudafort, R., Lopez, H., Skopek, J., Scott, M., Ryan, B., Guernsey, C., and Brake, D.: Initial results from a field campaign of wake steering applied at a commercial wind farm – Part 1, *Wind Energy Science*, 4, 273–285, <https://doi.org/10.5194/wes-4-273-2019>, 715 2019.

Fleming, P. A., Gebraad, P. M., Lee, S., van Wingerden, J.-W., Johnson, K., Churchfield, M., Michalakes, J., Spalart, P., and Moriarty, P.: Evaluating techniques for redirecting turbine wakes using SOWFA, *Renewable Energy*, 70, 211–218, <https://doi.org/10.1016/j.renene.2014.02.015>, special issue on aerodynamics of offshore wind energy systems and wakes, 2014.

Foti, D., Yang, X., and Sotiropoulos, F.: Similarity of wake meandering for different wind turbine designs for different scales, *Journal of 720 Fluid Mechanics*, 842, 5–25, <https://doi.org/10.1017/jfm.2018.9>, 2018.

Gebraad, P. M. O., Teeuwisse, F. W., van Wingerden, J. W., Fleming, P. A., Ruben, S. D., Marden, J. R., and Pao, L. Y.: Wind plant power optimization through yaw control using a parametric model for wake effects—a CFD simulation study, *Wind Energy*, 19, 95–114, <https://doi.org/10.1002/we.1822>, 2016.

Guilloré, A., Campagnolo, F., and Bottasso, C. L.: A control-oriented load surrogate model based on sector-averaged inflow quantities: 725 capturing damage for unwaked, waked, wake-steering and curtailed wind turbines, *Journal of Physics: Conference Series*, 2767, 032019, <https://doi.org/10.1088/1742-6596/2767/3/032019>, 2024.

Heck, K., Johlas, H., and Howland, M.: Modelling the induction, thrust and power of a yaw-misaligned actuator disk, *Journal of Fluid Mechanics*, 959, A9, <https://doi.org/10.1017/jfm.2023.129>, 2023.

Jensen, N.: A note on wind generator interaction, no. 2411 in Risø-M, Risø National Laboratory, ISBN 87-550-0971-9, 1983.

730 Jiménez, A., Crespo, A., and Migoya, E.: Application of a LES technique to characterize the wake deflection of a wind turbine in yaw, *Wind Energy*, 13, 559–572, <https://doi.org/10.1002/we.380>, 2010.

Kanev, S., Savenije, F., and Engels, W.: Active wake control: An approach to optimize the lifetime operation of wind farms, *Wind Energy*, 21, 488–501, <https://doi.org/10.1002/we.2173>, 2018.

Katic, I., Højstrup, J., and Jensen, N.: A Simple Model for Cluster Efficiency, in: EWEC'86. Proceedings. Vol. 1, edited by Palz, W. and Sesto, 735 E., pp. 407–410, A. Raguzzi, european Wind Energy Association Conference and Exhibition, EWEC '86 ; Conference date: 06-10-1986 Through 08-10-1986, 1987.

King, J., Fleming, P., King, R., Martínez-Tossas, L. A., Bay, C. J., Mudafort, R., and Simley, E.: Control-oriented model for secondary effects of wake steering, *Wind Energy Science*, 6, 701–714, <https://doi.org/10.5194/wes-6-701-2021>, 2021.

Lee, S., Churchfield, M. J., Moriarty, P. J., Jonkman, J., and Michalakes, J.: A Numerical Study of Atmospheric and Wake Turbulence Impacts 740 on Wind Turbine Fatigue Loadings, *Journal of Solar Energy Engineering*, 135, <https://doi.org/10.1115/1.4023319>, 031001, 2013.

Liew, J. and Larsen, G. C.: How does the quantity, resolution, and scaling of turbulence boxes affect aeroelastic simulation convergence?, *Journal of Physics: Conference Series*, 2265, 032049, <https://doi.org/10.1088/1742-6596/2265/3/032049>, 2022.

Liew, J., Urbán, A. M., and Andersen, S. J.: Analytical model for the power–yaw sensitivity of wind turbines operating in full wake, *Wind Energy Science*, 5, 427–437, <https://doi.org/10.5194/wes-5-427-2020>, 2020.

745 Liu, Y., Wang, Y., Wang, X., Zhu, J., and Lio, W. H.: Active Power Dispatch for Supporting Grid Frequency Regulation in Wind Farms
Considering Fatigue Load, *Energies*, 12, <https://doi.org/10.3390/en12081508>, 2019.

Medici, D. and Alfredsson, P. H.: Measurements on a wind turbine wake: 3D effects and bluff body vortex shedding, *Wind Energy*, 9, 219–236, <https://doi.org/10.1002/we.156>, 2006.

Meyer Forsting, A. R., Pirrung, G. R., and Ramos-García, N.: A vortex-based tip/smearing correction for the actuator line, *Wind Energy Science*, 4, 369–383, <https://doi.org/10.5194/wes-4-369-2019>, 2019.

Meyers, J., Bottasso, C., Dykes, K., Fleming, P., Gebraad, P., Giebel, G., Göçmen, T., and van Wingerden, J.-W.: Wind farm flow control: prospects and challenges, *Wind Energy Science*, 7, 2271–2306, <https://doi.org/10.5194/wes-7-2271-2022>, 2022.

Nocedal, J. and Wright, S.: Numerical optimization, Springer series in operations research and financial engineering, Springer, New York, NY, 2. ed. edn., ISBN 978-0-387-30303-1, http://gso.gbv.de/DB=2.1/CMD?ACT=SRCHA&SRT=YOP&IKT=1016&TRM=ppn+502988711&sourceid=fbw_bibsonomy, 2006.

NREL: floris v3, <https://github.com/NREL/floris>, <https://doi.org/10.5281/zenodo.6320841>, 2023.

Oudich, Y., Gyselinck, J., De Belie, F., and Kinnaert, M.: Providing power reserve for secondary grid frequency regulation of offshore wind farms through yaw control, *Wind Energy*, 26, 850–873, <https://doi.org/10.1002/we.2845>, 2023.

Sagaut, P.: Large Eddy Simulation for Incompressible Flows: An Introduction, Springer, Berlin, <https://doi.org/10.1007/b137536>, 2006.

760 Salavatidezfouli, S., Sheidani, A., Bakhshaei, K., Safari, A., Hajisharifi, A., Stabile, G., and Rozza, G.: Modal Analysis of the Wake Shed Behind a Horizontal Axis Wind Turbine with Flexible Blades, *Journal of Applied and Computational Mechanics*, 11, 826–843, <https://doi.org/10.22055/jacm.2024.47176.4670>, 2025.

Schottler, J., Mühle, F., Bartl, J., Peinke, J., Adaramola, M. S., Sætran, L., and Hölling, M.: Comparative study on the wake deflection behind yawed wind turbine models, *Journal of Physics: Conference Series*, 854, 012 032, <https://doi.org/10.1088/1742-6596/854/1/012032>, 2017.

765 Shapiro, C. R., Bauweraerts, P., Meyers, J., Meneveau, C., and Gayme, D. F.: Model-based receding horizon control of wind farms for secondary frequency regulation, *Wind Energy*, 20, 1261–1275, <https://doi.org/10.1002/we.2093>, 2017.

Silva, J. G., Doekemeijer, B. M., Ferrari, R., and van Wingerden, J.-W.: Active power control of wind farms: an instantaneous approach on waked conditions, *Journal of Physics: Conference Series*, 2265, 022 056, <https://doi.org/10.1088/1742-6596/2265/2/022056>, 2022.

770 Sinner, M., Spyrou, E., Bay, C. J., King, J., and Corbus, D.: Coordinated wind power plant and battery control for active power services, *Journal of Renewable and Sustainable Energy*, 15, 053 304, <https://doi.org/10.1063/5.0156464>, 2023.

Starke, G. M., Meneveau, C., King, J., and Gayme, D. F.: Yaw-Augmented Control for Wind Farm Power Tracking, in: 2023 American Control Conference (ACC), pp. 184–191, <https://doi.org/10.23919/ACC55779.2023.10156444>, 2023.

Tamaro, S. and Bottasso, C. L.: A New Wind Farm Active Power Control Strategy to Boost Tracking Margins in High-demand Scenarios, *775 in: 2023 American Control Conference (ACC)*, pp. 192–197, <https://doi.org/10.23919/ACC55779.2023.10156275>, 2023.

Tamaro, S., Campagnolo, F., and Bottasso, C. L.: On the power and control of a misaligned rotor – beyond the cosine law, *Wind Energy Science*, 9, 1547–1575, <https://doi.org/10.5194/wes-9-1547-2024>, 2024a.

Tamaro, S., Guilloré, A., Anand, A., Mühle, F. V., Campagnolo, F., and Bottasso, C. L.: Validation of induction/steering reserve-boosting active power control by a wind tunnel experiment with dynamic wind direction changes, *Journal of Physics: Conference Series*, 2767, 092 067, <https://doi.org/10.1088/1742-6596/2767/9/092067>, 2024b.

780 The MathWorks Inc.: MATLAB version: 9.13.0 (R2022b), 2022.

Thomas, J. J., Baker, N. F., Malisani, P., Quaeghebeur, E., Sanchez Perez-Moreno, S., Jasa, J., Bay, C., Tilli, F., Bieniek, D., Robinson, N., Stanley, A. P. J., Holt, W., and Ning, A.: A comparison of eight optimization methods applied to a wind farm layout optimization problem, *Wind Energy Science*, 8, 865–891, <https://doi.org/10.5194/wes-8-865-2023>, 2023.

785 Troldborg, N., Sørensen, J. N., and Mikkelsen, R.: Actuator Line Simulation of Wake of Wind Turbine Operating in Turbulent Inflow, *Journal of Physics: Conference Series*, 75, 012063, <https://doi.org/10.1088/1742-6596/75/1/012063>, 2007.

Vali, M., Petrović, V., Steinfeld, G., Y. Pao, L., and Kühn, M.: An active power control approach for wake-induced load alleviation in a fully developed wind farm boundary layer, *Wind Energy Science*, 4, 139–161, <https://doi.org/10.5194/wes-4-139-2019>, 2019.

van Kuik, G., Peinke, J., Nijssen, R., Lekou, D., Mann, J., Sørensen, J., Ferreira, C., van Wingerden, J., Schlipf, D., Gebraad, P., Polinder, H.,
790 Abrahamsen, A., van Bussel, G., Sørensen, J., Tavner, P., Bottasso, C., Muskulus, M., Matha, D., Lindeboom, H., Degraer, S., Kramer, O., Lehnhoff, S., Sonnenschein, M., Sørensen, P., Künneke, R., Morthorst, P., and Skytte, K.: Long-term research challenges in wind energy – a research agenda by the European Academy of Wind Energy, *Wind Energy Science*, 1, 1–39, <https://doi.org/10.5194/wes-1-1-2016>, 2016.

van Wingerden, J.-W., Pao, L., Aho, J., and Fleming, P.: Active Power Control of Waked Wind Farms, *IFAC-PapersOnLine*, 50, 4484–4491,
795 <https://doi.org/10.1016/j.ifacol.2017.08.378>, 20th IFAC World Congress, 2017.

Vermeer, L., Sørensen, J., and Crespo, A.: Wind turbine wake aerodynamics, *Progress in Aerospace Sciences*, 39, 467–510, [https://doi.org/10.1016/S0376-0421\(03\)00078-2](https://doi.org/10.1016/S0376-0421(03)00078-2), 2003.

Vollmer, L., Steinfeld, G., Heinemann, D., and Kühn, M.: Estimating the wake deflection downstream of a wind turbine in different atmospheric stabilities: an LES study, *Wind Energy Science*, 1, 129–141, <https://doi.org/10.5194/wes-1-129-2016>, 2016.

800 Wang, C.: Numerical simulation of wind farm control, Ph.D. thesis, Technische Universität München, <https://mediatum.ub.tum.de/1578640>, 2021.

Wang, J., Wang, C., Campagnolo, F., and Bottasso, C. L.: Wake behavior and control: comparison of LES simulations and wind tunnel measurements, *Wind Energy Science*, 4, 71–88, <https://doi.org/10.5194/wes-4-71-2019>, 2019.