

Load assessment of a wind farm considering negative and positive yaw misalignment for wake steering

Regis Thedin¹, Garrett Barter¹, Jason Jonkman¹, Rafael Mudafort¹, Chris Bay¹, Kelsey Shaler², and Jasper Kreeft³

¹National Renewable Energy Laboratory, Golden, CO, USA

²Shell International Exploration and Production, Houston, TX, USA

³Shell Global Solutions International B.V., The Hague, Netherlands

Correspondence: Regis Thedin (regis.thedin@nrel.gov)

Abstract. Wake steering strategies are employed to increase the overall power production of wind farms by deflecting wakes of upstream turbines away from downstream ones. The gain in net power comes at the expense of increased fatigue loads experienced by downstream turbines. In this work we investigate performance and fatigue loading characteristics of a small farm consisting of five aligned International Energy Agency Wind Technology Collaboration Programme 15 MW wind turbines. A parametric study is performed where, for every wind direction from -20 to 20 degrees, the yaw misalignment angle varies from -25 to 25 degrees. This setup allows us to investigate asymmetries and identify optimal conditions for a given wind direction. In general, we find that positive yaw configurations are preferred and that yaw configurations that result in attractive power differences when compared to a baseline no-yaw scenario (25%) come with significant increase in fatigue loading (we use standard deviation and damage-equivalent load (DEL) of the blade-root, low-speed shaft, and tower-base moments as proxies for fatigue load). We find that for any given positive wind inflow angle, yaw angles between -2.5 and 15 degrees yield power differences of 10–20% over a no-yaw baseline, and positive yaw is preferred because of lower fatigue loading. For any given negative wind inflow angles, positive yaw also results in lower magnitudes of standard deviation and DEL for the channels investigated. A small power loss of up to 2% is observed for some positive yaw angles under negative wind directions (as compared to symmetric negative yaw and positive wind cases), but improvements in terms of loads exceed 25% and may be enough to justify a positive yaw configuration under negative winds as well. We show that such behavior can be explained by partial waking and the direction of rotation of the rotor.

1 Introduction

Deflecting or dissipating wind turbine wakes within a larger array is a recognized strategy to increase wind farm energy yields. In the case of wake steering, upstream turbines are yaw-misaligned with the incoming flow such that their wakes are deflected away from downstream turbines. While the upstream turbine produces slightly less power than it would otherwise, the downstream turbines see higher inflow velocity and the net power production of the array increases. Exploring the benefits and operational dynamics of wake steering has been an active field of research over the past 10–15 years via simulation (Medici and Dahlberg, 2003; Jiménez et al., 2010; Fleming et al., 2014; Gebraad et al., 2016), wind tunnel experiments (Bastankhah

and Porté-Agel, 2019; Adaramola and Krogstad, 2011; Medici and Dahlberg, 2003; Medici and Alfredsson, 2006; Wagenaar et al., 2012), and field test trials (Fleming et al., 2017, 2019, 2020; Doekemeijer et al., 2020; Simley et al., 2021). Houck (2021) offers a nice review of the literature.

In their simplest form, wind farm controllers might dial-in optimal yaw settings to every turbine in the array based on a precomputed library of inflow conditions. With more sophistication, farm controllers might make collective adjustments based on real-time observations and turbine condition monitoring. Regardless of its level of sophistication, the farm controller will have to account for a fundamental asymmetry in wake deflection physics. The asymmetries are primarily due to the clockwise rotation standard of the rotor (when viewed “head-on”), but also likely a result of the atmospheric conditions (e.g., veer due to Coriolis forces). This makes positive-angle yaw deflections (counterclockwise when viewed from above) more effective than negative-angle displacements (clockwise from above). The physical mechanisms of this asymmetry are well explained by Bastankhah and Porté-Agel (2016) and Hulsman et al. (2022). A diagram of the conventions used is shown in Fig. 1.

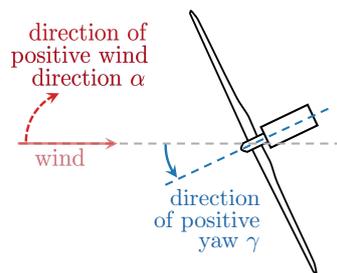


Figure 1. Wind turbine viewed from above, showing the conventions used. A positively yawed turbine is shown under aligned (zero degrees) wind.

That being said, the negative-angle yaw settings could still be used for inflow scenarios where the wind direction alignment relative to the turbine array creates partial wakes in the direction of negative yaw misalignment, so additional negative yaw deflection is the optimal approach.

As wake steering reaches commercial maturity, it is critical that its impact on operational loading is well understood so that farm control algorithms can be devised that maximize energy production while limiting turbine structural loads. The impact of wake steering on turbine loading has been, and continues to be, an important focus because intentional yaw misalignments and partial waking are not part of the turbine design standards. Furthermore, turbine manufacturers, project developers, and wind farm owner-operators will be wary of adopting wake steering strategies if they lead to increased maintenance costs or shorter turbine lifetimes.

The consensus in the literature is that for most load quantities of interest, the loads on the front-row turbines decrease with wake steering while the loads on the downstream turbines increase. This is due to the decrease in power production and rotor thrust in the front-row turbines, with a corresponding increase in the downstream turbines, relative to the no-wake steering scenario (Kanev et al., 2020; Shaler et al., 2022). This brings loading values into a more consistent level across the array, which could have maintenance planning and cost advantages. While this might be a general trend, there are exceptions, such

as tower or monopile yaw moments, which increase with any yaw misalignment. Because wake steering is generally applied
50 below rated wind speeds, any increase in loading is unlikely to approach the ultimate load operational envelope of the turbine,
but fatigue loading and turbine lifetime is of concern. The impact on fatigue and turbine lifetime will depend on the frequency
of wake steering (Shaler et al., 2022). There is also consensus that imperfect wake steering that exacerbates the presence of
partial waking could be even more harmful than not employing wake steering at all (Ciri et al., 2018; Stanley et al., 2020).

In terms of positive vs. negative yaw angles, one might think that turbine load behavior would mimic that of power pro-
55 duction, meaning that turbine loads (whether measured by mean loading, fatigue loading, or ultimate loading) would show
similar trends regardless of the direction of yaw, with some allowance for slight asymmetry due to the clockwise rotation of the
rotor. However, prior work suggests that this might not always be the case. Damiani et al. (2018) compared elastic simulations
with in situ instrumentation of a General Electric (GE) 1.5sle 1.5 MW wind turbine and noted asymmetries related to positive
and negative yaw directions. Ennis et al. (2018) similarly used field tests of a three-turbine array to demonstrate preferential
60 direction for yaw misalignments for blade flap and edge loading. Zalkind and Pao (2016) also reported asymmetric trends (with
respect to the yaw misalignment angle) in a simulation-based study of the National Renewable Energy Laboratory’s 5 MW
reference wind turbine. In these studies, some load statistics showed monotonically increasing or decreasing behavior when
sweeping from negative to positive yaw offset angles.

This work intends to understand the load asymmetries across positive vs. negative yaw angles to arrive at operational rec-
65 ommendations for wind farm owners. We learn from the range of conditions examined by Shaler et al. (2022) to focus our
computational simulations on a five-turbine array with a systematic exploration of incoming wind direction and yaw angle at
set values for freestream velocity, shear profile, and turbulence intensity. We ensure that the modeling fidelity and verification
are sufficient to capture the curled (“kidney bean shaped”) wake physics that are important to the asymmetries. We indeed
observe asymmetric loading and arrive at recommendations based on the “quadrants” of positive and negative incoming wind
70 direction angles (relative to the array alignment) compared to positive and negative yaw misalignments for wake steering.

2 Methodology

In this work, we investigate power performance and load response of a small wind farm subjected to positive and negative
yaw misalignments. A parametric study of different yaw angles under different wind directions is performed using the mid-
fidelity wind plant performance tool, FAST.Farm. A preprocessing calibration step is taken to ensure the FAST.Farm wake
75 characteristics, including its curled shape, match those obtained using high-fidelity large-eddy simulation (LES) results.

2.1 Computational models

The high-fidelity LES are executed using ExaWind’s AMR-Wind code (Sprague et al., 2020). AMR-Wind is a massively
parallel, incompressible flow solver tailored for wind energy applications. The primary application of AMR-Wind in this work
is to perform LES of an atmospheric boundary layer, and simulate turbine wakes using an actuator line model for the turbine
80 representation. AMR-Wind is only used for the calibration step, which will be described in Section 3.

After calibration, the majority of results in this work are obtained using FAST.Farm. FAST.Farm is a multiphysics engineering tool for predicting the power performance and structural load response of wind turbines within a wind farm. For each turbine, FAST.Farm uses individual OpenFAST instances to solve for the aero-hydro-servo-elastic dynamics. Information about the ambient wind are given as inputs to FAST.Farm—that is, it does not fully resolve all fluid flow equations. Wake deficits, advection, deflection, meandering, and merging are all handled by dynamic wake meandering models within FAST.Farm. The inflow conditions are given in terms of small, high-resolution boxes around each turbine, and one large, lower-resolution box around the entire farm. The high-resolution boxes are responsible for providing inflow conditions to OpenFAST for aeroelastic analysis, while the low-resolution box meanders the wake. Due to its computational cost and level of accuracy, FAST.Farm is considered an engineering-fidelity tool to support detailed design and loads analysis of wind turbines in a farm context. More details of FAST.Farm can be found in Jonkman et al. (2017).

The approach of using an engineering-fidelity tool (after a proper calibration against high-fidelity LES) is due to computational cost and speed. A single LES case executed using AMR-Wind and used for calibration takes up to about 2.5 days on a high-performance computing environment, using around 2,800 cores. That does not include a precursor simulation, which takes an additional 1 to 1.5 days on around 750 cores. Similarly, each FAST.Farm case executed in this work takes about 3 to 4 hours on five cores only. The calibration step is important to ensure the wake characteristics match those of the LES for the same condition; they were, however, the most computationally expensive step of this work. The calibration performed here allows us to reuse the values obtained in future studies.

2.2 Case study setup

For this study we devise a small wind farm consisting of five aligned International Energy Agency Wind Technology Collaboration Programme (IEA Wind) 15 MW reference wind turbines (Gaertner et al., 2020). All OpenFAST instances of turbines within FAST.Farm in this work use the same controller, based on the reference open-source controller ROSCO (Abbas et al., 2022). The turbines are spaced by 1 nautical mile, or approximately 7.7 rotor diameters, inspired by the projects in active development in Massachusetts waters. The goal is to perform a parametric study varying the wind direction and the yaw angle of the five turbines in the following fashion:

- Vary wind direction, α , from -10 to 10 degrees, in increments of 2 degrees.
- For each wind direction α , vary the reference yaw misalignment angle, γ , from -25 to 25 degrees in increments of 2.5 degrees. For each value of the reference γ , the yaw misalignment for the five turbines is such that $\gamma_{T1} = \gamma_{T2} = \gamma_{T3} = \gamma$, $\gamma_{T4} = \gamma/2$, and $\gamma_{T5} = 0$.

The discretizations of the wind direction and reference yaw misalignment angle were chosen such that they provide a balance between computational resources (given total number of cases) and results that show smooth trends, meaning no local minima or maxima in between angles. We limit the yaw misalignment to a magnitude of 25 degrees due to the higher structural loads that such conditions impose (Damiani et al., 2018; Shaler et al., 2022). The yaw angles for each of the five turbines are considered to be realistic, such as zero yaw angle for the last turbine in the row, based on prior literature where the yaw angles

of a row of five turbines have been investigated (Bastankhah and Porté-Agel, 2019). We keep the same magnitudes for positive
 115 and negative yaw to allow direct comparison, rather than using optimized angles for each scenario.

In the figures presented in this work, we refer to the different yaw cases simply by the reference yaw γ . Given the farm
 layout, instead of generating new wind fields for each wind direction, we opt to rotate the farm. This approach ensures that
 the wind fields for the leading turbine are the exact same in all cases, resulting in a more consistent turbine response across
 different scenarios (some quantities will be normalized by the leading turbine). An illustration of the complete setup, including
 120 the farm layout for different wind directions, is given in Fig. 2.

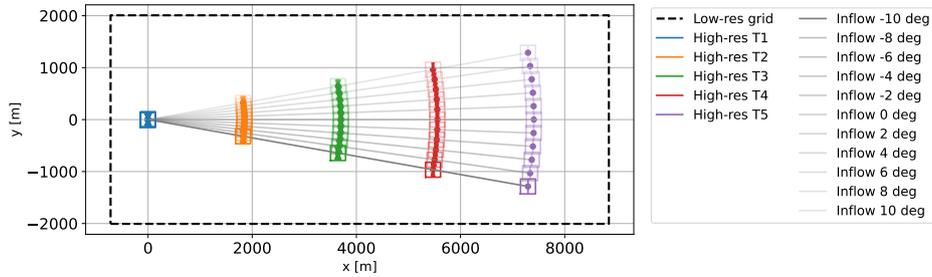


Figure 2. Illustration of the farm layout setup for FAST.Farm. The turbulence boxes used within FAST.Farm are shown as the low-resolution
 and all the high-resolution domains. Each gray line represents one inflow direction. All the yaw positions considered for each turbine are
 shown; note that the last turbine is always at zero yaw.

The layout shown in Fig. 2 includes all the cases at once. Each case, consisting of a certain wind direction α and yaw γ can
 be obtained through a trivial change in reference frames. Figure 3 demonstrates the change in reference frame. The flow field
 figures presented in this paper are shown in the style of Fig. 3(b).

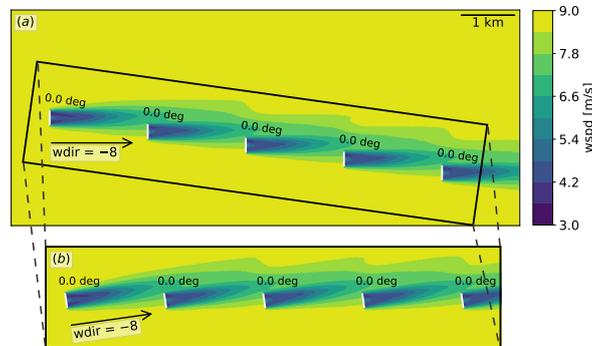


Figure 3. Change in reference frame to use the same wind fields for different wind directions. (a) Following reference of Fig. 2, a solution at
 wind direction $\alpha = -8$ and yaw angles $\gamma = 0$; (b) Rotated reference to a more intuitive layout given wind direction and yaw angles. All the
 flow field pictures presented here will be shown using the rotated frame of reference shown in (b).

2.3 Inflow generation

125 The focus of this study is on a single wind condition. Our interests lie in the differences observed in terms of power and loads for different wind directions and different yaw misalignment. The IEA Wind 15 MW reference wind turbine has a rated wind speed of 10.6 m/s; hence, we have selected a wind speed of 8.6 m/s, representing a below-rated, “Region II” operating point.

Turbulent wind fields are complex and difficult to model accurately. LES tools are able to resolve the larger eddies up to the grid resolution while the smaller, subgrid eddies are modeled. Such simulations require intensive computational resources.
130 The LES resolve a wide range of physics present in an atmospheric boundary layer, including temperature stratification at the capping inversion, veer effects due to the Coriolis force, and others, and are considered the high-fidelity “truth” in this work. Engineering-fidelity models of the wind inflow are available where the wind fields have the appropriate second-order statistics but are not able to recreate true turbulence eddies, their coherent structure, or more complex physics. One of the main underlying assumptions of these models is that buoyancy is not considered, leading to significant inaccuracies when used for
135 non-neutral scenarios.

For the calibration of FAST.Farm, we generated an inflow velocity field using the AMR-Wind LES model. That means the background inflow is exactly the same; naturally, the leading turbine, which is undisturbed by wakes, is subject to the exact same flow field. A key approximation used in FAST.Farm is the handling of the lateral component of the wind. The lateral component varies with height and is present in the LES¹. In FAST.Farm, the rotor-mean value, which is close to zero
140 (see Fig. 4), is used to convect and meander the wake to either positive or negative y direction. Because the passive-tracer calculation used by FAST.Farm to meander the wakes involves spatially averaging the transverse velocity components, the effect of veer on the wake gets averaged out, effectively neglecting veer in the wake calculation; such a scenario is typical of locations near the equator. (Wind veer is still accounted for in OpenFAST in the load calculation.)

Using the same flow field from the LES in FAST.Farm allows a direct comparison between the two codes and ensures that
145 the parameters for the dynamic wake meandering model within FAST.Farm are as accurate as possible for the conditions under investigation. As noted above, we expect asymmetries in the LES results of positive and negative yaw due to the veer that will not be present in FAST.Farm. After the calibration step, for the parametric study of turbine loads based on inflow wind direction and yaw angles, we model the flow fields using TurbSim (Kelley and Jonkman, 2005), a synthetic turbulence generator based on the Kaimal spectrum model with Davenport’s exponential coherence model, which assumes neutral atmospheric stability.
150 We tuned the TurbSim inputs to mimic the LES turbulence statistics. While the flow fields generated using TurbSim contain time-varying v (a result of the coherence model applied and needed for proper meandering of the wake), the mean value of v is zero, removing the veer in the flow altogether.

In generating the LES inflow, we specify our desired wind speed (8.6 m/s) at the desired height (150 m), as well as parameters such as ground heat flux (0 K-m/s for neutral stability), aerodynamic surface roughness (0.75 m), capping inversion height

¹In the LES, a background forcing (pressure gradient) is applied so that it drives the flow to match a user-specified velocity vector at a certain height. Typically, and as done in this work, the height chosen is the hub height. That means that the chosen velocity vector (8.6, 0, 0) m/s is matched at 150 m, as evidenced by the vertical profiles presented in Fig. 4. Because of this approach, the temporal mean v component of the wind is, by definition, zero at 150 m, and the rotor-mean value is close to zero.

155 (750 m), and strength (10 K across 100 m). Our domain is large to minimize deep array effects once we add the turbines, coming to an extent of $10 \times 7.5 \times 1$ km. The turbulence intensity and shear exponents are not inputs but rather a result of the combination of physical drivers, and they require an iterative process (varying the aforementioned parameters) to achieve target statistics. The values given are the final ones used such that we obtained our target quantities. Our target quantities include a shear exponent of 0.2 and a mean TI of 10% at the hub height. Such high shear value is an attempt to model the wind profile of a stable boundary layer, while maintaining statistics of a neutral boundary layer. To obtain this shear exponent, a relatively high value for the aerodynamic surface roughness was needed, and we find that the combination of parameters given above with an aerodynamic surface roughness value of 0.75 m resulted in the shear profile we aimed for. For the precursor part, the simulation is executed for 1,800 s after it reaches a quasi-stationary state. The simulation has a 10 m uniform spatial resolution and a fixed time step of 0.2 s, resulting in a maximum CFL number of approximately 0.5. The one equation for kinetic energy turbulence model is used. The resulting atmospheric boundary layer for the period of interest has a mean shear exponent of 0.196, a mean TI of 10%, and a friction velocity of 0.54 m/s. A low nonzero veer is obtained at 0.011 degrees/m. The final vertical profiles of the boundary layer are shown in Fig. 4.

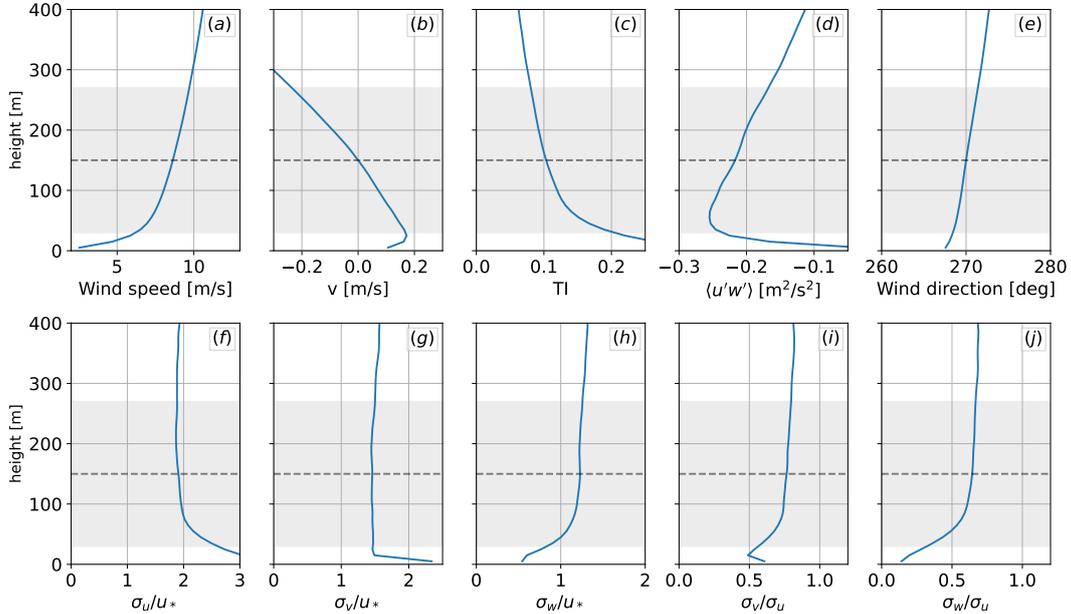


Figure 4. Planar- and time-averaged vertical profiles of several quantities of a neutral atmosphere simulated using LES. Rotor-swept zone shown for reference with gray shading. (a) Horizontal wind speed; (b) lateral component of the wind, v ; (c) turbulence intensity; (d) vertical flux of the streamwise component of the wind; (e) wind direction, where 270 degrees is westerly winds (left to right in Fig. 2); (f,g,h) variance of the streamwise, lateral, and vertical component of the wind normalized by the friction velocity, respectively; (i,j) ratio of the variance of the lateral and vertical components of the wind, respectively, to the streamwise component.

The precursor simulation with the inflow presented in Fig. 4 is then used on subsequent simulations that model the IEA Wind 15 MW reference wind turbine as actuator lines. For this simulation, two refinement boxes are used near the turbines, resulting

170 in a 2.5 m near-turbine spatial resolution. A temporal resolution of 0.2 s is used, which is limited by the actuator point present at the tip of the blade, ensuring proper use of the actuator line model. These simulations are then used to tune the FAST.Farm model, as will be discussed in Section 3.

The canonical condition obtained through LES is also specified within TurbSim after the calibration step for the batch runs. That includes the mean wind speed vector (8.6, 0, 0) m/s at 150 m, coherence in all three components of the wind, 10% 175 turbulence intensity, and 0.2 shear exponent. The lack of veer is purposeful due to the approximation within FAST.Farm. We acknowledge that while using TurbSim is appropriate for the conditions of interest of this work, it would incur larger model errors if the condition of interest included non-neutral atmospheric stratification.

2.4 Quantities of interest

Wind farm performance and loads are the main interest in this work. With the power gains due to wake steering, the turbine 180 components are typically subjected to increased fatigue loads. The selected load channels that will be investigated are outlined in Table 1 and processed in terms of quantities used as a proxy for fatigue loads.

Table 1. Description of the load channels of interest.

Channel	Description
Blade root bending moment	Magnitude of the combined blade root out-of-plane bending moment (flapwise) and in-plane bending moment (edgewise)
Low-speed shaft (LSS) bending moment	Magnitude of the combined low-speed shaft bending moment in the two transverse directions at the rotating hub
Tower base bending moment	Magnitude of the combined fore-aft and side-to-side bending moment at the tower base
Tower base torsional moment	Magnitude of the torsional moment at the tower base

Fatigue is the damage that accumulates in structural components over time due to cyclic loadings. Such damaging effect can be modeled using the S-N curve, which indicates the amount of cycles a structure can endure until it fails for a given cyclic stress. The slope of this curve is the Wöhler exponent. We can frame the concept of damage related to fatigue loads in terms 185 of an equivalent fatigue load. A load of equivalent fatigue damage, commonly referred to as damage equivalent load (DEL) is a constant-amplitude load range that would, over the same number of cycles, cause an equivalent amount of damage as the original load series of variable amplitude.

To compute DELs, we use rainflow counting to extract the ranges and frequencies of the different load cycles. To be able to use our signals to compute DEL, we also need to correct it for the mean loads, which is accomplished by applying the Goodman correction. In other words, the Goodman correction takes into account the mean loads to provide a more accurate calculation of the DELs. While the results shown here use the Goodman correction, we note that such correction has not changed the results in any meaningful way, as the loading cycles resulted in stresses that were significantly lower than the ultimate stresses of the material of each component of the turbine. That is to be expected since the turbine is operating far from the edges of its operational envelope. The Goodman-corrected DELs are then used to directly compare the different simulation scenarios. No lifetime fatigue calculations are shown because it is not applicable to the single condition under investigation.

As noted in Table 1, the magnitudes of two orthogonal signals are used for the tower-base, low-speed shaft (LSS), and blade-root bending moments. We compute load roses for each pair of these signals and come up with single signals for each sector around the load rose. We use sectors of 10 degrees and obtain the time series. The entire workflow of obtaining both the standard deviation and DEL are then executed for each sector, and the highest value is taken as the standard deviation or DEL of that signal. This approach is taken for three pairs of signals: tower-base fore-aft and side-to-side bending moments; blade-root in-plane and out-of-plane bending moments (flapwise and edgewise); and LSS bending moment across the two orthogonal directions.

The short-term DELs, referred to as simply DELs, are computed as

$$\text{DEL} = \left(\frac{1}{T} \sum_k n_k L_k^m \right)^{1/m} \quad (1)$$

where n_k is the number of cycles of some load L_k related to load bin k (after applying the Goodman correction), m is the slope of the Whöhler curve of the material, and T is the total length of a given simulation representing a 1-Hz-equivalent cycle (that is, T is the number of 1-Hz cycles over the simulation). The bins k are obtained using the load ranges obtained using rainflow counting of the Goodman-corrected time series of stresses. We first convert loads to stresses using material properties and then apply the Goodman correction on the stresses. The resulting stresses are converted back to loads and used in the DEL calculations.

In our load analysis, we assume axial stresses only, with the bending moments being the contributing load toward this axial stress in each turbine component. We also assume the geometry of the blade root is of uniform circular cross section. We do consider material ultimate strength separately for each turbine component and pick the worse value (from tension and compression), yielding possibly conservative results. We do not perform further analysis to take the anisotropy of composite materials into account for our load analysis. The ultimate stresses used to correct for mean loads using Goodman were 1,047 MPa for the composite blade root, 814 MPa for ASTM 4140 drivetrain steel used for the LSS, and 450 MPa for ASTM A572 steel used for the tower. We have used a Wöhler exponent of 4 for steel components and 10 for the composite blade.

3 Calibration of the curled wake model

To ensure that wake meandering characteristics, and thus loading of downstream turbines, are as accurate as they can be when running analysis using FAST.Farm, we first perform a calibration against the LES solutions. The calibration encompasses varying key parameters in the wake models and compares the resulting quantities with LES. The LES runs included the turbines using the actuator line model, and three yaw conditions were executed: positive, negative, and no yaw. The positive yaw case had a reference yaw $\gamma = 20$ degrees, (meaning $\gamma_{T1} = \gamma_{T2} = \gamma_{T3} = 20$ degrees, $\gamma_{T4} = 10$ degrees, and $\gamma_{T5} = 0$). The negative yaw had the same magnitudes (with the sign flipped), and all turbines were aligned (zero yaw) for the no-yaw case. Note that these simulations include the full physics considered and resolved by the LES.

We performed the workflow described in this section for both the polar wake and the curled wake models of FAST.Farm. The polar wake model is a simpler axisymmetric model, and two constants were used for tuning. The constants are related to the wake propagation angle and initial offset that are proportional to yaw misalignment (respectively called `C_HWkDf1_xY` and `C_HWkDf1_OY` in FAST.Farm). The curled wake model (Martínez-Tossas et al., 2021; Branlard et al., 2023) is a more realistic model that takes into consideration the curled vorticity present in skewed flow and the rotor rotation, yielding a curled wake that resembles a kidney bean shape. For the curled wake model, two additional constants are considered in the calibration: one related to the decay rate of the spanwise velocity components (`k_VortexDecay` in FAST.Farm), and another related to the scaling of the eddy viscosity, which effectively controls the amount of diffusion in the model (`k_vCurl` in FAST.Farm). For our cases of interest, we perform sweeps on all of these constants for the IEA Wind 15 MW reference wind turbine under neutrally stratified inflow generated in LES (described above). Our approach consisted of setting a series of sweeps, analyzing the results, and setting up another series of sweeps around the best-performing values, narrowing the range of values closer to the optimal ones observed in prior sweeps. In this way, we effectively avoid local minima and increase the resolution of the final parameters without having to execute an unnecessarily large number of cases. A table showing all sweeps and ranges of each parameter for the curled wake calibration is shown in Table 2. A total of 4,500 FAST.Farm simulations were executed and analyzed to come up with the optimal parameters for the curled wake model that are used in the rest of this work. A similar set of cases were used to calibrate the polar wake model.

The quantities of generated power and lateral wake center behind each turbine are used for the comparison and to determine which combination of wake parameters performed best. Transverse, vertical planes of inflow are saved at 2, 3, 4, 5, and 6 diameters downstream of each turbine for each of the three LES cases and each FAST.Farm case shown in Table 2. From those planes, wake centers are identified using a one-dimensional Gaussian fit (Quon et al., 2020). For illustration purposes, an example of a time series between LES and FAST.Farm for the no-yaw case is shown in Fig. 5. Note how, as expected, the error due to the dynamic wake meandering model increases with increasing downstream distances.

In contrast to prior studies (e.g. Shaler and Jonkman (2021)), we devised a simple wake center error metric that takes into account time-accuracy and not just mean wake center. We compare LES and FAST.Farm at every second and then get the mean error. The metric that was minimized is given by $\Delta\bar{y} = \overline{|y_{FF} - y_{LES}|}/D$ where y corresponds to the time series of the lateral location of the wake and D is the rotor diameter. A similar power metric is used, $\Delta\bar{P} = |\overline{P}_{FF} - \overline{P}_{LES}|/|\overline{P}_{LES}|$.

Table 2. Range of the sweeps of curled wake modeling parameters within FAST.Farm for calibration procedure. Each case was executed under three different yaw conditions and compared with its respective LES solution.

	Sweep 1		Sweep 2		Sweep 3		Sweep 4	
	min	max	min	max	min	max	min	max
C_HWkDfl_OY [m/deg]	0.1	0.5	0.0	0.3	0	0.3	0	0
C_HWkDfl_xY [deg ⁻¹]	-0.008	-0.002	-0.006	0	-0.006	0	0	0
k_VortexDecay [-]	0	1	0	0.4	0	0.4	0	0.04
k_vCurl [-]	1	5	1	5	1	5	1	5
Number of cases	1260		1134		1944		162	

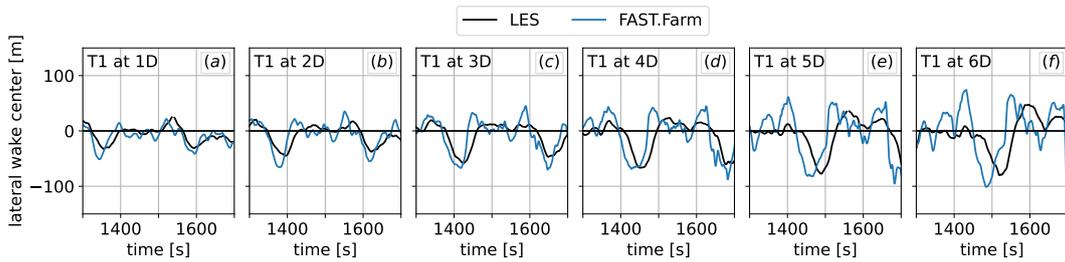


Figure 5. Illustration of the time series for the lateral wake center for both FAST.Farm and LES. The FAST.Farm data shown are related to optimal values for the polar wake model.

Here, we skip much of the details of earlier sweeps and provide two intermediate results to give the reader a sense of the quantities and cases compared. For the third sweep indicated in Table 2, we can plot the four-dimensional wake center difference $\Delta\bar{y}$ dataset as shown in Fig. 6. Note how for every combination of the four parameters, a zero-valued C_HWkDfl_xY is consistently better (darker shades within each subplot). Also, for each subplot, minimal variation of C_HWkDfl_OY is observed. From that, we picked our two wake deflection parameters to be zero on the fourth sweep. Additionally, observing each row, results improved as we move toward the bottom rows (lower k_vortexDecay), suggesting we should focus on lower values for this variable.

There is a balance between choosing the case that yields the best performance in terms of wake center and the case that is best for matching power. In Fig. 6, little variation is observed across the different columns of subplots (varying k_vCurl). However, if we create the same plots for the power metric, it is clear an optimal k_vCurl exists—see Fig. 7.

Figures 6 and 7, however, only show a single yaw case. Another balance has to be considered across all three yaw scenarios, since this is the goal of this study. On Sweep 4, we can reduce the number of variables we are varying and make ensemble-averaged plots, considering three downstream planes of the first turbine (at 4D, 5D, and 6D) and average them together. These average plots are shown in Fig. 8, representing the fourth and final sweep considered in this work.

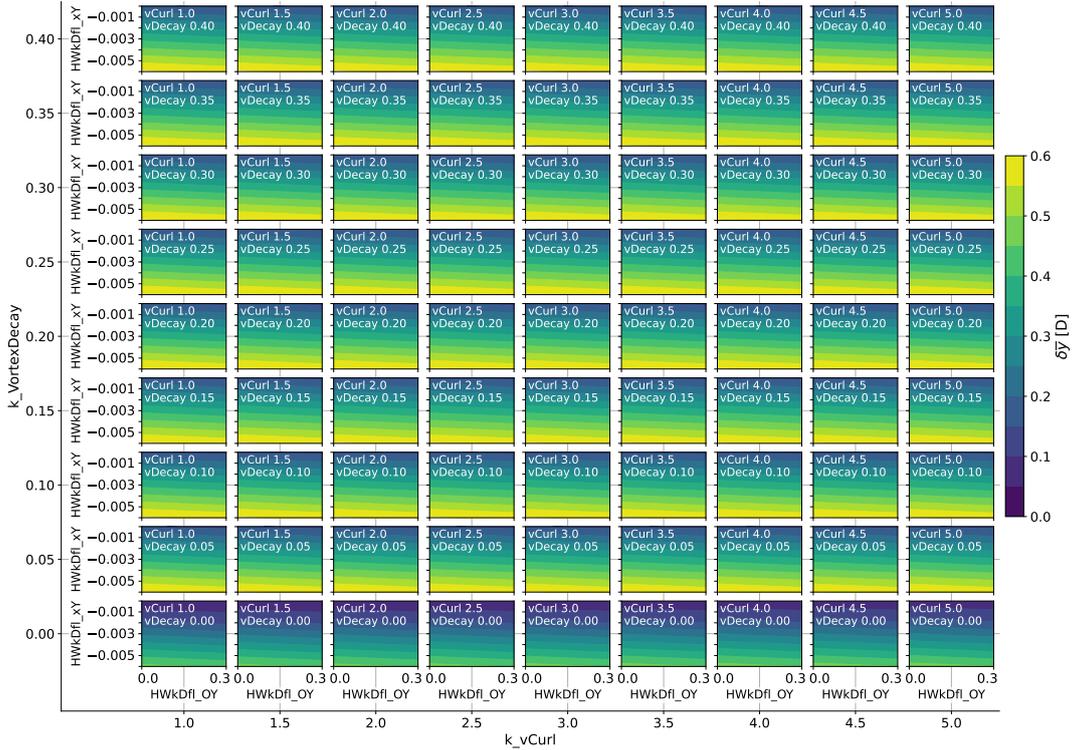


Figure 6. Differences between the wake center obtained through FAST.Farm and those obtained using LES for the zero-yaw case. Shown are the values related to Sweep 3 from Table 2. Each small subplot is a function of the two wake deflection parameters while the outermost axes vary in the other two curled wake parameters. Darker shades mean such constants performed best.

Different optimal values exist between the different yaw configurations. This is a result of the LES for positive and negative yaw having nonsymmetric results due to the veer, while FAST.Farm asymmetries of results are due to the rotation of the rotor only, not because of veer.

Nonetheless, with the goal of picking single values for all yaw scenarios (also to not overfit the LES data), the final values chosen strike a balance between all three yaw cases (positive, negative, and zero yaw) and both quantities of interest (power and wake center). For instance, wake location is much less sensitive to the value of k_{vCurl} than power, so we prioritize values in which power matched better. The optimal k_{vCurl} for power varies depending on the yaw configuration. That is because the LES solutions have asymmetrical results due to the height-varying veer that is not considered in the FAST.Farm simulations. In this work we prioritize yawed turbines, so we focus on the results from Fig. 8(e) and (g). We pick the final value as 5 because the negative yaw scenario is less sensitive to its value than the positive yaw case. Final values are given in Table 3. In the final table, for the purposes of generality, we also provide the optimal values for the polar wake model (steps not shown in this article). We note that this work assumes that the errors between the LES and the FAST.Farm model are constant across all cases, such that they are a bias and not a measure of uncertainty, allowing us to perform analysis focusing on the relative differences.

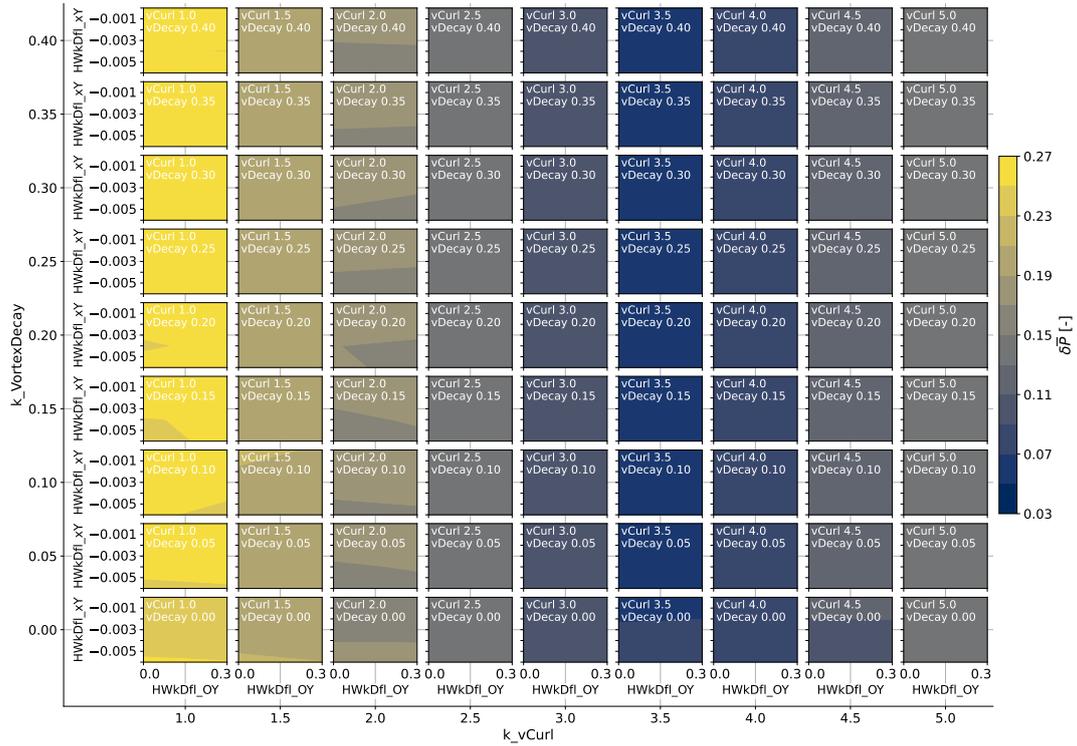


Figure 7. Differences between total power obtained through FAST.Farm and those obtained using LES for the zero-yaw case. Note how for this sweep, an optimal k_vCurl around 3.5 exists.

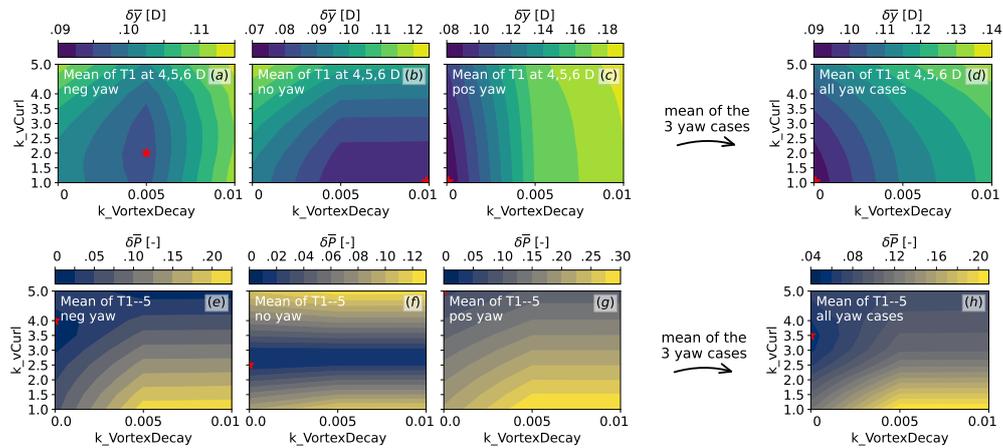


Figure 8. Differences in wake position (a–d) and power (e–g) between FAST.Farm and LES for each yaw condition, as indicated in each subplot. The red mark represents the optimal point of each scenario and each quantity of interest. The mean of each quantity is shown in (d) and (h).

Table 3. Final calibrated FAST.Farm parameters for the IEA Wind 15 MW fixed-bottom reference wind turbine. Default values are also shown for reference.

	Polar		Curled	
	Default	Calibrated	Default	Calibrated
$C_{\text{HWkDfl_OY}}$ [m/deg]	0.3	0.3	0	0
$C_{\text{HWkDfl_xY}}$ [deg ⁻¹]	-0.004	-0.006	0	0
$k_{\text{VortexDecay}}$ [-]	—	—	0	0
k_{vCurl} [-]	—	—	2	5

280 We note that the calibration procedure outlined above applies to the fixed-bottom IEA Wind 15 MW reference wind turbine for the conditions of interest in this work. The reader is reminded that these values are subject to change depending on the conditions of interest, and a tuning exercise should always be carried out for the farm layout and atmospheric conditions of interest. For illustration purposes, we finish this section by showing a flow field picture comparing the LES and the equivalent FAST.Farm case with the calibrated parameters for the curled wake model (Fig. 9).

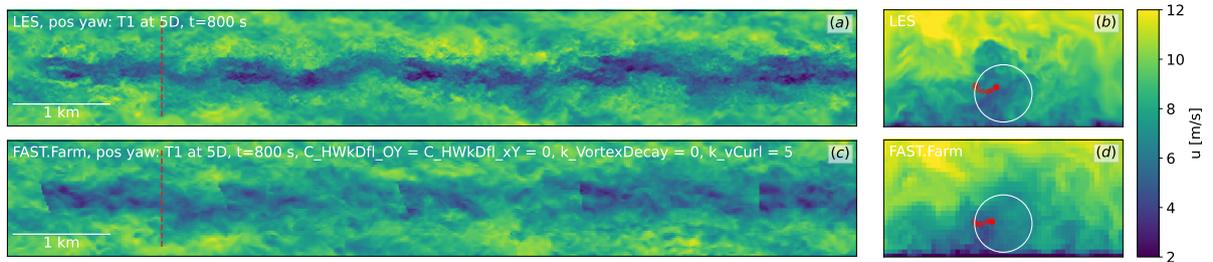


Figure 9. A snapshot of the positive-yaw case of LES and FAST.Farm for the calibrated curled wake solution. Note how the background flow is the same. The wake plane at $5D$ of the first turbine, indicated by the red dotted line, is shown on the right. The wake center at the time instant shown, as well as a 50 s history, is indicated by red dots on the wake plane panels.

285 4 Results

We present the results of all the cases in terms of contour plots, or, “heat maps”, which includes sweeps in terms of both wind direction and yaw angle. In the figures that follow, the vertical axis varies with respect to the reference yaw γ , and the horizontal axis represents the wind direction.

The reference conditions no yaw ($\gamma = 0$) and aligned wind ($\alpha = 0$) are highlighted by black lines, splitting the contour plots into four quadrants. Each of the four quadrants represent a combination of positive or negative yaw and wind direction. The intersection of the reference conditions is what we define as the baseline (reference) case; that is, no yaw and aligned wind ($\alpha = \gamma = 0$). We use this baseline case to normalize results and to give percent differences with respect to it. We acknowledge that,

ideally, each wind direction should have its own baseline case, and that such a single baseline case may not be representative of a true baseline for $\alpha \neq 0$; however, we have decided to use a single baseline value for ease of comparison.

295 4.1 Effects on wind farm power

The power differences are shown for each turbine and each condition in Fig. 10(a–e). For each turbine, we normalize the power by the power of the leading turbine at the baseline no-yaw, aligned-wind scenario. Such normalization means that non-leading, downstream turbines will generally have a normalized power value lower than unity due to the upstream wake effects. Figure 10(f) shows an average of the five turbines.

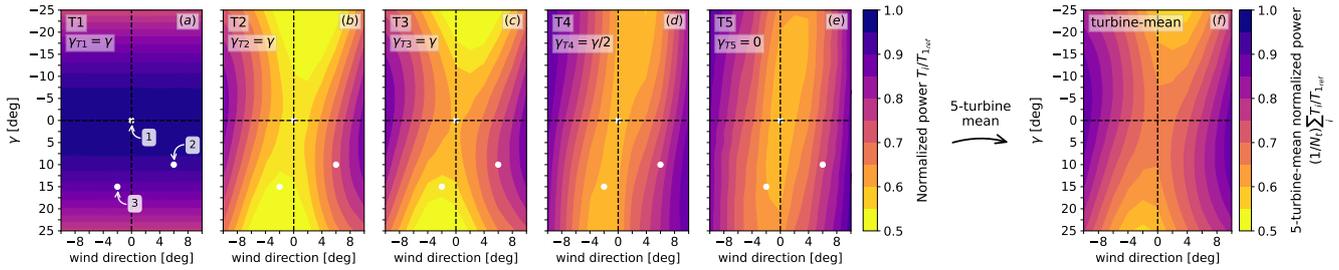


Figure 10. Normalized power for each turbine under each configuration investigated. For each turbine, (a–e), the power of turbine i is normalized by the baseline power of the leading turbine, $P(T_i/T_{1,ref})$. Turbine-mean results are shown in (f). The same data at points indicated by 1, 2, and 3 are shown in Fig. 11.

300 Figure 10 contains a lot of information, and it can be difficult to understand all of it at first. To provide the reader a more intuitive way of understanding this figure and subsequent figures, we show the same data for the three indicated points differently. Taking the values of the contours from Fig. 10 of the three points, we obtain Fig. 11, which shows the power down the row. The power normalization choice becomes more intuitive now. Point 2 generates more total farm-wide power than the baseline case (Point 1) and Point 3.

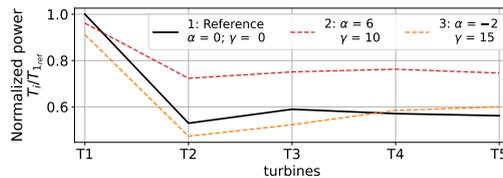


Figure 11. Normalized power down the row for selected cases, showing typical behavior of power loss on waked turbines. As indicated in the legend, each curve is related to each of the points 1, 2, and 3 from Fig. 10.

305 The average power of each case is shown in Fig. 10(f), but it is difficult to visualize the relative differences with respect to the baseline case. Therefore, we take the turbine-mean normalized power from Fig. 10(f) and normalize by its baseline case, shown in Fig. 12. Using such normalization allows us to see which farm configurations provide a positive power difference with respect to the baseline case, where no wake steering strategy is employed.

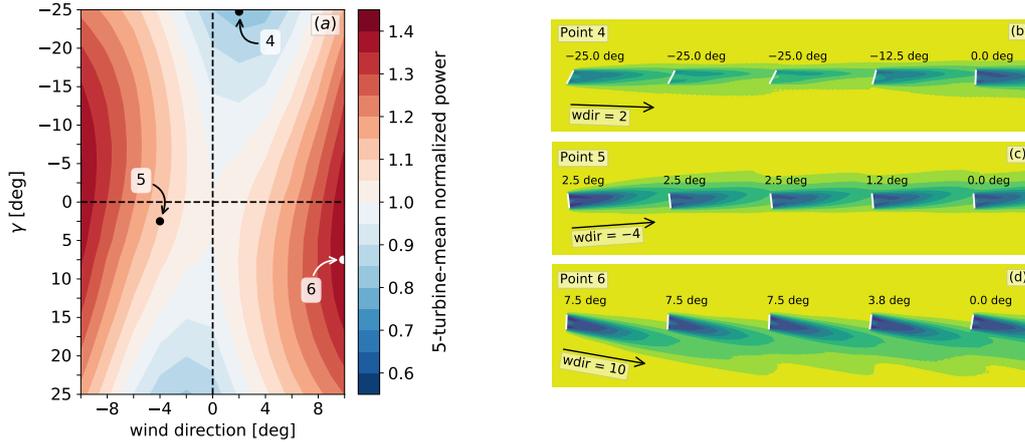


Figure 12. Farm-mean power (Fig. 10(f)) normalized by the power of the baseline no-yaw aligned wind case (Point 1 in Fig. 10(a)), (a). Power variation exceeds 35% of the baseline power. Time-averaged flow field of points 4, 5, and 6 indicated are shown on right, (b,c,d).

Figure 12 shows that power differences can exceed 35% in certain scenarios (shown by Point 6) as well as have a negative impact (e.g., Point 4). To relate the positive power differences to the actual configuration, we visualize three arbitrary points of the domain. For Point 4, Fig. 12(b), the turbines are highly yawed in the negative direction while the wind is at a slightly positive ($\alpha = 2$) inflow direction. In this case, as the turbines yaw to steer the wake away from downstream turbines, the wind direction brings the wake directly back into the rotor of downstream turbines, making wake steering highly ineffective. The loss in power is noted by the blue color. Such scenarios are not the goal but can be observed in scenarios of unintended yaw misalignment. On the other hand, the condition noted by Point 5, Fig. 12(c), is a realistic combination of wind direction and yaw angle, and it yields a power difference of about 13% with respect to the baseline case. Finally, the regions of dark red are the conditions of highest power gain, which are related to highly skewed flow, naturally diverting most of the wake by itself. This is clearly illustrated in Fig. 12(d), on a condition of $\alpha = 10$ degrees with $\gamma = 7.5$ degrees. While the positive power differences obtained by dark red areas are operationally attractive, as will be shown in the next section, this comes with significantly increased fatigue loads, rendering such a condition unrealistic. This is an important observation, as it showcases the need for performance studies with wake steering alongside structural fatigue load analysis.

We note that a wind farm operator that has control over the turbine configurations (but not the wind) can use charts like the ones presented in Figs. 10 and 12 to estimate power changes given certain yaw angle configurations.

While Fig. 12(a) appears symmetric at first, we can compute the differences between the positive and negative yaws in terms of power. Because there is zero veer in these flow fields, and thus no effect on the wake, the asymmetries presented next are due to the rotation of the rotor only. Here we introduce the idea of splitting the results in terms of the four quadrants, as separated by the dashed lines at $\alpha = 0$ and $\gamma = 0$. To establish a direct comparison, we split Fig. 12(a) in half horizontally, at $\gamma = 0$. We obtain the differences between positive and negative yaw by taking the bottom two quadrants (positive yaw) and comparing against a mirrored version of the top two quadrants (negative yaw). Figure 13 shows the approach. Panel (a) shows

330 the bottom two quadrants and panel (b) shows the top two quadrants, mirrored. Panel (c) shows the differences in terms of absolute percentage when compared to the baseline case.

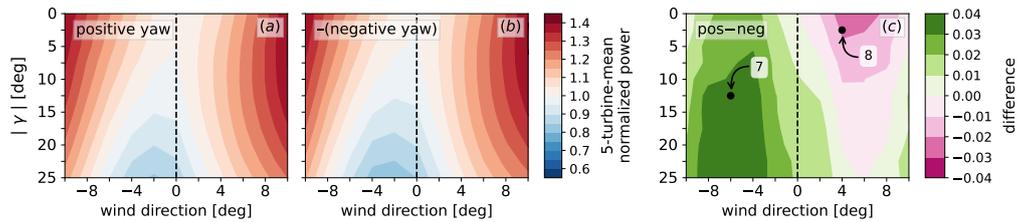


Figure 13. Asymmetry differences between positive and negative yaw cases. (a) Positive yaw contour map, the same as the bottom part of Fig. 12; (b) negative yaw contour map, mirrored from the top part of Fig. 12; (c) difference in terms of power between positive and negative yaw. In (c), for winds and yaw of opposite sign, positive yaw is up to $\sim 4\%$ better (green on the left side of (c)), while for winds and yaw of the same sign, negative winds and yaw yields up to $\sim 2\%$ more power (pink part on the right side of (c)). Points 7 and 8 in (c) are shown in Figs. 14 and 15, respectively.

Regarding Fig. 13(c), the green left side means negative wind and positive yaw yields more power than its symmetric counterpart (positive wind and negative yaw), for all conditions. In other words, the bottom-left quadrant outperforms the top-right quadrant. The mostly-pink right side means the bottom-left quadrant underperforms the top left quadrant—that is, for wind and yaw of the same sign, positive values are up to 2.2% worse. For aligned wind ($\alpha = 0$ line), a positive yaw always results in more power than a negative yaw of the same magnitude. More intuitive visual schematics for Points 7 and 8 of Fig. 13(c) are presented in Figs. 14 and 15, respectively.

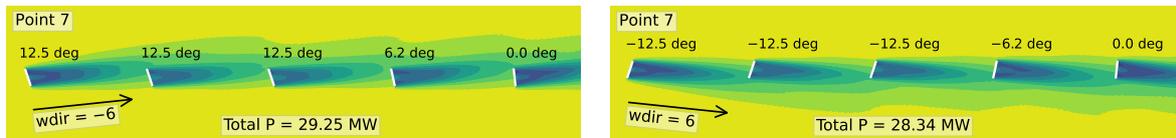


Figure 14. Visual schematic of the dark green point noted as 7 in Fig. 13(c). Negative wind direction with positive yaw angles generates 3.2% more power than its symmetric scenario (positive wind direction and negative yaw angle).

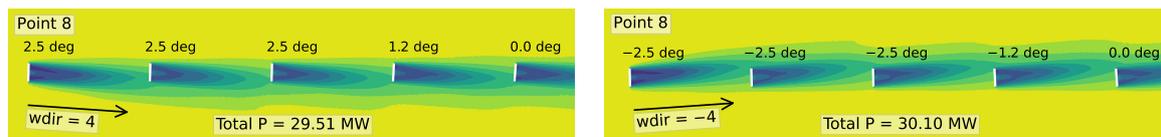


Figure 15. Visual schematic of the dark pink point noted as 8 in Fig. 13(c). Positive wind direction with positive yaw angles generates 2.2% less power than its symmetric scenario (negative wind direction and negative yaw angles).

4.2 Effects on wind farm loads

Load characteristics for each turbine and farm-mean are presented in this section. We will focus on the load channels outlined
340 in Table 1, processed as discussed in Section 2.4. In all the plots shown in this section, the color scale has two values. One
is the absolute value of the quantity being shown, and the other is a percent change from the baseline (indicated by “BL”),
which is the leading turbine of the no-yaw, aligned-wind case. The percent scale allows us to identify relevant changes while
the absolute scale captures the severity of the change. We will focus on the areas of positive net differences in power to assess
their corresponding impact on loads and relate our findings to the four quadrants of the contour plots.

345 The times-series mean loads are shown in Fig. 16, where each row shows one channel for each turbine as well as a turbine-
mean load map. First, and as expected, the leading turbine experiences vastly different wind characteristics due to its unyawed
condition, and for any given yaw misalignment γ , the loads remain constant across wind directions. The tower-base torsional
moment exhibits asymmetries with respect to the yaw angle and has the lowest mean at around $\gamma = 12$ degrees. The blade-
root and tower-base mean bending moments are highest for the baseline case. The mean LSS bending moments in a rotating
350 frame of reference are nearly zero for all cases (as expected). Downstream, the mean blade-root and tower-base loads for
waked turbines follow more or less the trends of the power map. The magnitude of the mean for the tower torsional moment
shows signs of asymmetry for the waked turbines too, in which the conditions within the bottom-right quadrant (positive wind
direction and positive yaw) result in the lowest mean loads.

Next, we show the standard deviation of the load channels in Fig. 17 in the same fashion. Standard deviation is often more
355 intuitive than DEL, can be considered as a proxy for fatigue loads, and offers more insight into the asymmetries than the mean
loads. For the blade-root and LSS bending moments, for any given wind direction and any turbine, the load standard deviations
monotonically increase from positive to negative yaw angle. Such monotonic behavior has been observed in both simulation
and measured data (Damiani et al., 2018) and is attributed to gravity loading being relieved by aerodynamic loading as the yaw
goes from negative to positive. For both channels, the bottom-right quadrant (positive wind direction and positive yaw) incurs
360 the lowest standard deviation. Reductions of 4 to 6% with respect to the baseline can be observed for the blade-root bending
moment, and up to 20% for the LSS. Analogously, $\sim 4\%$ higher load standard deviations for the blade-root bending moment
and up to 30% for the LSS are seen in the worst-performing quadrant (negative wind direction and negative yaw). For any
given wind direction for these two channels, the more positive the yaw, the better. For example, for aligned-wind situations, a
certain positive yaw, γ , will generally reduce the standard deviations with respect to no-yaw, while a negative yaw, $-\gamma$, will
365 increase it.

For any given wind direction, non-monotonic (with respect to the yaw angles), more complex trends in standard deviation are
seen for the tower-base bending and torsional moments. For the tower-base bending moment, the conditions of low standard
deviation (yellow regions in turbine-mean map, panel (3) of Fig. 17) are the same conditions of overall power reductions
with respect to the baseline and are thus undesirable. As mentioned, however, these conditions, while undesirable, might be
370 met during an unintended yaw misalignment scenario and are worth showing. Excluding those conditions from our analysis,
the same bottom-right quadrant (positive wind direction and positive yaw) is the best performing, and the top-left quadrant

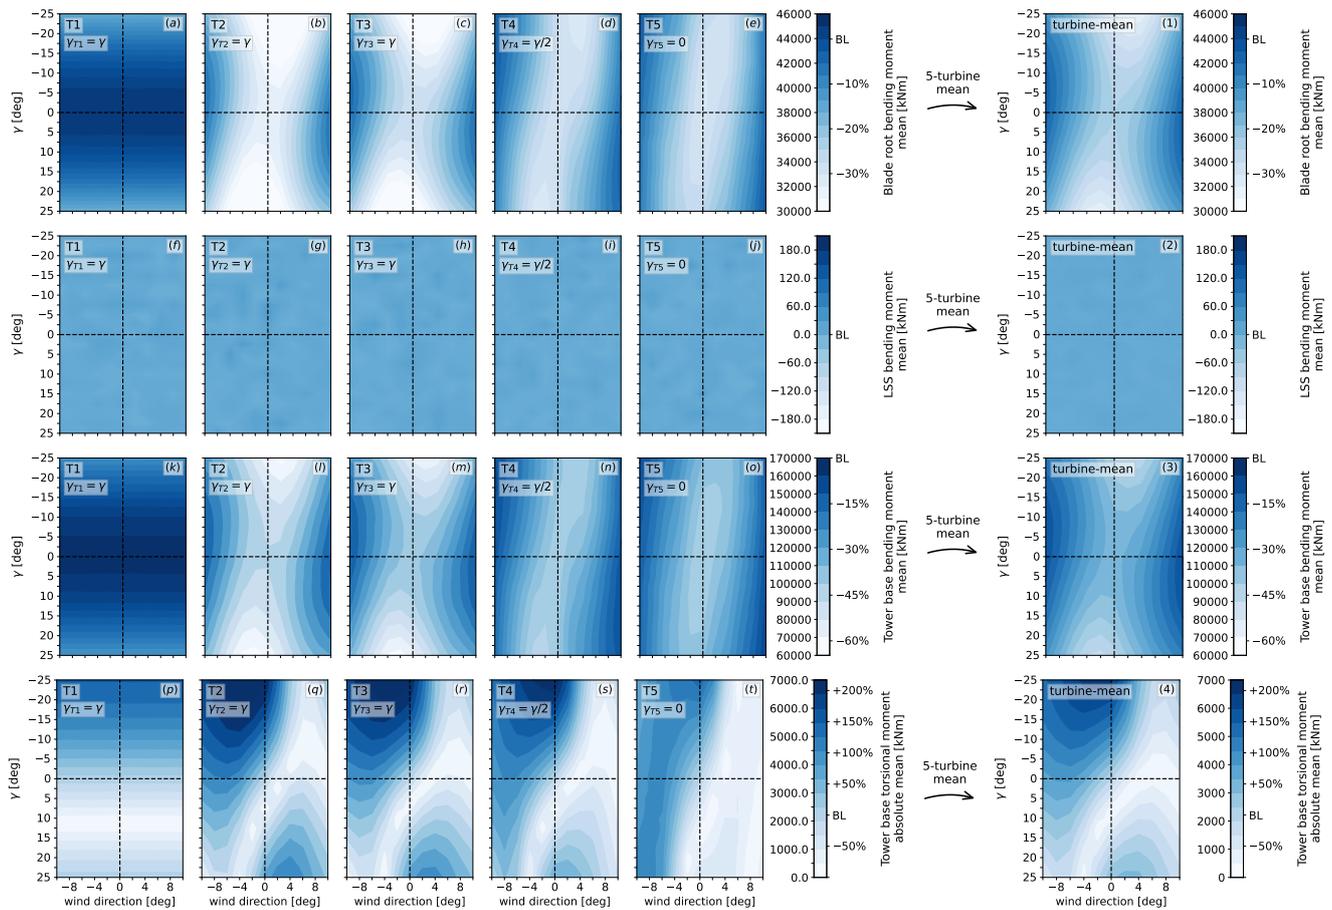


Figure 16. Mean of the load channels of interest for each turbine. One channel per row; different scales are used for each channel. Note how some channels increase monotonically with respect to the yaw. Low-speed shaft bending moment mean is nearly zero, as expected.

(negative wind direction and negative yaw) is the worst performing. Regarding the tower-base torsional moment, areas of low standard deviation are observed for negative wind directions, where positive yaw condition is better than negative. Between the two symmetric top-right and bottom-left quadrants for all channels shown, the bottom-left quadrant (negative wind direction and positive yaw) generally yields lower values of standard deviation, suggesting the combination of negative wind direction and positive yaw incur lower fatigue damage. In general, the LSS bending moment and the tower-base bending moment are the two channels with the largest sensitivity, resulting in larger increases/decreases when compared to baseline cases.

The two channels noted above with significant standard deviation are the same channels with the highest relative variation in short-term DEL, as shown in Fig. 18. In general, the trends observed in the standard deviations are very similar to those seen here for DELs, with a clear correlation between the two quantities. An exception is the tower-base torsional moment, where the downstream turbines have significantly lower values of DELs when compared to the leading turbine. The turbine-mean contours are similar to the standard deviations, and the cases with the lowest DELs are those from the bottom-right quadrant

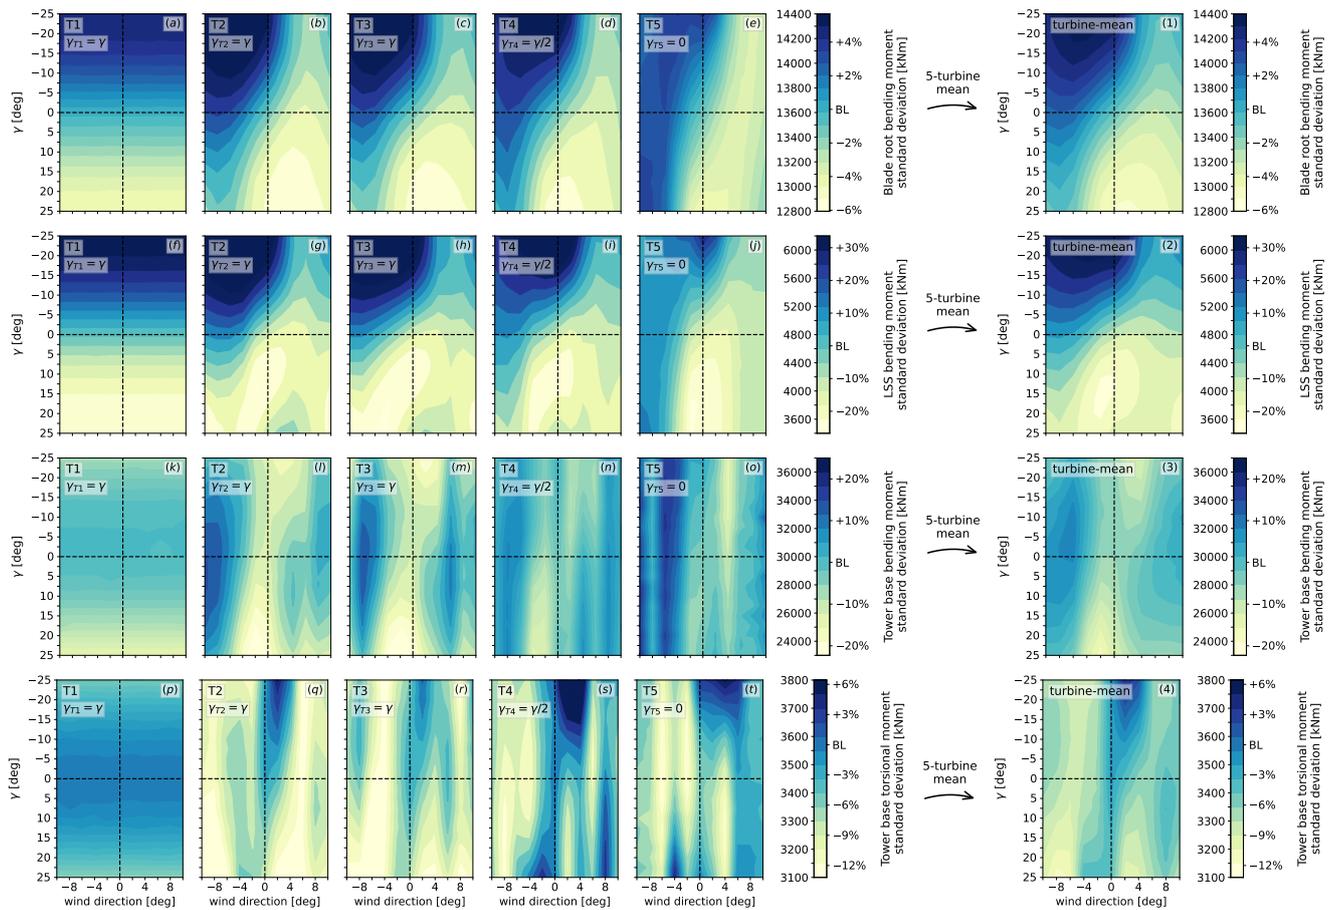


Figure 17. Standard deviation of the load channels of interest for each turbine. One channel per row; different scales are used for each channel. The standard deviation of low-speed shaft (f–j) and tower-base (k–o) bending moments are the most sensitive to the yaw and wind direction conditions investigated.

for the blade-root, LSS, and tower-base bending moments, when excluding the areas of net loss in power (the blue shades of Fig. 12(a)).

385 4.3 Asymmetries and partial waking

Partially waked conditions are known to play a role in load fluctuations experienced by the blades. If the wake is prominently hitting the rising blade, then an increase in load fluctuations will be observed. Similarly, if the wake is prominently reaching the falling blade, lower fluctuations in loads can be seen. This was explained nicely by Stanley et al. (2020) and is confirmed here with the additional context of positive versus negative yaw angles.

390 To help explain the differences observed in terms of symmetry, we pick two pairs of cases. The first is with wind direction ± 2 degrees and $|\gamma| = 7.5$ degrees, in which a variation of plus or minus 2% in the standard deviation of the blade-root bending

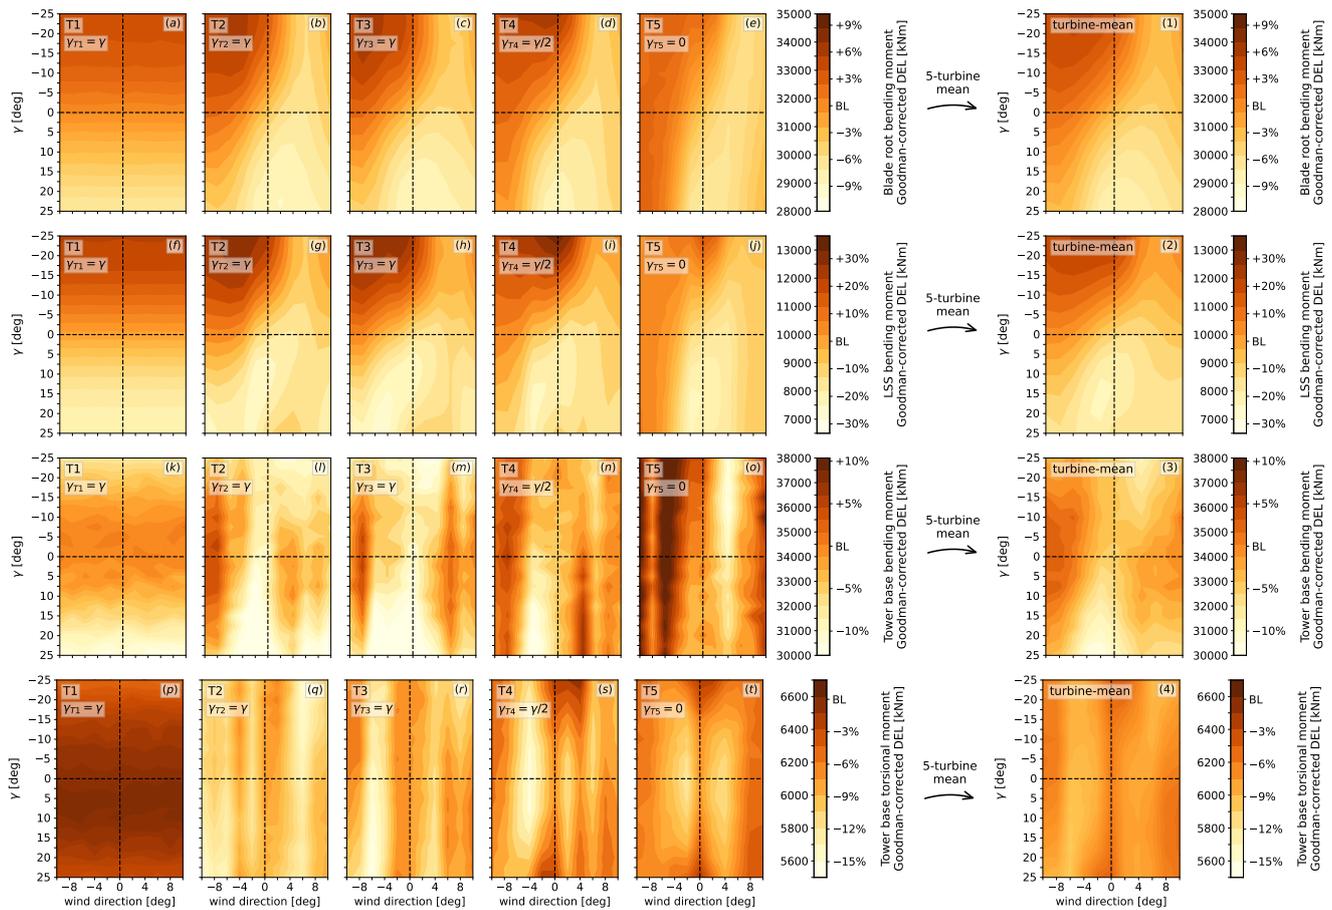


Figure 18. Short-term DELs including Goodman correction for the load channels of interest for each turbine. Note different scale for each channel.

moment loads is observed (Fig. 17(a–e)). A second case is more extreme, and the differences are more evident: wind direction ± 4 degrees and $|\gamma| = 20$ degrees, which gives a load variation of more than 6% with respect to the baseline case. We will use these two pairs to highlight partial waking and asymmetry in loads. The two scenarios are illustrated in Figs. 19 and 20.

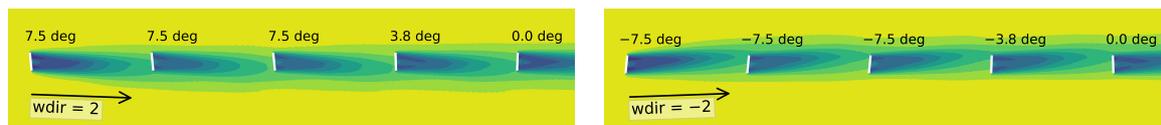


Figure 19. Flow field of symmetric cases where partial waking plays a role.

395 During partial wake conditions, the blades passing through the waked part of the rotor are subject to smaller aerodynamic loading due to reduced freestream wind speed. In both Figs. 19 and 20, the positive yaw and positive wind direction cases have the wake reaching the downstream turbines more on the falling side. At the falling side of the blade, gravitational loads have the

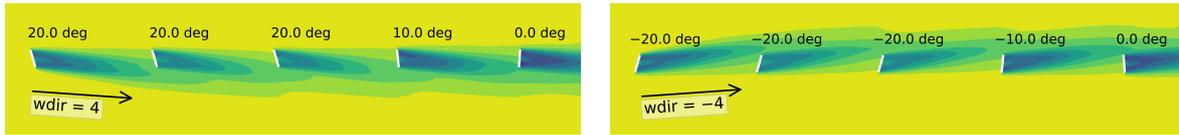


Figure 20. Flow field of symmetric cases where partial waking plays a more significant role. The yaw angles and wind direction combinations might not be realistic.

same sign as the aerodynamic loading. Because gravitational loads have constant amplitude, the smaller aerodynamic load on the falling side results in overall smaller loads, that is, cycles that have smaller peaks. In this situation of positive yaw, the rising blades experience unwaked conditions, and thus higher loads, which are now opposite to the sign of the gravitational loads, resulting in smaller peaks of load magnitude for the rising blades. The opposite, analogous physics happen under negative yaw and negative wind direction conditions, resulting in overall larger peaks. For negative yaw cases, the wake is hitting the blades that are rising, reducing aerodynamic loads when aerodynamics and gravity are opposite, resulting in larger magnitudes. Similarly, the unwaked, falling blades experience higher-speed wind that produces aerodynamic loads with the same sign as gravitational loads, again resulting in cycles with larger peaks. This behavior is illustrated in Fig. 21 for a small time series shown in terms of revolutions of the rotor (a revolution starts with the blade pointing up) and an aggregate for the entire 20 min simulated in this work.

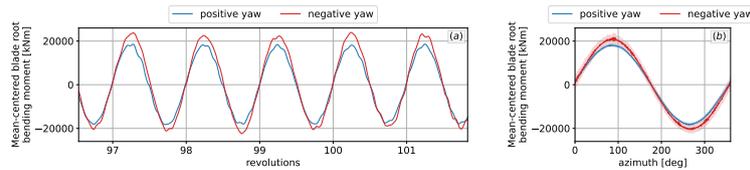


Figure 21. Mean-centered blade-root bending moments of the same symmetric case shown in Figs. 19 and 20. (a) Illustration of a short period of the time-series; (b) revolution-averaged quantities, with shaded region showing one standard deviation from the mean.

5 Discussion

We discuss the results in terms of the four quadrants of the contour plots shown. Figure 22 provides schematics of each of the four quadrants for reference.

Given the quadrants in Fig. 22, a summary of the findings is given below for each quadrant:

- Figure 22(a):
- Represents the top-left quadrant in the contour plots. Worst-performing quadrant in terms of loads; worst-performing quadrant overall.
 - Consistent gain in power when compared to baseline case, up to 35%. Values are analogous to symmetric Fig. 22(d) and the exceptions described therein.
 - Small gain in power (1–2%) compared to equivalent symmetric cases (bottom-right quadrant).

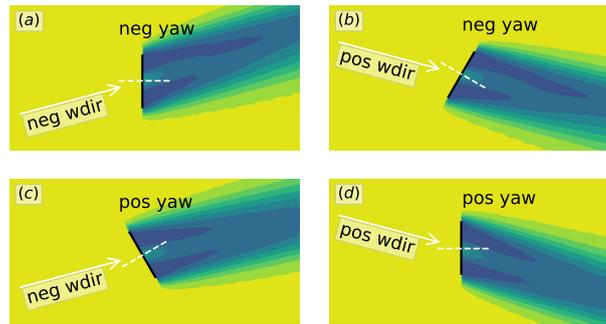


Figure 22. Schematic of the cases investigated. The figures are arranged in such a way that they correspond to the each of the four quadrants from prior images– (a) represents the top left quadrant; (b) the top right; (c) the bottom left; and (d) the bottom right

- Significantly worse loads compared to its symmetric case. Blade loads standard deviation up to 9% higher and DEL up to 18% higher; LSS standard deviation up to 50% higher and DEL up to 50% higher; tower-base bending moment standard deviations up to 10% higher, and DEL up to 8% higher. Tower-base torsional loads are about the same or sometimes slightly (less than 5%) better than their counterpart.

420

Figure 22(b): – Represents the top-right quadrant in the contour plots.

- Small values of wind direction ($\alpha \leq 4$ degrees) and very negative values of yaw ($\gamma \leq -15$ degrees) are undesirable, as the combination of wind and yaw result in the wake being steered directly into downstream turbines. These cases are the ones with a decrease in power with respect to the baseline case; thus, loads and DEL information are of interest only during periods of unintended yaw misalignment.

425

- Positive power differences of up to 40% compared to the baseline case in the high wind direction $\alpha > 8$, small yaw angle $|\gamma| < 5$. Power losses are observed in undesirable configurations mentioned in the prior bullet point.

- Power losses compared to its symmetric case can be observed for any given wind direction. The losses are the highest for wind where $3 < \alpha < 9$ and $\gamma < -10$, up to 4%. They are the lowest for small wind directions $\alpha \leq 2$ and small yaw misalignment angles, up to $\gamma = -5$, where the losses are less than 1%.

430

- For any given positive wind direction, blades and low-speed shaft loads (in terms of both standard deviation and DEL) increase monotonically from lowest values at no-yaw conditions (for this quadrant). Tower-base bending and torsional moment standard deviation and DEL increase slightly depending on the magnitude of the yaw angle and wind direction.

435

- Regarding its symmetric case, blade and tower bending moment are about 3% lower, LSS loads are 5–10% higher, depending on the combination of yaw angle and wind direction. Tower torsional moments are slightly lower (less than 3%) than its symmetric cases.

- Figure 22(c): – Represents bottom-left quadrant in the contour plots.
- 440 – Positive power differences of up to 38% compared to baseline case. Decrease in power can be observed under undesirable scenarios as noted in Fig. 22(b).
- Power differences with respect to symmetric scenarios are between 1 and to 4%. Findings are analogous to findings from Fig. 22(b) listed above.
- Figure 22(d): – Represents the bottom-right quadrant in the contour plots. Best-performing quadrant in terms of loads; best-performing quadrant overall.
- 445 – When compared to the baseline case, for yaw angles $\gamma < 15$ degrees, a power gain of up to 35% is observed for wind direction $\alpha = 10$ degrees, and up to 20% gain for wind direction $\alpha = 4$ degrees. For very large yaw angles $\gamma > 18$ degrees and small wind directions $\alpha < 2$ degrees, the undesirable situation of power loss is observed.
- 450 – Compared to equivalent symmetric scenario: 1–2% loss in power.
- For positive yaw angles, up to 5% lower standard deviation and 9% lower DEL is seen for the blade-root bending moment and up to 20% lower standard deviation and 25% lower DEL for the LSS (both with respect to their baseline case). Standard deviation for the tower-base bending moment is within 5% of its baseline, and DELs 2%, depending on the wind direction. The tower-base torsional moment standard deviation is up to 9% lower, and DEL is at least 5% lower with respect to the baseline.
- 455 – Compared to its symmetric counterpart, large benefits in terms of standard deviation and DELs for the blade-root bending moment and the LSS bending moment are present. Smaller differences are observed for the tower loads.

The discussions above are given in terms of yaw angles and wind direction. We can interpret the results from a different point of view. Operators can only control the yaw, so it is useful to summarize the findings in terms of wind directions. Considering only the regions of net positive power gain (red areas in Fig. 12(a)):

460

- For any given positive wind direction: yaw of $-2.5 < \gamma < 15$ degrees yields increases in power of 10–20% over baseline. Positive yaw will incur lower blade-root, LSS, and tower-base torsional standard deviation and DEL. Tower-base bending moment means and standard deviation follow the power map (higher power leads to higher loads). It is overall better to keep the turbine positively yawed.
- 465
- For any given negative wind direction: Positive yaw results in lower magnitudes for blade-root, LSS, tower-base bending, and tower torsional standard deviation and DEL values. With regard to tower-base torsion, negative yaw angles are preferred for inflow angles between 0 and -4 degrees, due to the loss in power if positive yaw is enforced. A loss in power of up to 2% is observed for positive yaws, but the gains in terms of loads may be significant enough to justify a positive yaw configuration for negative winds too.
- 470

5.1 Limitations

The model and approach used here have some limitations, and the findings reported in this work should be considered alongside them.

475 The wind farm used in this study consists of a single row of turbines, in an attempt to model a single row of a realistic farm in an independent way. With the current layout, however, more complex interplay between rows of turbines is not captured. For some conditions (wind directions), applying a yaw offset would not steer the wake in a meaningful way, as the wake might naturally move away from downstream turbines due to the wind direction. Figure 12(d) illustrates this limitation well: the yaw offset incurs an increased fatigue loading with no significant power advantage. With a more complex farm layout, that would not necessarily be true, as other rows of turbines would be located where the wake currently is. While the sweep of the angles
480 used in this work are relatively high, results for the cases where both the wind and yaw angles are closer to the extremes of the range selected should be interpreted with caution if nearby rows of turbines are to be included or considered.

This work seeks to identify trends and relative differences in power and fatigue loading given different configurations and conditions. It is not the goal of this work to give absolute values in terms of gains or losses. The tool of choice, FAST.Farm, is an engineering-fidelity tool that comes with modeling errors. In all of our analysis and comparisons, we attempt to quantify
485 relative differences.

Currently, FAST.Farm does not consider the height-varying cross-flow component of the wind in the wake calculation. Because of that, our synthetic turbulence-generated flow fields were created with a zero mean cross-flow component. The results presented here do not consider veer. As seen by the LES cases used for calibration, it is expected that the inclusion of a wind veer typical of the northern hemisphere is likely to further accentuate the findings discussed in this work. A deeper study
490 on the interplay of atmospheric boundary layer physics (especially the wind veer), wake flow physics, and dynamics relative to positive and negative yaw misalignment could help to more deeply explain some of the nuanced findings in this work. A higher-fidelity study would also help to confirm that the trends observed in FAST.Farm, and the conclusions we derive, remain consistent as the physics are better resolved.

6 Conclusions

495 In this work we investigated yaw misalignment in a systematic way regarding positive and negative yaws with respect to wind directions varying from -10 to 10 degrees. We sought to analyze data trends and physical driving mechanisms for power and loads asymmetries that have been previously reported.

In general, we found that for positive wind directions, a positive yaw yields more power. For these conditions, lower positive yaw angles result in larger power gains. For the blade-root bending moment and LSS bending moment, the standard deviation
500 and DELs are consistently lower for positive yaw angles when compared to the negative counterparts of same magnitude. For the tower-base loads, the trends are more complex, but in general positive yaw still results in lower overall loads. For negative wind directions, positive yaw results in slightly less power than negative yaw, but the gains in terms of loads are significant.

Blade-root and LSS bending moments have significantly lower loads in terms of standard deviation and DELs. Tower-base loads are comparable between positive and negative yaws.

505 The interplay of gravity and aerodynamic loading is the physical mechanism driving some of the asymmetries observed. This interplay can be triggered in at least two ways observed in our simulations: first, by yaw angles that present either the rising or falling blade to the incoming wind and second, by partial waking that reduces the inflow wind velocity on either the rising or falling half of the rotor plane. The latter mechanism was evidenced by flow field visualization and time-series analysis of blade-root bending moment channels.

510 For the wind condition and the farm layout investigated in this work, and any wind direction, yawing the turbine in the positive direction is better than in the negative direction. While small power losses may incur ($\sim 2\%$ with respect to its symmetric counterpart, see Fig. 13(c)), gains in terms of loads may be significant for farm operators. Finally, we note that the workflow used in this work could be easily adapted to investigate other conditions and farms of interest.

Code and data availability. Simulations were performed using AMR-Wind (as of mid-2023, before version numbering has been introduced, 515 available at github.com/exawind/amr-wind) and FAST.Farm version 3.3.0 (available at github.com/OpenFAST/openfast). Data shown in this work are available from the authors upon request.

Author contributions. RT: Conceptualization, methodology, investigation, formal analysis, software, validation, writing (original draft). GB: Conceptualization, supervision, writing (original draft, review and editing). JJ: Conceptualization, supervision, writing (review and editing) RM: Validation, writing (review and editing) CB: Validation, writing (review and editing) KS: Conceptualization, writing (review and editing) 520 JK: Conceptualization, supervision, writing (review and editing).

Competing interests. The authors declare no competing interests.

Acknowledgements. This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the 525 article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes. This work was funded by Shell Global Solutions International B.V. under agreement ACT-1800038-5. The research was performed using computational resources sponsored by the Department of Energy's Office of Energy Efficiency and Renewable Energy and located at the National Renewable Energy Laboratory.

References

- 530 N. J. Abbas, D. S. Zalkind, L. Pao, and A. Wright. A reference open-source controller for fixed and floating offshore wind turbines. *Wind Energy Science*, 7(1):53–73, 2022.
- M. Adaramola and P.-Å. Krogstad. Experimental investigation of wake effects on wind turbine performance. *Renewable Energy*, 36(8): 2078–2086, 2011. ISSN 0960-1481. <https://doi.org/10.1016/j.renene.2011.01.024>.
- M. Bastankhah and F. Porté-Agel. Wind farm power optimization via yaw angle control: A wind tunnel study. *Journal of Renewable and*
535 *Sustainable Energy*, 11(2), 2019.
- M. Bastankhah and F. Porté-Agel. Experimental and theoretical study of wind turbine wakes in yawed conditions. *Journal of Fluid Mechanics*, 806:506—541, 2016. <https://doi.org/10.1017/jfm.2016.595>.
- E. Branlard, L. A. Martínez-Tossas, and J. Jonkman. A time-varying formulation of the curled wake model within the FAST.Farm framework. *Wind Energy*, 26(1):44–63, 2023.
- 540 U. Ciri, M. A. Rotea, and S. Leonardi. Effect of the turbine scale on yaw control. *Wind Energy*, 21(12):1395–1405, 2018. <https://doi.org/10.1002/we.2262>.
- R. Damiani, S. Dana, J. Annoni, P. Fleming, J. Roadman, J. van Dam, and K. Dykes. Assessment of wind turbine component loads under yaw-offset conditions. *Wind Energy Science*, 3(1):173–189, 2018.
- B. M. Doekemeijer, S. Kern, S. Maturu, S. Kanev, B. Salbert, J. Schreiber, F. Campagnolo, C. L. Bottasso, S. Schuler, F. Wilts, T. Neumann,
545 G. Potenza, F. Calabretta, F. Fioretti, and J.-W. van Wingerden. Field experiment for open-loop yaw-based wake steering at a commercial onshore wind farm in Italy. *Wind Energy Science Discussions*, 2020:1–22, 2020. <https://doi.org/10.5194/wes-2020-80>.
- B. L. Ennis, J. R. White, and J. A. Paquette. Wind turbine blade load characterization under yaw offset at the SWiFT facility. In *Journal of Physics: Conference Series*, volume 1037, page 052001. IOP Publishing, 2018.
- P. Fleming, J. Annoni, J. J. Shah, L. Wang, S. Ananthan, Z. Zhang, K. Hutchings, P. Wang, W. Chen, and L. Chen. Field test of wake steering
550 at an offshore wind farm. *Wind Energy Science*, 2(1):229–239, 2017. <https://doi.org/10.5194/wes-2-229-2017>.
- P. Fleming, J. King, K. Dykes, E. Simley, J. Roadman, A. Scholbrock, P. Murphy, J. K. Lundquist, P. Moriarty, K. Fleming, et al. Initial results from a field campaign of wake steering applied at a commercial wind farm—part 1. *Wind Energy Science*, 4(2):273–285, 2019.
- P. Fleming, J. King, E. Simley, J. Roadman, A. Scholbrock, P. Murphy, J. K. Lundquist, P. Moriarty, K. Fleming, J. van Dam, C. Bay,
555 R. Mudafort, D. Jager, J. Skopek, M. Scott, B. Ryan, C. Guernsey, and D. Brake. Continued results from a field campaign of wake steering applied at a commercial wind farm – Part 2. *Wind Energy Science*, 5(3):945–958, 2020. <https://doi.org/10.5194/wes-5-945-2020>.
- P. A. Fleming, P. M. Gebraad, S. Lee, J.-W. van Wingerden, K. Johnson, M. Churchfield, J. Michalakes, P. Spalart, and P. Moriarty. Evaluating techniques for redirecting turbine wakes using SOWFA. *Renewable Energy*, 70:211–218, 2014.
- E. Gaertner, J. Rinker, L. Sethuraman, F. Zahle, B. Anderson, G. E. Barter, N. J. Abbas, F. Meng, P. Bortolotti, W. Skrzypinski, et al. IEA
560 wind TCP task 37: definition of the IEA 15-megawatt offshore reference wind turbine. Technical report, National Renewable Energy Lab.(NREL), Golden, CO (United States), 2020.
- P. M. O. Gebraad, F. W. Teeuwisse, J. W. van Wingerden, P. A. Fleming, S. D. Ruben, J. R. Marden, and L. Y. Pao. Wind plant power optimization through yaw control using a parametric model for wake effects—a CFD simulation study. *Wind Energy*, 19(1):95–114, 2016. <https://doi.org/10.1002/we.1822>.
- D. R. Houck. Review of wake management techniques for wind turbines. *Wind Energy*, n/a(n/a), 2021. <https://doi.org/10.1002/we.2668>.

- 565 P. Hulsman, M. Wosnik, V. Petrović, M. Hölling, and M. Kühn. Development of a curled wake of a yawed wind turbine under turbulent and sheared inflow. *Wind Energy Science*, 7(1):237–257, 2022. <https://doi.org/10.5194/wes-7-237-2022>.
- Á. Jiménez, A. Crespo, and E. Migoya. Application of a LES technique to characterize the wake deflection of a wind turbine in yaw. *Wind Energy*, 13(6):559–572, 2010. <https://doi.org/10.1002/we.380>.
- J. M. Jonkman, J. Annoni, G. Hayman, B. Jonkman, and A. Purkayastha. Development of FAST.Farm: A new multi-physics engineering
570 tool for wind-farm design and analysis. In *35th Wind Energy Symposium*, page 0454, 2017.
- S. Kanev, E. Bot, and J. Giles. Wind farm loads under wake redirection control. *Energies*, 13(16), 2020. <https://doi.org/10.3390/en13164088>.
- N. D. Kelley and B. J. Jonkman. Overview of the turbsim stochastic inflow turbulence simulator. Technical report, National Renewable Energy Lab.(NREL), Golden, CO (United States), 2005.
- L. A. Martínez-Tossas, J. King, E. Quon, C. J. Bay, R. Mudafort, N. Hamilton, M. F. Howland, and P. A. Fleming. The curled wake model:
575 a three-dimensional and extremely fast steady-state wake solver for wind plant flows. *Wind Energy Science*, 6(2):555–570, 2021.
- D. Medici and P. H. Alfredsson. Measurements on a wind turbine wake: 3D effects and bluff body vortex shedding. *Wind Energy*, 9(3): 219–236, 2006. <https://doi.org/10.1002/we.156>.
- D. Medici and J. Å. Dahlberg. Potential improvement of wind turbine array efficiency by active wake control (AWC). In *Proceedings of European Wind Energy Conference*, pages 65–84, Madrid, Spain, 2003.
- 580 E. Quon, P. Doubrawa, and M. Debnath. Comparison of rotor wake identification and characterization methods for the analysis of wake dynamics and evolution. In *Journal of Physics: Conference Series*, volume 1452, page 012070. IOP Publishing, 2020.
- K. Shaler and J. Jonkman. FAST.Farm development and validation of structural load prediction against large eddy simulations. *Wind Energy*, 24(5):428–449, 2021.
- K. Shaler, J. Jonkman, G. E. Barter, J. J. Kreeft, and J. P. Muller. Loads assessment of a fixed-bottom offshore wind farm with wake steering.
585 *Wind Energy*, 25(9):1530–1554, 2022.
- E. Simley, P. Fleming, N. Girard, L. Alloin, E. Godefroy, and T. Duc. Results from a wake-steering experiment at a commercial wind plant: investigating the wind speed dependence of wake-steering performance. *Wind Energy Science*, 6(6):1427–1453, 2021.
- M. A. Sprague, S. Ananthan, G. Vijayakumar, and M. Robinson. ExaWind: A multifidelity modeling and simulation environment for wind energy. In *Journal of Physics: Conference Series*, volume 1452, page 012071. IOP Publishing, 2020.
- 590 A. P. J. Stanley, J. King, C. Bay, and A. Ning. A model to calculate fatigue damage caused by partial waking during wind farm optimization. *Wind Energy Science Discussions*, 2020:1–34, 2020. <https://doi.org/10.5194/wes-2020-117>.
- J. Wagenaar, L. Machiels, and J. Schepers. Controlling wind in ECN’s scaled wind farm. *Proc. Europe Premier Wind Energy Event*, pages 685–694, 2012.
- D. S. Zalkind and L. Y. Pao. The fatigue loading effects of yaw control for wind plants. In *2016 American Control Conference (ACC)*, pages
595 537–542. IEEE, 2016.