



Field test of wake steering at an offshore wind farm

Paul Fleming¹, Jennifer Annoni¹, Jigar J. Shah², Linpeng Wang³, Shreyas Ananthan², Zhijun Zhang³, Kyle Hutchings², Peng Wang³, Weiguo Chen³, and Lin Chen³

¹National Wind Technology Center, National Renewable Energy Laboratory, Golden, CO 80401, USA

²Research & Development, Envision Energy USA Ltd, Houston, TX 77002, USA

³Research & Development, Envision Energy Ltd, Shanghai, 200051, China

Correspondence to: Paul Fleming (paul.fleming@nrel.gov)

Received: 16 January 2017 – Discussion started: 6 February 2017

Revised: 4 April 2017 – Accepted: 7 April 2017 – Published: 8 May 2017

Abstract. In this paper, a field test of wake-steering control is presented. The field test is the result of a collaboration between the National Renewable Energy Laboratory (NREL) and Envision Energy, a smart energy management company and turbine manufacturer. In the campaign, an array of turbines within an operating commercial offshore wind farm in China have the normal yaw controller modified to implement wake steering according to a yaw control strategy. The strategy was designed using NREL wind farm models, including a computational fluid dynamics model, Simulator fOr Wind Farm Applications (SOWFA), for understanding wake dynamics and an engineering model, FLOW Redirection and Induction in Steady State (FLORIS), for yaw control optimization. Results indicate that, within the certainty afforded by the data, the wake-steering controller was successful in increasing power capture, by amounts similar to those predicted from the models.

1 Introduction

Wind farm control is an active field of research in which the controls of individual turbines co-located within a wind farm are coordinated to improve the overall performance of the farm. One objective of wind farm control is improving the power production of wind farms by accounting for the wake interactions between nearby turbines.

In one wind farm control concept, turbines are yawed to introduce a deflection of the wake away from downstream turbines. This method has been referred to as “controlling the wind” (Wagenaar et al., 2012) and “yaw-based wake steering” (Fleming et al., 2014b). High-fidelity simulations of wake steering have shown the potential of this technique. Jiménez et al. (2010) used computational fluid dynamics (CFD) simulations to demonstrate the wake deflection capability of wind turbines and provided a model of this deflection. In Fleming et al. (2014b), they used National Renewable Energy Laboratory (NREL)’s CFD-based Simulator fOr Wind Farm Applications (SOWFA) to investigate the capabilities of wind turbines to redirect wakes. In Vollmer et al. (2016), the behavior of wake steering in different atmo-

spheric conditions was investigated, also using CFD. Finally, in Fleming et al. (2014a), simulations of two-turbine wind farms, again using SOWFA, were used to show that through wake steering, the net power of the two turbines is increased when the upstream turbine applies an intentional yaw misalignment.

Based on high-fidelity simulations, there appears to be good opportunities for improved power performance of wind farms with significant wake losses. Recent efforts have focused on the design of lower-fidelity, controller-oriented models, and controllers based on these models that use wake steering to actively improve power. In Gebraad et al. (2014), the FLOW Redirection and Induction in Steady State (FLORIS) model is described and used to determine optimal yaw settings for a model six-turbine wind farm. Set points for a particular wind speed and direction are determined by optimizing the yaw angles of the turbines using FLORIS, and these set points are used in SOWFA simulations. The results from SOWFA agree with the predictions from FLORIS, and total power capture is increased by 13%. This work is carried further, and FLORIS is used to assess the overall improvement from control, first for one speed and over a wind

rose of directions in Fleming et al. (2015) and then to determine the overall annual energy production in Gebraad et al. (2016). These studies indicate a good potential for improved overall annual power production for wind farms experiencing significant wake losses. It should be noted that even greater benefits can be yielded if future wind farms are designed for active control of wakes, rather than using large inter-turbine spacings to avoid wake losses. This combined optimization of wind farm control and wind farm system engineering is a subject of active research. These simulation studies have demonstrated a theoretical potential of wind farm control. However, it is often noted that issues arising in implementation in real conditions might undermine the positive results. Inaccuracies in the control-oriented models or high-fidelity simulations have been cited as a potential issue. Additional modeling of constantly changing wind direction could improve the comparison between simulation and field testing.

Some field testing of wake steering has been performed to date to understand the potential of this wind farm control strategy. In Wagenaar et al. (2012), wake steering is implemented at a scaled wind farm; however, the results are inconclusive. Wind tunnel testing of wake steering is performed in Schottler et al. (2016) and Campagnolo et al. (2016), and the results are encouraging because in each case that wake steering is observed, overall power capture is improved for two-turbine cases. The results are in alignment with earlier simulation studies conducted by Jiménez et al. (2010), Gebraad et al. (2014), and Fleming et al. (2014a) in the amount of improvement and the asymmetric relationship of yaw misalignment and power improvement. On the full scale, a nacelle-mounted lidar is used to observe wake deflection on a utility-scale turbine in Trujillo et al. (2016). Finally, in Fleming et al. (2016), two ongoing studies of wake steering are presented, which are both part of a multiyear US Department of Energy Atmosphere to Electrons project. In one of the field studies at the Scaled Wind Farm Test Facility (SWiFT), two 27 m diameter (D) turbines (intentionally aligned in the dominant wind direction at a spacing of 5 D) were used to perform a comprehensive test of wake steering. A second ongoing experiment was done, in which a rear-mounted lidar was used to monitor a wake of a utility-scale turbine, which is set to hold a specific yaw misalignment for a prolonged period of time. Initial results from both campaigns are in accord with the model predictions described earlier.

In the present study, a wind farm controller that performs wake steering is designed and implemented for an operating commercial offshore wind farm in China with Envision turbines. The control strategy is simple and based on the approach used to control the simulated wind farm in Gebraad et al. (2014). The test was run for several months and data were collected and compared from time periods when the controller was operating and not operating. The results and analysis indicate a successful improvement in power production. Additionally, the data provide important validation of the models, specifically SOWFA and FLORIS, used in the




Figure 1. This figure demonstrates the workflow of this project. In particular, this project started by running CFD models of the turbines in this field study to understand the wake characteristics. The FLORIS model was tuned to those CFD simulations. A control look-up table was generated that contained the optimal yaw settings for the control turbine. A field test was conducted with the controller on and off. The data were analyzed and the results are presented in this paper. LUT – look-up table; SCADA – supervisory control and data acquisition.

design of the controller, and can be generalized to similar CFD and control-oriented models. There are some qualifications to these results that will be fully considered in the text.

The contributions of this paper include the results and analysis from a wind farm test of wake-steering control. The project is a collaborative project between the NREL and Envision Energy, a smart energy management company and turbine manufacturer. The positive results motivate further efforts into the design and development of such control. Additionally, this paper provides evaluation on the performance of wind farm control modeling tools in their ability to predict the effect of wind farm control strategies.

2 Project overview

The goal of the project was to implement a test of yaw-based wake steering at an operating wind farm. The project was broken down into a workflow illustrated in Fig. 1, which shows the stages of work and the structure by which this paper is organized.

The first stage of work is the selection of a wind farm for use in the experiment. The Longyuan Rudong Chaojiandai offshore wind farm in Jiangsu, China, consists of turbines from multiple manufacturers undergoing various phases of construction. A selected portion of the site was studied for this effort, consisting of 25 Envision EN136/4 MW turbines incorporating a high-speed three-stage gearbox and induction generator (see Fig. 2a).

From the wind farm, a subset of turbines was selected for implementing the experiment. These turbines form the front two rows of turbines for winds coming from the northeast. The arrangement of these turbines and their names to be used throughout this paper are shown in Fig. 2.

For the campaign, a single turbine was selected to be the controlled turbine. This is turbine C1, indicated in Fig. 2. The turbine is shown to wake three particular turbines: D1, from a wind direction of 340° at a distance of 7 D; D2, from a wind direction of 51° at a distance of 8.6 D; and D3, from a wind direction of 81° at a distance of 14.3 D. A control




Figure 2. (a) Rudong wind farm used in the field study. (b) Turbine locations. The wake-steering control strategy is implemented to mitigate the wake interactions between C1, on the one hand, and D1, D2, and D3, on the other. As indicated by the diagram, C1 is the control turbine, R1 is the reference turbine, and D1, D2, and D3 are downstream turbines waked by C1. Turbines O1 and O2 are other turbines not directly used in the study but whose wakes are noticed in certain directions.

strategy was designed to improve the summed power of turbine C1 and a downstream turbine (i.e., C1 and D1, C1 and D2, and C1 and D3) through yaw misalignment of turbine C1. This approach provides three tests of wake steering at different inter-turbine distances. Of these, the pairing with turbine D1, at 7D, is most promising. This is the distance used in Fleming et al. (2014a), which showed a potential net improvement of 4.5 % in power capture in a study with the NREL 5 MW reference turbine for the waking wind direction. Finally, a reference turbine which is not waked in any of the experimental directions is chosen to provide reference signals. This is turbine R1 in Fig. 2.

In the initial phase of work, Envision shared with NREL a FAST model of the turbines (FAST is an aero-servo-elastic wind turbine simulation tool maintained by NREL; Jonkman and Buhl Jr., 2005). Additionally, Envision provided the layout of the turbines used in the campaign. Finally, Envision provided several months of supervisory control and data acquisition (SCADA) data used to evaluate and tune models in advance of the campaign. As shown in Fig. 1, these inputs were then included in the derivation of the models used in the yaw control strategy, which will be described in the next section.

3 Modeling

3.1 SOWFA

The first model to be used in this study was the SOWFA model (Churchfield and Lee, 2014). SOWFA is a wind farm simulation tool, which models the atmospheric boundary layer using CFD and then models the turbines using embedded FAST models (in the version of SOWFA used in this work). The turbines and flow interact through a two-way actuator line coupling.

SOWFA has been validated against wind farm SCADA data (see, for example, Churchfield et al., 2012). However,

validating SOWFA is not the focus of this campaign. The primary use of SOWFA in this study is generating data sets of wakes simulated from the Envision turbine, which can be used to tune the FLORIS model (and repeating the procedure that was done for the NREL 5 MW reference turbine in Gebraad et al., 2014).

Using the FAST model of the turbine, a suite of simulations is assembled. A wind condition of 8 m s^{-1} wind and 6 % turbulence (Gebraad et al., 2014) was used for this study. Simulations were then run with a single turbine operating with various amounts of yaw misalignment, as well as for two turbine cases, wherein the upstream turbine has various yaw misalignments and the downstream turbine is placed in various positions downstream and cross-stream. The trends in power for the single-turbine case and two-turbine case provide important data with which to tune the control-oriented FLORIS model, described in the next section.

3.2 FLORIS

Using a FAST model of the turbine and the data sets produced by SOWFA, it is possible to obtain a FLORIS model that can predict the power of turbines in a farm in steady state, including wake redirection. As discussed in Gebraad et al. (2014), the FLORIS model is primarily based on the Jensen model (Jensen, 1984) and the Jiménez model (Jiménez et al., 2010). In particular, FLORIS identifies three different wake zones with separate wake recovery parameters to capture the wake characteristics. In addition, FLORIS includes the Jiménez model to incorporate the wake deflection caused by yaw misalignment. FLORIS contains a set of parameters to be tuned for a given turbine, including parameters describing the turbine and wake behavior.

Some parameters of the FLORIS model can be set directly using the FAST model. Examples of this include the rotor radius and the table of power and thrust coefficients by wind

Table 1. FLORIS wake tuning.

Parameter	NREL 5 MW	Envision 4 MW
k_d	0.17	0.26
k_e	0.05	0.063
WD_{init}	-4.5°	-3.7°
pP	1.88	1.43

speed, which can be obtained by running steady wind simulations in FAST.

The remaining parameters are tuned in this work through an optimization routine that minimizes the error between the power outputs simulated in SOWFA and those predicted by FLORIS. As a design rule, the smallest set of parameters should be adjusted away from their default settings that produce a reasonable fit. Through iteration, it was determined that adjustments to four parameters gave a good approximation. The parameters that we focus on are described below.

The parameter pP relates yaw misalignment to reduction in power by

$$L_{yaw} = \cos(\gamma)^{pP}, \quad (1)$$

where γ is the yaw misalignment of the turbine and L_{yaw} is the fraction of power relative to a non-yawed baseline. The higher pP is, the more quickly a turbine loses power by misalignment and the less likely it is that wake steering will work because the downstream turbine needs to recover more power to compensate for power lost by the upstream turbine.

The parameter k_e determines the rate at which a wake expands and recovers to the free-stream velocity. A larger k_e value indicates a faster wake recovery to free stream. Standard values in the literature range from 0.05 to 0.1. Similarly, k_d , based on the Jiménez model, describes the rate at which a deflected wake reverts to the free-stream direction. A larger k_d parameter indicates that the wake is less sensitive to yaw misalignment. Standard values in the literature range from $k_d = 0.1$ to $k_d = 0.3$. Finally, WD_{init} describes an initial wake deflection angle without steering, which is important for capturing the asymmetry of wake steering. This asymmetry is likely caused by the combination of the rotation of the turbine and the shear layer in the atmospheric boundary layer.

The results of the tuning optimization are presented in Table 1, which shows both the default values obtained from tuning to the NREL 5 MW reference turbine as well as the newly obtained values.

Among these, the lower value of pP is interesting, as 1.43 is below the value obtained for the NREL 5 MW turbine (1.88) and other experimental results. For example, the coefficient is fit to wind tunnel tests in Medici (2005) to be 2.0. However, 1.43 is an attractive number as it implies that wake steering can be performed with less losses incurred on the upstream yawed turbine.

4 Control design and field test

Given a completed FLORIS model, it is now possible to derive a set of yaw misalignments for turbine C1 that will optimize power for the pairs of turbines (D1, D2, and D3 downstream) by wind direction. Wind speed is not used as an input as inspection indicated minimal sensitivity. It is important to note that there is not much benefit at very low and very high wind speeds, suggesting that it is sufficient to enable and disable the controller by wind speed rather than scheduling.

Some constraints were placed on the optimization. First, for turbine loading and safety reasons, the maximum yaw misalignment was limited to 25° . Second, it was decided for this experiment to limit the controller to positive yaw misalignment angles (in our nomenclature this is rotating the turbine counterclockwise from the wind when viewed from above). This is because this has been demonstrated to be more effective (see, for example, Fleming et al., 2014a, where positive yaw misalignments yield higher power increases as compared to negative yaw misalignments), and including negative misalignments might raise loads and require a nontrivial transition from positive to negative yaw misalignments near the wake crossover point, i.e., when it becomes more beneficial to redirect the wake from the right side of the downstream turbine to the left side of the downstream turbine. Using the tuned FLORIS model and these constraints, it is possible to now derive an optimal table of yaw misalignments for turbine C1. This is shown in Fig. 3.

The optimal yaw misalignment angles for turbine C1 are shown in the upper left of Fig. 3. The dashed lines indicate the directions in which turbine C1 wakes one of the three downstream turbines (see Fig. 2). Near the fully waked directions, the misalignments are largest and taper down as less deflection is needed to remove the partial wake overlap situations. The power loss of turbine C1, shown in normalized power, is indicated in the middle left plot. The overall “plant” gain for these four turbines is then shown in the lower left. Finally, the right panels show the normalized power of the three downstream turbines with and without wake steering. Based on the percent improvement, the sum power is expected to increase for all three pairings, meaning the gains downstream exceed the losses upstream. However, this is least so for turbine D3 as at 14.3 D, the baseline wake loss is much less, which is expected.

Using this table of offsets by wind direction, engineers at Envision modified the yaw controller of turbine C1 to deliver these offsets. Note that this offset tracking must be done within the limits of yaw control actuation, and for safety reasons the offset was disabled in sustained winds above 10 m s^{-1} . We note here also that the purpose of the experiment was to demonstrate the principle of wake steering and not produce a fully optimized closed-loop control implementation, which is expected to be the subject of future work. As will be discussed, the present controller offsets correctly in




Figure 3. FLORIS optimal yaw misalignment results. The dashed vertical line indicates the direction in which the downstream turbine is fully waked by turbine C1. See Fig. 2 for details on the precise wind directions. Specifically, the first dashed line (from left to right) refers to the direction of turbine C1 fully wakening downstream turbine D1. The second dashed line refers to the direction of turbine C1 fully wakening downstream turbine D2. Finally, the third dashed line refers to the direction of turbine C1 fully wakening downstream turbine D3. Wakes from turbines other than C1 are also evident from the dashed lines. Note that for the power subplots, the power is given normalized to the power of the turbine in yawed and un-waked conditions

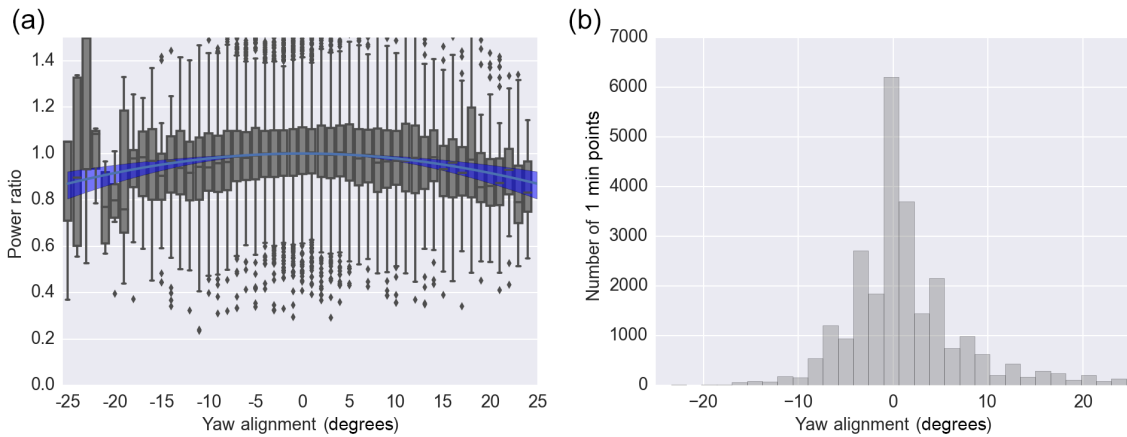


Figure 4. Cosine exponent fit. Panel (a) shows the range of data over various yaw misalignments and the power ratio. The blue line indicates a cosine function with the fit exponent value of $pP = 1.41$, and the banded region indicates the confidence interval as described in the text. Panel (b) indicates the number of points included in each yaw alignment position.

the average sense, with a wide variation of offsets occurring dynamically.

Following the implementation of the controller, the experimental campaign was then run in two phases. In the first phase, the wind farm was operated normally while data relevant for this campaign were collected. This phase lasted 4 months, from 3 April 2016 to 5 August 2016. The second phase used the controller designed above and was run an additional 4 months, from 5 August 2016 to 2 December 2016.

It is regrettable that the campaign could not be run longer. Naturally, only a portion of the data features the wind di-

rections of primary interest. Additionally, it should be noted that it would be better to alternate the controller on and off regularly throughout the campaign to better compare conditions when the controller is on than when it is off. However, Rudong is a commercial wind site with internal and external restrictions, and it was necessary to run the test within these constraints. It is important to note that there are limitations with field testing and data collection. In particular, the results are impacted because of the relatively short window of data collection. Further, the sequential testing pattern opens the possibility of confounding influences such as seasonal




Figure 5. Statistical analysis of the SCADA vane data from turbine C1, with the offset controller off and on. This graph shows the controller of turbine C1 operating to redirect the wake away from the downstream turbines. The black and red lines indicate the average yaw offset positions from SCADA data for the baseline and offset strategy, while the dashed lines are those of the look-up table (LUT) from FLORIS. Finally, the red and black regions show the range of points from the 25th through 75th percentiles for the baseline and offset SCADA data.

variation in atmospheric conditions. The results of these limitations will be discussed more in the analysis section.

5 Results and analysis

In this section, the results of the campaign are presented and analyzed. As a first step, the behavior of the upstream turbine can be analyzed independently. Specifically, the pP term, which describes the loss of power against yaw misalignment, can be derived from the experimental data and compared with the value computed from SOWFA and used in FLORIS.

The data are collected at a 20 Hz rate from the campaign and are reduced to 1 min averages. Next, we define a power ratio, which is the power of turbine C1 (the controlled turbine) divided by the power of turbine R1 (the reference turbine). We then consider this ratio as a function of the yaw misalignment of turbine C1. This misalignment is measured by the wind vane of turbine C1. Note that this calculation was repeated, wherein the misalignment is the difference between the yaw angle of turbine R1 and C1, without substantially affecting the results. The value of pP is determined as the result of a minimization problem between the data and Eq. (1). Additionally, to provide a form of confidence interval, a separate pP is similarly derived for every day of data, and the 25th and 75th percentile values are used as a range of confidence. The data, fit, confidence interval, and amount of available data are shown in Fig. 4.

Using this method, the determined value of pP was 1.41, which is close to the value derived from SOWFA. This is the first encouraging result, as the value from SOWFA initially appeared lower than expected. From a wake-steering perspective, it is useful that the upstream turbine will lose less

power when operating in yawed conditions with a smaller exponent.

Figure 5 shows the performance of the controller in obtaining the desired offsets. In the figure, for the baseline and optimized control strategies, the mean offset of turbine C1 is shown, as well as a shaded region including points (1 min averages) from the 25th through 75th percentile of observed points. Finally, the initially prescribed pattern of FLORIS is shown. What can be observed is that implementing the offset strategy in real conditions where the turbine must track an ever-changing wind direction within the bounds of the yaw controller results in the sharp narrow peak of the optimal strategy being smoothed and spread. This implies that the actual control is somewhere between the baseline and optimum in terms of performance.

Finally, Figs. 6, 7, and 8 compare the overall performance across the directions for the three two-turbine pairings. To compute this performance, the following procedure is used. First, for each wind direction (in 5° bins) and for each turbine a power curve of the form

$$P = \min(aNv^3, P_{\text{rated}}) \quad (2)$$

is computed for each day using an error-minimization technique, where a is the fitted value and N is the nominal value, which includes the air density, coefficient of power, rotor area, and efficiency losses. v is the wind speed measured by the reference turbine R1.

The value a represents a scalar gain on the power curve below rated and also effectively shifts the rated wind speed. N is selected such that a turbine operating normally and unyawed would on average have an a value of 1.0.

This method, in which the data are used to derive a reduction value to the power curve, was selected after significant

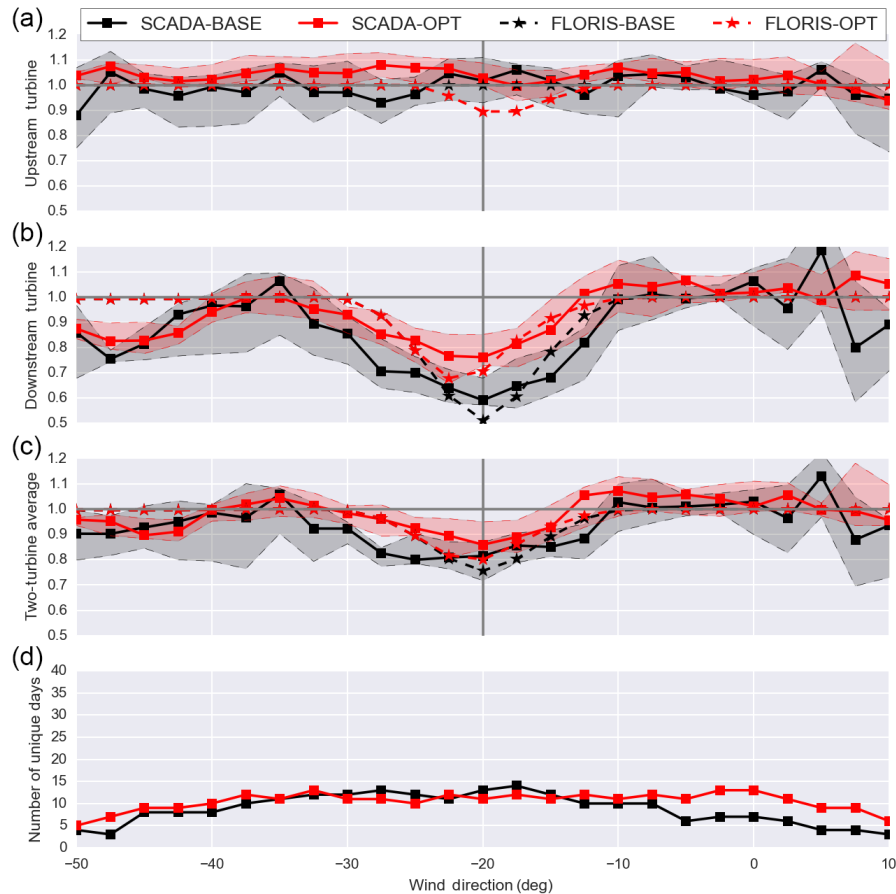


Figure 6. This figure shows the results fitting power curves for the turbine C1 to turbine D1 pair. FLORIS-BASE refers to the results from the FLORIS model with no wake steering, and FLORIS-OPT refers to the results from the FLORIS model with the optimal yaw set points. SCADA-BASE refers to the field test where the controller was off, and SCADA-OPT refers to the field test where the controller was on. The banded region shows the range of fit values from the 25th to 75th percentile of fitted power curves. In particular, (a) shows turbine C1 and the resulting power with the controller on (red) and off (black). Panel (b) shows the downstream turbine, i.e., turbine D1, with the offset controller of C1 on (red) and off (black). Panel (c) shows the overall power between the pair of turbines. There is a significant power gain when the controller is on versus when it is off. Finally, (d) indicates the number of days used to compute the statistics provided in the top three plots.

effort to combine the available data into a complete analysis. Note that alternative methods, such as directly comparing the power production of the turbines was performed and gives similar results. The power curves are fit on a per-day basis (and not for example globally or through random bootstrapping) to help visualize the day-to-day variation in performance.

The computed value a is the amount by which the power curve is reduced below nominal as observed for non-waked and non-yawed turbines. From these values computed for each available day, for each control setting, the 25th, 50th, and 75th percentile values are determined and used as the middle fit and confidence range. The reason for not computing the power curve globally was to limit the impact of particular outlier days and give some indication of the range of results observed. Where the range is small and the number of

days of data collected large, the trends converge. Finally, the power predictions of FLORIS, for the same conditions, are overlaid on the plots for comparison. The power values from FLORIS were computed by simulating the same conditions and then using the same fitting approach to include some of the effects such as blurring between wind directions in the 5° bins.

For a first analysis, consider the turbine C1–D1 pair in Fig. 6. In particular, the results shown in Fig. 6 show the results from FLORIS with no wake steering (FLORIS-BASE, shown as the dashed black line with star symbols), FLORIS with optimal yaw set points (FLORIS-OPT, shown with the dashed red line with star symbols), the field test with no wake steering (SCADA-BASE, shown with the solid black line with squares), and the field test with the optimal yaw set points (SCADA-OPT, shown with the solid red line with

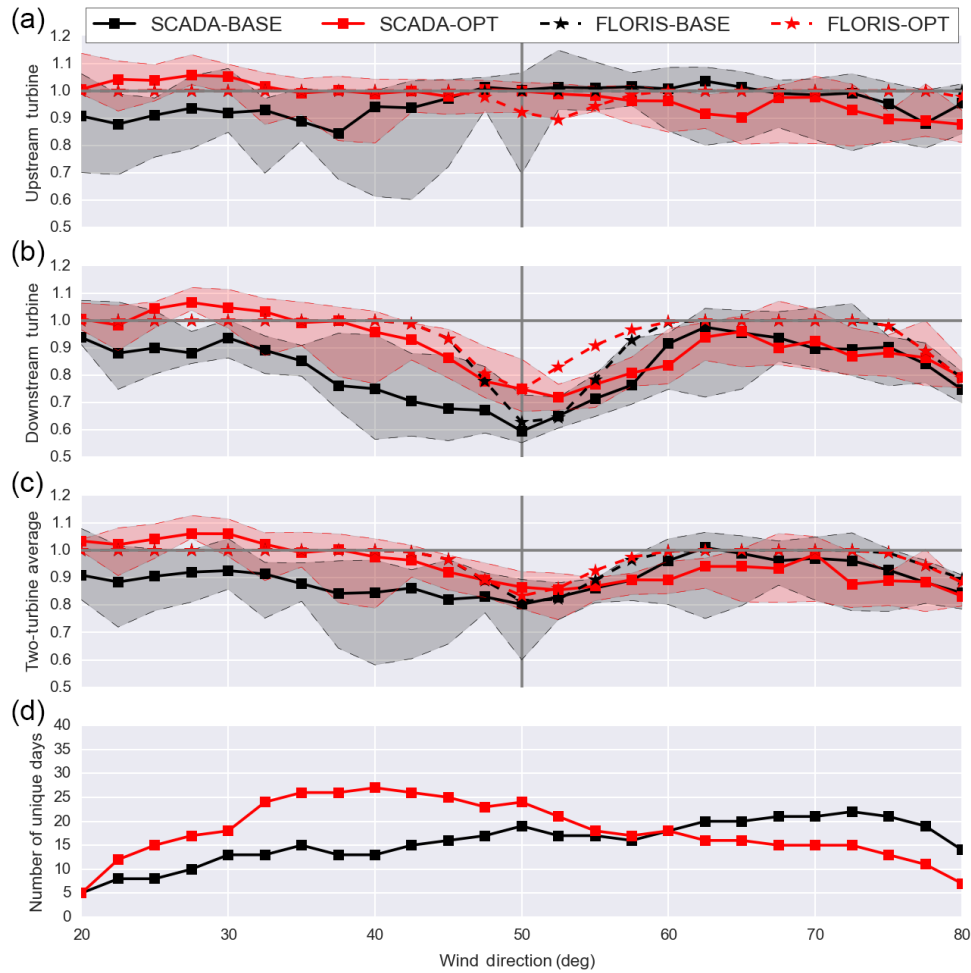


Figure 7. This figure shows the results of the turbine C1 to turbine D2 pair. In particular, (a) shows turbine C1 and the resulting power with the controller on (red) and off (black). Panel (b) shows the downstream turbine, i.e., turbine D2, with the controller on (red) and off (black). Panel (c) shows the overall power between the pair of turbines. Finally, (d) indicates the number of days used to compute the statistics provided in the top three plots.

squares). For the upstream turbine C1, it can be noted that no trend in power change can be strongly observed. This is fairly consistent with the limited power loss of the lower pP exponent and the fact the achieved yaw misalignments were, on average, not very large. The results for the 7D downstream turbine D1 are encouraging. A clear increase in power production is observed in the wind directions near the primary wake direction of -20° (or 340°). In the case of the main wake direction, the power is increased from 0.59 to 0.76 (an increase of 29%), which is less than the gain predicted by FLORIS of approximately 40%. However, it is evident that the smearing that results from various wind directions, which was evident in Fig. 5, is impacting the results as well. The overall pattern is again of a less dramatic gain; however, it is spread over a wider area. This result fits with the analysis reported in Gaumond et al. (2013), which observed that the Jensen model (and by extension Jensen-based models like

FLORIS) can predict wind farm power production more accurately if wind direction measurements are assumed to be uncertain. Incorporating this uncertainty into the FLORIS model should yield a better fit in that the wake deficits would appear more widely spread, as they do in the SCADA data. Incorporating this uncertainty into control design is the subject of ongoing research.

Turning to the power average of the two turbines, the lack of loss upstream and a good improvement downstream leads to a better than expected return for the pair. In interpreting Fig. 6, it is perhaps useful to focus as much on the shaded regions (which indicate the region containing the 25th to 75th percentile of fitted power curves) as on the solid lines with squares, which are the 50th percentile. Where the banded regions overlap least is where we see a persistent change in performance. Large changes in the line, for example around 5° in Figure 6b are not yet meaningful as the regions are

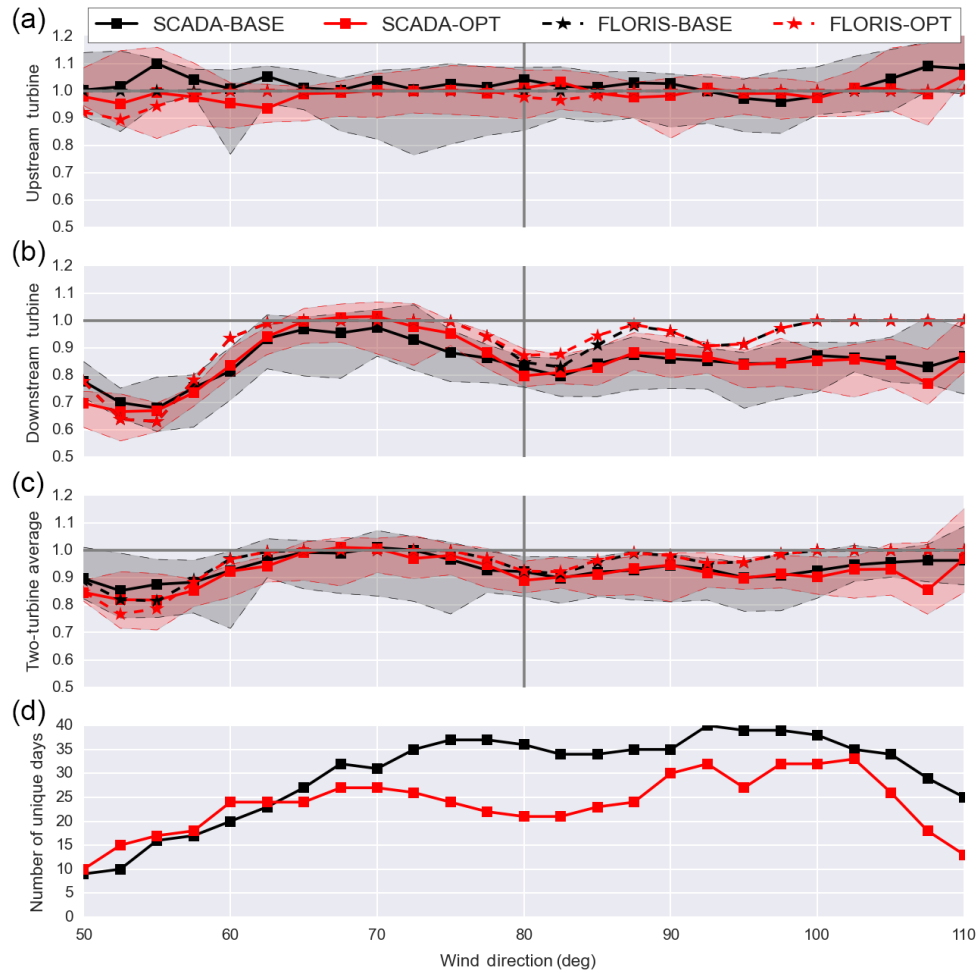


Figure 8. This figure shows the results of the turbine C1 to turbine D3 pair. In particular, (a) shows turbine C1 and the resulting power with the controller on (red) and off (black). Panel (b) shows the downstream turbine, i.e., turbine D3, with the controller on (red) and off (black). Panel (c) shows the overall power between the pair of turbines. There is no noticeable power gain with the controller on or off. This is likely because of the increase in spacing between turbines C1 and D3 (14.3 D) in comparison to the spacing between turbine C1 and D1 (7 D). Finally, (d) indicates the number of days used to compute the statistics provided in the top three plots.

completely overlapping. That the banded regions of the two-turbine average contain significant nonoverlapping regions around the wake control direction is probably the main positive finding of this paper. It suggests that the power improvement is consistent.

Finally, Fig. 6d shows the number of separate days available for computing the data in each bin. It is useful to note that these do not indicate full days but that some data were collected on a given day, which was in this direction, and a separate power curve computed. A value of 10, for example, indicates that 10 days were used for a given bin for a given controller (sharing no data points), and these values were used to produce the statistics shown in the plot.

In the wind directions in which wake control is not active (to the left and right of the plot), there is no persistent trend between the baseline and optimized control on the down-

stream turbine. This is also mostly true where the turbine is waked by the reference turbine R1 at -47° . This helps to confirm that the change observed is caused by wake deflection, rather than a difference in atmosphere between testing periods being the cause of the underlying changes in wake behavior. This observation could be made for turbine D2 and 80° in Fig. 7, when it is waked by turbine O2 (refer to Fig. 2), and turbine D3 at 55° in Fig. 8, when it is waked by reference turbine R1. This last case is especially compelling, as the spacing is 8.5 D, which is similar to the space between C1 and D2 and there is a comparable amount of data. However, unlike D1 and D2 when they are behind the controlled turbine C1, there is no improvement in power production.

For the pairing of turbine C1 and D2 in Fig. 7, the spacing is now 8.5 D and wake steering is expected to become more challenging. Nevertheless, the combination yields an

improvement in the main wake direction (50°) for the downstream turbine and the two as a pair. When stepping away from the main wake direction of 50° in either direction, however, things become ambiguous. To the left, we see the amount of data for the baseline case grows smaller and the spread in results for the downstream turbine D2 grows larger (observing the large gray regions). The power is low despite no wake (although this is a wind direction in which the inflow to turbine D2 runs in between turbine R1 and turbine C1.) Therefore, it is probable that a lack of data plays a part, and we would expect little change here. To the right, the trend in power goes negative for the downstream turbine and the pair. Nevertheless, it should be noted that the spread of results, indicated by the bands, is completely overlapping, and so the significance of this impact cannot be established.

Finally, observing the results of the turbine C1–D3 pair in Fig. 8, at a spacing of 14 D, little improvement is expected and basically none is observed. It is useful to note that turbine D3 has the most data collected and the results seem best converged. As noted earlier, when turbine D3 is waked by turbine R1 at 54° , it is the deeper wake, having a spacing of 8.5 D, and no noticeable change in wake loss occurs (the power of the downstream turbine is the same in the baseline and optimized cases); this points to wake steering being the primary cause of change in power for the earlier cases at 7 and 8.5 D between the baseline and optimized cases.

Also, it is interesting to consider what is happening west of 80° for the downstream turbine D3. FLORIS predicts a return to full power, followed by a dip around 95° when turbine O2 is upstream. However, what is actually observed is a reduction in power basically across the whole range. Considering Fig. 2, this is a range of wind direction without an obvious single-turbine wake but with four turbines still upstream. Notice that unlike an explanation of shallower but more spread loss from wind direction uncertainty, this deficit is deeper than what is predicted by FLORIS. This deficit speaks to an unmodeled deep array effect that may prove important to include when FLORIS is used to model and design controllers for multi-turbine arrays.

6 Conclusions and future work

This study provided several encouraging, albeit qualified results. The main result was that for the directions and spacings (7 D and 8.5 D) expected to produce an improvement in power for the pair of turbines, such an improvement is observed. The most easily interpreted results come from the closer 7 D spacing, whereas the 8.5 D spacing has some changes that are partly caused by limited data availability. Another good result was the observed agreement between the lower than expected power loss with yaw function predicted by SOWFA and the loss derived experimentally. This result is positive because it provides another form of validation for SOWFA (data sets of utility turbines operating misaligned

are not commonly available for testing), and the low power loss value makes wake steering in general more successful.

In this paper, it was discussed that the primary limitations were constraints placed on the amount of data that could be collected at a commercial wind farm and constraints placed on the wind turbine yaw controller's ability to control. In the case of the pairing with turbine D3, the control and optimal cases are very close in midpoint and range, suggesting that 30–40 separate days of testing is a good target per controller. Had it been possible to toggle between control set points, it is very likely that this number could be reduced although toggling creates some issues of transition. On the controller side, it would appear that the implemented controller was sufficient to secure power gains; however, it could be that a more advanced controller could achieve more.

Also an important concern is the impact this control will have on turbine loads, both of upstream and downstream turbines. This question was outside the scope of the present campaign, and the turbine load data were not available for this joint study. However, there are ongoing research campaigns which seek to understand and quantify the impacts to loading. Already published are studies which consider the impact on loads from operating in yaw using aero-servo-elastic wind turbine models (Zalkind and Pao, 2016) and computational fluid dynamics (Schulz et al., 2016). Additionally, there is a currently ongoing field test in which a utility-scale wind turbine which is highly instrumented with load sensors is operated intentionally in yaw to assess impacts (Fleming et al., 2017). Finally, in order to systematically understand the impact on loads of turbines downstream from wake steering, new engineering tools are being developed, which include mid-fidelity models of wakes, such that loads can be simulated, but which are computationally inexpensive enough for load suites to be computed (Jonkman et al., 2017).

More generally, the design of closed-loop control systems for wake steering remains an open topic of research. The present method can be regarded as open loop because the actual wake to be controlled is never observed or estimated by the controller. Research that uses lidar to track and control wakes (Raach et al., 2016) or uses estimation techniques (Doekemeijer, 2016) may very well improve upon these first results. This is also the subject of ongoing multiyear research projects in the United States (A2e, 2016) and Europe (CL-Windcon, 2016).

Data availability. No data sets were used in this article.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. The authors would like to acknowledge and thank Matthew Churchfield for his support in the use of SOWFA and the development of the simulations. Additionally, the authors thank Pieter Gebraad for his work on this effort during his tenure at NREL.

The US Government retains, and the publisher, by accepting the article for publication, acknowledges that the US Government retains, a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work or to allow others to do so, for US Government purposes.

Edited by: S. Aubrun

Reviewed by: two anonymous referees

References

- A2e: Atmosphere To Electrons, <http://energy.gov/eere/wind/atmosphere-electrons>, last access: 15 December 2016.
- Campagnolo, F. P., Vlaho Bottasso, C. L., and Croce, A.: Wind tunnel testing of wake control strategies, in: American Control Conference (ACC),
- Churchfield, M., Lee, S., Moriarty, P., Martinez, L., Leonardi, S., Vijayakumar, G., and Brasseur, J.: A large-eddy simulation of wind-plant aerodynamics, in: 50th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition, Nashville, TN, Nashville, Tennessee, USA, 2012. American Automatic Control Council (AACC), 513–518, 2016.
- Churchfield, M. and Lee, S.: NWTC Information Portal (SOWFA), <https://nwtc.nrel.gov/SOWFA> (last access: 1 February 2017), 2014.
- CL-Windcon: CL-Windcon project kicked-off in Pampolna, <http://www.ifb.uni-stuttgart.de/news/CL-Windcon-Kickoff>, last access: 15 December 2016.
- Doekemeijer, B. M.: Enhanced Kalman filtering for a 2D CFD Navier-Stokes wind farm model, Ph.D. thesis, Delft University of Technology, 2016.
- Fleming, P. A., Gebraad, P. M. O., Lee, S., van Wingerden, J. W., Johnson, K., Churchfield, M., Michalakes, J., Spalart, P., and Moriarty, P.: Simulation comparison of wake mitigation control strategies for a two-turbine case, *Wind Energy*, 18, 2135–2143, 2014a.
- 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}, *Renew. Energ.*, 70, 211–218, doi:10.1016/j.renene.2014.02.015, 2014b.
- Fleming, P. A., Ning, A., Gebraad, P. M. O., and Dykes, K.: Wind plant system engineering through optimization of layout and yaw control, *Wind Energy*, 19, 329–344, 2015.
- Fleming, P., Churchfield, M., Scholbrock, A., Clifton, A., Schreck, S., Johnson, K., Wright, A., Gebraad, P., Annoni, J., Naughton, B., Berg, J., Herges, T., White, J., Mikkelsen, T., Sjöholm, M., and Angelou, N.: Detailed field test of yaw-based wake steering, *J. Phys. Conf. Ser.*, 753, 052003, doi:10.1088/1742-6596/753/5/052003, 2016.
- Fleming, P., Annoni, J., Scholbrock, A., Quon, E., Dana, S., Schreck, S., Raach, S., Haizmann, F., and Schlipf, D.: Full-Scale Field Test of Wake Steering (Under Review), in: Wake Conference, Visby, Sweden, 2017.
- Gaumond, M., Réthoré, P.-E., Ott, S., Peña, A., Bechmann, A., and Hansen, K. S.: Evaluation of the wind direction uncertainty and its impact on wake modeling at the Horns Rev offshore wind farm, *Wind Energy*, 17, 1169–1178, 2013.
- Gebraad, P., Teeuwisse, F., van Wingerden, J., Fleming, P., Ruben, S., Marden, J., and Pao, L.: Wind plant power optimization through yaw control using a parametric model for wake effects – a CFD simulation study, *Wind Energy*, 19, 95–114, 2014.
- Gebraad, P., Thomas, J. J., Ning, A., Fleming, P., and Dykes, K.: Maximization of the annual energy production of wind power plants by optimization of layout and yaw-based wake control, *Wind Energy*, 20, 97–107, doi:10.1002/we.1993, 2016.
- Jensen, N. O.: A note on wind generator interaction, Tech. Rep. Risø-M-2411, Risø National Laboratory, 1984.
- Jiménez, Á., 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, 2010.
- Jonkman, J. M. and Buhl Jr., M. L.: FAST Manual User's Guide, NREL report No. NREL/EL-500-38230, 2005.
- Jonkman, J. M., Annoni, J., Hayman, G., Jonkman, B., and Purkayastha, A.: Development of FAST, Farm: A New Multi-Physics Engineering Tool for Wind-Farm Design and Analysis, in: 35th Wind Energy Symposium, 0454, 2017.
- Medici, D.: Experimental studies of wind turbine wakes: power optimisation and meandering, Ph.D. thesis, KTH, Mechanics, qC 20101018, 2005.
- Raach, S., Schlipf, D., Borisade, F., and Cheng, P. W.: Wake redirecting using feedback control to improve the power output of wind farms, in: American Control Conference (ACC), American Automatic Control Council (AACC), 1387–1392, 2016.
- Schottler, J., Hölling, A., Peinke, J., and Hölling, M.: Wind tunnel tests on controllable model wind turbines in yaw, in: 34th Wind Energy Symposium, 1523, 2016.
- Schulz, C., Letzgus, P., Lutz, T., and Krämer, E.: CFD study on the impact of yawed inflow on loads, power and near wake of a generic wind turbine, *Wind Energy*, 20, 253–268, 2016.
- Trujillo, J. J., Seifert, J. K., Würth, I., Schlipf, D., and Kühn, M.: Full-field assessment of wind turbine near-wake deviation in relation to yaw misalignment, *Wind Energ. Sci.*, 1, 41–53, doi:10.5194/wes-1-41-2016, 2016.
- 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 Energ. Sci.*, 1, 129–141, doi:10.5194/wes-1-129-2016, 2016.
- Wagenaar, J. W., Machielse, L., and Schepers, J.: Controlling wind in ECN's scaled wind farm, in: EWEA Annual Meeting, Copenhagen, Denmark, 2012.
- Zalkind, D. S. and Pao, L. Y.: The Fatigue Loading Effects of Yaw Control for Wind Plants, in: American Control Conference, Boston, MA, 2016.