

SANDIA REPORT

SAND2016-4484

Unlimited Release

Printed May 2016

Visualizing Wind Farm Wakes Using SCADA Data

Shawn B. Martin, Carsten H. Westergaard, Jonathan R. White, and B. Karlson

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185-0932 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.



Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@osti.gov
Online ordering: <http://www.osti.gov/scitech>

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Rd
Alexandria, VA 22312

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.gov
Online order: <http://www.ntis.gov/search>



Visualizing Wind Farm Wakes Using SCADA Data

Shawn B. Martin¹, Carsten H. Westergaard², Jonathan R. White², and B. Karlson²

¹Software Systems R&D, ²Wind Energy Technologies

Sandia National Laboratories

P.O. Box 5800

Albuquerque, New Mexico 87185

smartin@sandia.gov, chweste@sandia.gov, jonwhit@sandia.gov, bkarlso@sandia.gov

Abstract

As wind farms scale to include more and more turbines, questions about turbine wake interactions become increasingly important. Turbine wakes reduce wind speed and downwind turbines suffer decreased performance. The cumulative effect of the wakes throughout a wind farm will therefore decrease the performance of the entire farm. These interactions are dynamic and complicated, and it is difficult to quantify the overall effect of the wakes. This problem has attracted some attention in terms of computational modelling for siting turbines on new farms, but less attention in terms of empirical studies and performance validation of existing farms.

In this report, Supervisory Control and Data Acquisition (SCADA) data from an existing wind farm is analyzed in order to explore methods for documenting wake interactions. Visualization techniques are proposed and used to analyze wakes in a 67 turbine farm. The visualizations are based on directional analysis using power measurements, and can be considered to be normalized capacity factors below rated power. Wind speed measurements are not used in the analysis except for data pre-processing. Four wake effects are observed; including wake deficit, channel speed up, and two potentially new effects, single and multiple shear point speed up. In addition, an attempt is made to quantify wake losses using the same SCADA data. Power losses for the specific wind farm investigated are relatively low, estimated to be in the range of 3-5%.

Finally, a simple model based on the wind farm geometrical layout is proposed. Key parameters for the model have been estimated by comparing wake profiles at different ranges and making some ad hoc assumptions. A preliminary comparison of six selected profiles shows excellent agreement with the model. Where discrepancies are observed, reasonable explanations can be found in multi-turbine speedup effects and landscape features, which are yet to be modelled.

ACKNOWLEDGMENTS

This work is supported and made possible by the Department of Energy (DOE) Wind and Water Power Program. Wind farm SCADA data was provided by a strategic industrial partner.

CONTENTS

1. Introduction.....	11
2. Materials and Methods	12
2.1. Data Corrections	13
Wind Speed Screening	13
Met Tower Correction.....	13
Nacelle Direction Correction	14
Pitch Correction.....	15
Power Correction.....	16
2.2. Calculations	17
Turbine-Turbine Power Curve Pair Analysis.....	17
Directional Power Performance and Variance Plots	18
Wind Shadow Calculation.....	19
Power Observer Calculation.....	20
2.3. Visualization	22
Colored Polar Plots	22
2.4. Model	23
Fitted Parameters.....	23
Wake Model	26
3. Results.....	30
3.1. Visualization	30
Detailed Wake Analysis	31
Wake Deficit Profiles	32
Channel Speedup Profiles	34
Shear Point Speedup Profiles	35
Shear Speedup for Multiple Upstream Turbines.....	36
3.2. Wake Loss Quantification	38
3.3. Modelling Wakes.....	40
4. Conclusions and Future Work	43
5. References.....	44
Distribution	45

FIGURES

Figure 1. Wind Farm Layout. The relative positions of the turbines are shown, with turbines numbered from 1-67, and the met tower marked M. The site wind rose is shown in the upper left. Turbines in close proximity are connected by lines: turbines within 5 rotor diameters are connected using red lines; turbines between 5 and 6 rotor diameters are connected using blue lines; and turbines between 6 and 7 rotor diameters are connected using black lines. Icons were taken from the Map Icons Collection (<http://mapicons.nicolasmollet.com>) and are licensed under Creative Commons Attribution (3.0).....11

Figure 2. Met Tower Wind Direction Correction. On the left, power variability of turbine 64 is plotted against the met mast measured wind direction θ_m for winds from the south. A peak in power variability is seen at approximately 169° . In reality, this peak should occur at the geographical bearing from turbine 64 to turbine 67, which is shown as a red line, occurring at 180° . Therefore, the wind direction offset $\theta_f = \theta_t - \theta_m \approx 11^\circ$. On the right, power variability is plotted against the corrected wind direction, showing an alignment between the peak power variability and the geographic bearing. Throughout this figure, the power variability curves were computed across wind direction bins of 1° 13

Figure 3. Nacelle Direction Correction for Turbine 6. On the left (a), the residual $\theta_n - \theta_m$ of the nacelle wind direction measurement θ_n versus the met tower measurement θ_m is shown through time for turbine 6. In addition to the additive biases that can be observed, there are several noticeable changes in the measurements at different time points. In the middle (b), the time periods are clustered using a change detection algorithm so that an additive correction can be applied. The separation of the last two groups (green and blue) is due to a period of erratic measurements which can be observed in the uncorrected data as a solid vertical line. On the right (c), additive corrections are applied to the time period clusters and individual measurements significantly different from the mean are removed. 14

Figure 4. Pitch Correction. On the left, the pitch schedule is shown as the variation in mean blade pitch against wind speed. Data points more than one standard deviation from the pitch schedule (indicated by the dashed lines) are removed. On the right, the power curve of the corrected data is shown (power is shown on a normalized scale). Note that several abnormal operating modes were removed, including de-rated periods. 15

Figure 5. Power Corrections. On the left (a), the median power curve was computed using wind speed bins of 1 m/s, shown as a red line. Any data greater outside of the 60th percentile was then removed. This correction directly removes abnormal modes of turbine operation from consideration for further analysis. On the right (b), the full collection of median power curves (67) exhibit a large degree of uniformity over the wind farm, indicating that the power corrected data is suitable for estimating power losses due to wake effects on the farm. 15

Figure 6. Wake Effect for Turbine 7. On the left (a), power curves are shown for the upwind turbine 6 and the downwind turbine 7 (given westerly winds). These curves show no wake effect because the individual nacelle wind speed measurements were used, which are relative. On the right (b), power curves are shown for the same two turbines, this time using the upwind nacelle sensor to measure wind speed. These curves show a wake effect, because the wind speed is now absolute for both turbines. Throughout this figure, the power curves were obtained by averaging the nacelle power measurements over bins with width of 1 m/s. 16

Figure 7. Normalized Instant Power Plots. These plots show normalized instant power averages over time versus wind direction for turbines 6 and 7. The plots are arranged to mirror the positions of the turbines in the wind farm, so that turbine 6 is west of turbine 7. Thus, given a westerly wind, turbine 6 is upwind and turbine 7 is downwind. In this case, turbine 6 is overperforming relative to turbine 7, so that a bump is observed in the plot for turbine 6 at the angle 270° . Similarly, turbine 7 is underperforming due to the wake effect, and a dip is observed in the

plot for turbine 7 at the angle 270° . Both plots are obtained by averaging normalized instant power over time within wind direction bins of 1°17

Figure 8. Power Variability Plots. These plots show power variability over time against wind direction for turbines 6 and 7. They are again arranged to mirror the positions of the turbines in the wind farm. Wake effects can be observed as increased variability, which manifest as bumps on the plots. Thus there are bumps at 90° (easterly wind) for turbine 6 and 270° (westerly wind) for turbine 7. Both plots are obtained by averaging power variability over time within wind direction bins of 1°18

Figure 9. Wind Shadow Estimates. On the left (a), wind shadows are computed for wind blowing out of the north-east using $k = 0.075$ and $x^\infty = 15$ (rotor diameters). Turbines are marked with blue dots and wind shadows are shown using gray quadrilaterals. Turbines not in wind shadows are highlighted with red circles. On the right (b), wind shadows are computed using the same wind direction but with an extremely high $k = 0.4$ and $x^\infty = 20$. The wind front computed in the two examples shows how more or less conservative wind front estimates can be obtained by varying k and x^∞19

Figure 10. Power Observer versus Actual Power. Shown here is the power observer calculation (red) compared to the actual power measurements (blue). The curves are totaled over all turbines in the farm. For this estimate, $k = 0.075$ and $x^\infty = 15$ (rotor diameters).20

Figure 11. Wake Effect Visualizations. On the left (a), visualization is shown using normalized instant power, and on the right (b) visualization is shown using power variability. In both cases, the rose plots are positioned in place of the turbine icons seen in Figure 1. Further, the rose plots are colored according the radial magnitude. For the instant power plots, over-performing turbine directions are colored red, while under-performing turbine directions are colored blue. Grey circles show average performance (instant normal power value of 1). For the power variability plots, high variability directions are colored red, and low variability directions are colored blue.22

Figure 12. Google Earth Visualization. This visualization shows the wind farm wake effects using the normalized instant power plots, complete with the labels from Figure 1. Grey circles show average performance (instant normal power value of 1). The wind farm terrain imagery (not shown) can also be examined for correlations between performance and local topography.22

Figure 13. Peak Deficits. The peak deficit for all of the profiles plotted as an average of the nearest five points to the bearing against Jensen's Park model to the 3rd power, with $k = 0.075$ and $CT = 0.7$. The profiles from each bin are colored coded.23

Figure 14. Wake Profiles from 3D to 12D. On the top (a), the average normalized power wake profiles for the 1.25D bins are shown. There are approximately 16 profiles in each bin. On the bottom (b), power variance of the wake profiles for the 1.25D bins is shown.24

Figure 15. Fitted Parameters. On the top (a), we show wake deficit and centerline for the normalized power wake profiles. The fitted values outside the wake are given by

$a(xD) = 1 + 25/D \times \text{Weibull}(x/D, 10)$. In the middle (b), we show the measured peak power variability above the ambient level, with a fitted value. On the bottom (c), we show curves fitting wake width in linear coordinates for power and power variability. For reference, the Park model wake width is plotted for $k = 0.075$. The power performance fit is $wx = 0.11 \times x/D + 0.5$ and the variability fit is $bx = 0.15 \times x/D + 0.5$ 25

Figure 16. Initial Model. On the top (a), individual wake components for turbine 11 are modelled as Gaussian bells. On the bottom (b), all wakes for turbine 11 are modelled as the product of the individual Gaussian bells. Three of the wake directions have been adjusted a few degrees to align with the experimental data. 27

Figure 17. Cosine Model. On the top (a), individual wake components for turbine 11 are modelled using a Gaussian bell and a damped cosine overlay. On the bottom (b), all wakes for turbine 11 are modelled as the product of the individual Gaussian bells using the damped cosine overlay. 27

Figure 18. Power Variability Model. On the top (a), individual wake components for turbine 11 are modelled using Gaussian bells. On the bottom (b), all wakes for turbine 11 are modelled as the sum of the individual Gaussian bells. 28

Figure 19. Overall Wind Farm Performance. On the left (a), the power averaged over all turbines versus wind direction is shown. As expected, this polar plot is highly correlated with prevailing wind directions (NW and S), although there is also an interesting spike in the NE direction. In the middle (b), the average power variability versus wind direction is shown. The power variability is aligned with prevailing winds as well as a large number of closely positioned E-W turbine pairs. On the right (c), a radial histogram of counts for turbine-turbine pairs within 7 rotor diameters is shown. 29

Figure 20. Average Turbine Performance over Wind Farm. On the left, the wind direction averaged normalized instant power is shown for each turbine, where red indicates high performance, and blue indicates low performance. On the right, the wind direction averaged power variability is shown for each turbine, where red indicates high variability, and blue indicates low variability. 30

Figure 21. Wake Effects. On the left (a), three wake effects can be observed using instant normalized power plots for turbines 8, 10, and 11. Wake deficits can be seen as dips in the power production when turbine 11 is in the shadow of turbines 8 or 10; a speed up channel can be seen as a peak in the power production when turbine 11 is facing the midpoint of turbines 8 and 10; and shear point speedups can be seen when turbine 11 is tangent to the wake of turbine 10 or 8. On the right (b), the corresponding variability in power is shown. 31

Figure 22. Normalized Instant Power for Waked Turbines. The x -axis gives the wind direction degree offset normalized by the distance between the turbines (e.g. $xD = \theta^\circ \pi 180^\circ / (xD)$, where θ° is the offset in degrees, x is the distance between the two turbines, and D is the rotor diameter). For example, when the offset is 0° the wind is blowing straight from the upwind to the downwind turbine. The y -axis shows the normalized instant power for the downwind turbine. The turbine

pairs selected are given in the legend, where the notation $d \leftarrow u (r)$ indicates the downwind turbine (d), the upwind turbine (u) and the distance (r) in rotor diameters. See also Figure 1 for turbine positions.....32

Figure 23. Power Variability for Waked Turbines. As in Figure 22, the x -axis gives the wind direction degree offset from the upwind turbine. The y -axis gives the power variability of the downwind turbine.....32

Figure 24. Power Deficit Versus Distance. Maximum instant normalized power and minimum power variance for a downstream turbine were collected for 854 turbine pairs. On the left (a), the maximum normalized instant power is plotted versus distance behind the upwind turbine. The semi-empirical wind velocity deficit from Equation 4 is shown as a solid curve. On the right (b), the minimum power variance for the downstream turbine versus distance is shown.....33

Figure 25. Normalized Instant Power for Channel Speedup. The x -axis gives the wind direction degree difference from the downstream turbine to the midpoint between the two upstream turbines. The y -axis gives the normalized instant power. The channel speed up effect can be seen as the peak at 0° , flanked by wake deficits on either side of the peak.....34

Figure 26. Power Variability for Channel Speedup. As in Figure 25, the x -axis gives the wind offset from the midpoint between the two upstream turbines. The y -axis gives the power variability. The channel speed up is associated with low variability.....34

Figure 27. Normalized Instant Power for Shear Speedup. The x -axis gives the wind direction degree offset from the bearing between and upwind and downwind turbine pair, and the y -axis show the normalized instant power. Turbines were selected so that the region between 0° to 90° from the downwind turbine is undisturbed (no turbines nearby). The shear point speedup can be seen as improved power production between 15° and 20° , as indicated by the dotted line at 17.5° .
.....35

Figure 28. Power Variability for Shear Speedup. As in Figure 28, the x -axis gives the wind direction degree offset from the upwind turbine. The y -axis shows power variability.....35

Figure 29. Normalized Instant Power Increase from Multiple Upstream Turbines. The normalized instant power profiles are shown for 6 turbines offset from an upstream row of turbines. The upstream turbine rows are all arranged from north to south.....36

Figure 30. Power Variability from Multiple Upstream Turbines. The power variability profiles are shown for 6 turbines offset from an upstream row of turbines. The upstream turbine rows are arranged North to south.....36

Figure 31. Percent Gain Parameter Search. Percent gain (G) is shown as a function of parameters in the range $k = 0.01$ to $k = 0.5$ and $x^\infty = 5$ to $x^\infty = 30$ rotor diameters.....37

Figure 32. Region II Percent Gain Parameter Search. The percent gain of the wind farm is uniformly higher using wind restricted speeds from 6-10 m/s, as compared to wind speeds from

4-20 m/s (Figure 31). Parameters are again in the range $k = 0.01$ to $k = 0.5$ and $x^\infty = 5$ to $x^\infty = 30$ rotor diameters.....38

Figure 33. Region III Percent Gain Parameter Search. The percent gain of the wind farm is minimal using wind restricted speeds from 10-20 m/s, as compared to wind speeds from 4-10 m/s (Figure 31 and Figure 32). Parameters are again in the range $k = 0.01$ to $k = 0.5$ and $x^\infty = 5$ to $x^\infty = 30$ rotor diameters.....39

Figure 34. Power roses for turbines 11, 8 and 10. In this figure, we show the power roses as computed using the wake model from Section 2.4. On the left(a), we show the power rose for turbine 11; in the middle (b), the power rose for turbine 8; and on the right (c), the power rose for turbine 10.....40

Figure 35. Power variability roses for turbines 11, 8 and 10. In this figure, we show the power variability roses as computed using the wake model from Section 2.4. On the left(a), we show the power variability rose for turbine 11; in the middle (b), the power variability rose for turbine 8; and on the right (c), the power variability rose for turbine 10.....41

Figure 36. Power roses for turbines 39, 40, and 54. In this figure, we show the power roses as computed using the wake model from Section 2.4. On the left(a), we show the power rose for turbine 39; in the middle (b), the power rose for turbine 8; and on.....41

1. INTRODUCTION

The wind energy used by a turbine to produce electrical power causes a reduction in wind speed behind the turbine, also known as a wake or wind shadow [1]. The wake behind a turbine will dissipate with distance, but will affect nearby downwind turbines. For a large wind farm, these effects will accumulate, resulting in potentially significant aggregate power loss over the entire farm. In addition, wakes are turbulent, and can stress downwind turbines, possibly affecting the reliability and lifetime of turbines within the farm.

To mitigate the effects of turbine wakes in a wind farm, it is important to carefully site turbines during the design phase [1]. Research into siting wind turbines has employed semi-empirical numerical models of turbine wakes [2-5], as well as more exact models based on the Navier-Stokes equations [6]. Power losses due to wake effects have also been modeled [7,8]. While these models provide valuable insight and guidelines, they are nevertheless incomplete and cannot address every practical concern in the siting of turbines when building a wind farm. Further, there are numerous existing wind farms, for which re-siting is not an option, but which might yield greater power production with improved understanding of on-site wake effects.

An alternative to using computational models for understanding and mitigating wake effects is to analyze data gathered from existing farms using techniques from statistics and data mining. Such approaches have yielded results in the areas of turbine failure prediction and condition monitoring, see for example [9-13], but have only been applied to limited degree in wake analysis [14].

In this report, Supervisory Control and Data Acquisition (SCADA) data is analyzed with the goal of documenting wake effects on a functioning industrial scale wind farm. SCADA data from this farm is summarized, corrected, transformed, and analyzed in order to provide an estimate of potential power loss due to wake effects. Finally, a simple model based on the wind farm geometrical layout is proposed.

2. MATERIALS AND METHODS

Data was gathered over a 1.5 year period in 2012 and 2013 from the SCADA system at an onshore wind farm in the United States. The wind farm included 67 horizontal axis, three bladed, variable pitch turbines, along with one meteorological (met) tower. In this study, analysis was performed on subset of the data collected from the turbines: nacelle wind speed, nacelle direction (position), rotor speed, blade pitch, and power output. The met tower collected data on temperature, air pressure, wind speed, and wind direction. The layout of the wind farm is shown in Figure 1, along with a wind rose showing the prevailing wind directions.

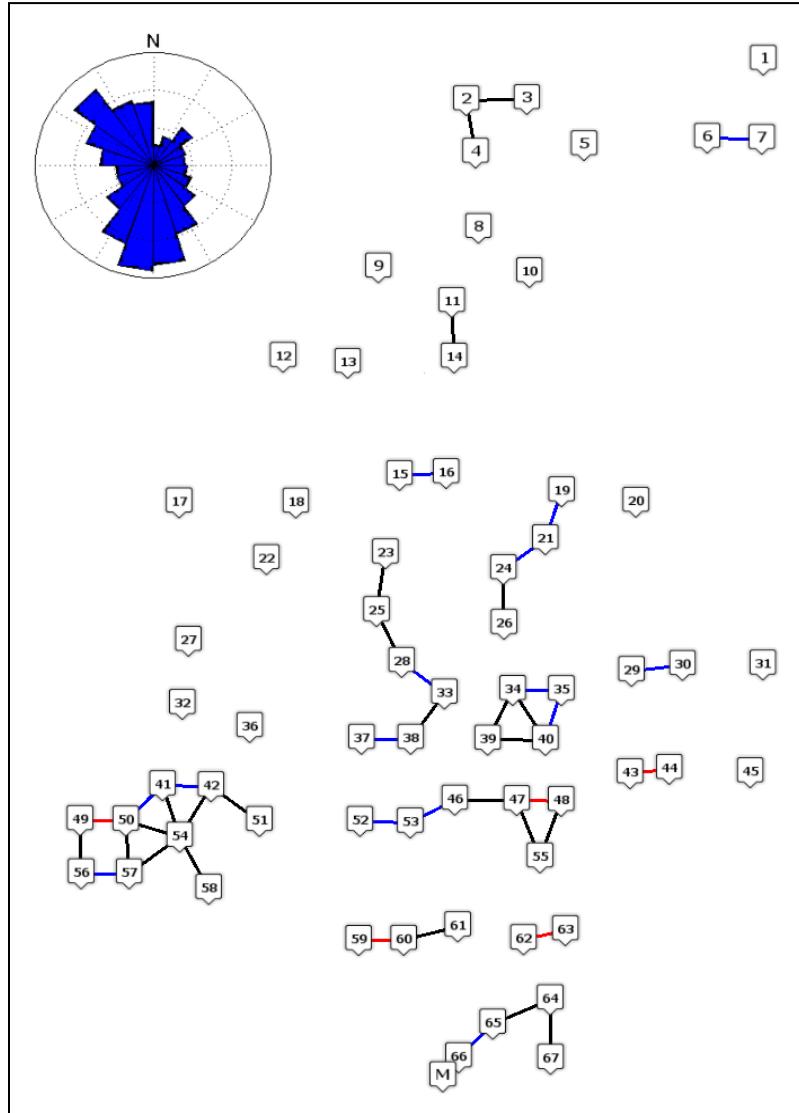


Figure 1. Wind Farm Layout. The relative positions of the turbines are shown, with turbines numbered from 1-67, and the met tower marked M. The site wind rose is shown in the upper left. Turbines in close proximity are connected by lines: turbines within 5 rotor diameters are connected using red lines; turbines between 5 and 6 rotor diameters are connected using blue lines; and turbines between 6 and 7 rotor diameters are connected using black lines. Icons were taken from the Map Icons Collection (<http://mapicons.nicolasmollet.com>) and are licensed under Creative Commons Attribution (3.0).

Data was collected continuously (every 2 seconds) but was summarized over 10 minute intervals prior to analysis. For each variable collected (e.g. wind speed), the minimum, maximum, average and standard deviation over the 10 minute interval was computed. After summarization, there were approximately 61,000 time points per turbines available for analysis.

2.1. Data Corrections

Wind Speed Screening

We required an operational turbine and wind speed to lie in the range of 4-20 m/s, yielding approximately 46,000 time points per turbine, corresponding to a site average wind speed of approximately 8 m/s.

Met Tower Correction

Due to sensor inaccuracies, various corrections were performed on the wind direction measurements. Following [14], the met tower was considered to have the most accurate wind direction sensor, but was adjusted for systematic bias by comparing the met tower sensor data with data gathered from two nearby turbines, as described next.

Given the measured wind direction at the met tower, denoted θ_m , the goal is to find an offset θ_f such that

$$\theta_t = \theta_m + \theta_f, \quad (1)$$

where θ_t is the true wind direction. To find θ_f , it is necessary to estimate the true wind direction θ_t . Fortunately, θ_t can be estimated using the geographical bearing (known exactly) of two turbines near the met tower, and comparing that bearing to the power variability of the downwind turbine. In theory, the power variability of the downwind turbine will peak when the wind direction is exactly aligned with the geographical bearing, providing an estimate of θ_t .

The power variability of a turbine is computed as

$$P_v = \sigma_p / \mu_p, \quad (2)$$

where μ_p is the power produced by the turbine, and σ_p is the standard deviation of the power produced, both taken over the ten minute intervals. For the purpose of comparing power variability against wind direction, the power variability is averaged over time for a given wind direction, where the wind directions are binned in 1° intervals.

For the wind farm under investigation, wind direction at the met tower was compared with the power variability of nearby turbine 64. Specifically, power variability was computed for turbine 64 when turbine 64 was downwind from turbine 67. Since turbine 64 is due north (bearing 180°) of turbine 67, the greatest power variability in turbine 64 should be seen in periods of southerly winds (bearing 180°). By comparing the power variability of turbine 64 with the wind direction at the met tower, the true wind direction θ_t can be estimated, as illustrated in Figure 2.

According to this analysis, the measured wind direction θ_m was offset from the true wind direction θ_t by an angle $\theta_f \approx 11^\circ$.

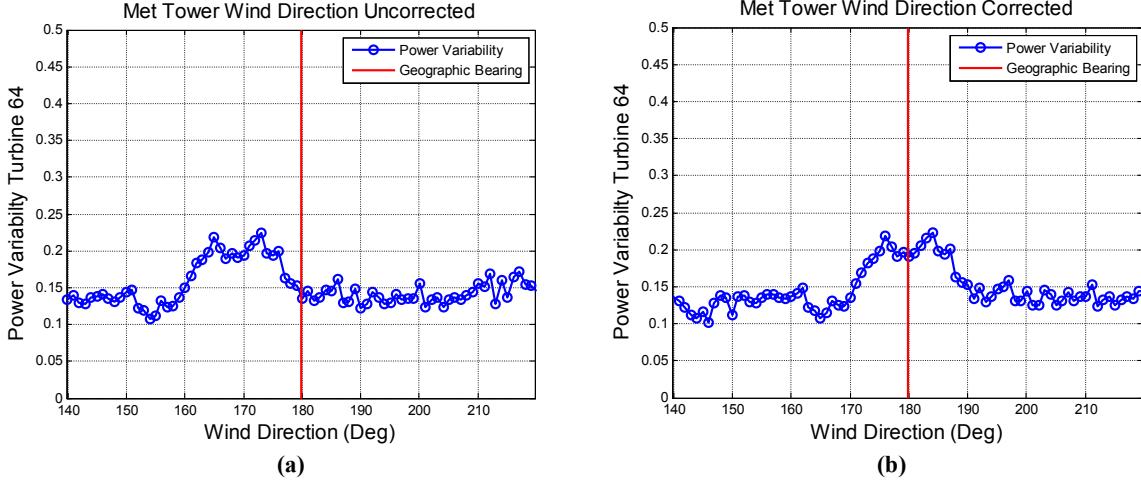


Figure 2. Met Tower Wind Direction Correction. On the left, power variability of turbine 64 is plotted against the met mast measured wind direction θ_m for winds from the south. A peak in power variability is seen at approximately 169° . In reality, this peak should occur at the geographical bearing from turbine 64 to turbine 67, which is shown as a red line, occurring at 180° . Therefore, the wind direction offset $\theta_f = \theta_t - \theta_m \approx 11^\circ$. On the right, power variability is plotted against the corrected wind direction, showing an alignment between the peak power variability and the geographic bearing. Throughout this figure, the power variability curves were computed across wind direction bins of 1° .

Nacelle Direction Correction

The nacelle direction sensors have a number of potential sources of error. First, they are not typically well maintained because the direction is not often used by the turbine controllers. Second, the typically slow changes in yaw and long periods of inaction cause the sensors to experience accuracy loss in relative position. These losses can be jumps in the position or slow deterioration. Although the direction may (in some turbines) be reset by the passage of a switch, some sites have very monotonous wind direction and the turbines rarely trigger the reset. Further, calibration to magnetic north is not always performed, or is sometime re-programmed with turbine software updates. This means that there can be instant jumps in directional determination at discrete times. Finally, if no service records are available, sensors may be replaced, after which absolute direction is lost.

The largest obstacle to overcome in correcting the nacelle wind direction is the identification of time period clusters showing significant relative change in wind direction against the corrected met tower data, now considered to be ground truth. An example of this behavior for turbine 6 is shown in Figure 3(a). Note that the direction measurements are present only when a turbine is operational and wind speeds exceed 4 m/s, as per the initial data screen.

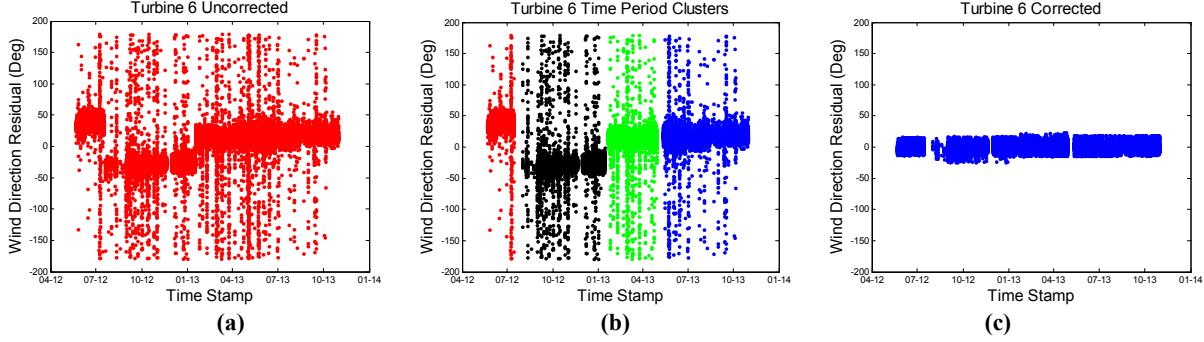


Figure 3. Nacelle Direction Correction for Turbine 6. On the left (a), the residual $\theta_n - \theta_m$ of the nacelle wind direction measurement θ_n versus the met tower measurement θ_m is shown through time for turbine 6. In addition to the additive biases that can be observed, there are several noticeable changes in the measurements at different time points. In the middle (b), the time periods are clustered using a change detection algorithm so that an additive correction can be applied. The separation of the last two groups (green and blue) is due to a period of erratic measurements which can be observed in the uncorrected data as a solid vertical line. On the right (c), additive corrections are applied to the time period clusters and individual measurements significantly different from the mean are removed.

Each time period cluster is identified using a change detection algorithm [15]. The change detection algorithm proceeds sequentially through each time series searching for time points where the moving average experiences a change beyond a given threshold. For the wind direction measurements, a moving average over 750 time points was used with a threshold of 20° . The results of the change detection algorithm applied to wind direction measurement data from turbine 6 is shown in Figure 3(b).

After the time period clusters were identified for a given nacelle, the wind direction based on the nacelle measurements were corrected using additive offsets, as was previously done for the met tower data. Using Equation 1, an offset θ_f is applied to the nacelle measurements for each time cluster such that the nacelle wind direction θ_n is equal on average to the met tower wind direction θ_m for that time period cluster.

Finally, individual wind direction measurements for each nacelle were discarded if greater than one standard deviation from the mean. An example of the corrected data for a nacelle is shown in Figure 3(c). The nacelle correction left approximately 34,000 time points per turbine for further analysis.

Pitch Correction

The final correction compared blade pitch versus wind speed to remove unusual instances of turbine operation. Normally, blade pitch should respond predictably to variations in wind speed, as dictated by the turbine controller. Unusual blade pitch response therefore indicates unusual turbine operation.

To perform this correction, instances of unusual turbine operation were removed if the mean blade pitch was more than one standard deviation from the mean blade pitch schedule, identified empirically by computing mean pitch vs. wind speed, using wind speed bins of 1 m/s. The blade pitch correction for turbine 6 is shown in Figure 4. Pitch correction left approximately 32,000 time points per turbine for further analysis, or approximately 222 days in operation.

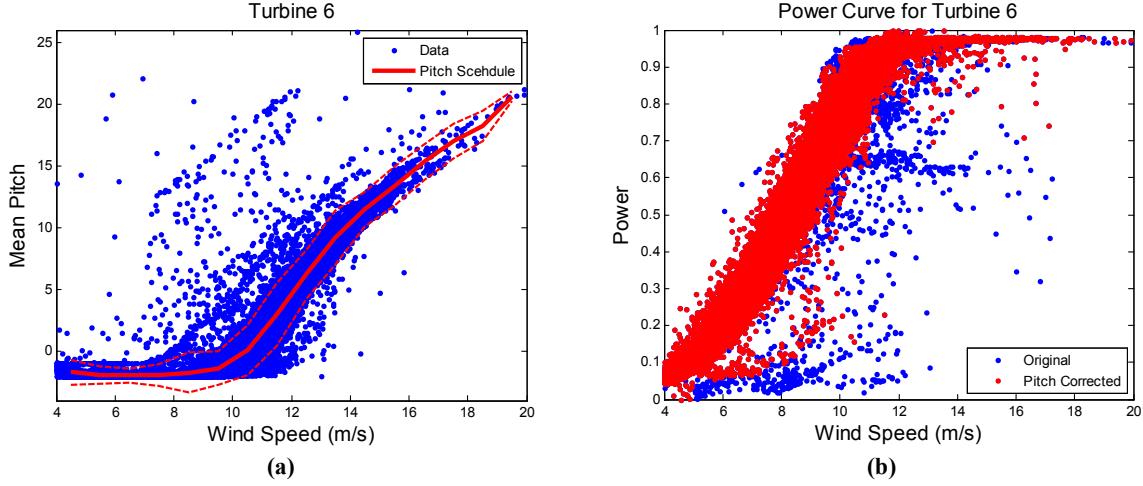


Figure 4. Pitch Correction. On the left, the pitch schedule is shown as the variation in mean blade pitch against wind speed. Data points more than one standard deviation from the pitch schedule (indicated by the dashed lines) are removed. On the right, the power curve of the corrected data is shown (power is shown on a normalized scale). Note that several abnormal operating modes were removed, including de-rated periods.

Power Correction

An alternative to pitch correction is power correction. In this correction, data outside the range of 60 percentile is removed from the median power curve (power vs. wind speed), where the median is computed using bins on the wind speed of 1 m/s. This correction removes power limited turbines as well as other abnormal modes of operation. An example is shown in Figure 5(a). The power correction left approximately 24,000 time points per turbine for analysis (note that we do not use pitch and power corrections together). The full collection of resulting median power curves shows a high degree of uniformity, as can be seen in Figure 5(b). This indicates that the power corrected data contains measurements of the wind farm in modes of normal operation, and is hence suited for use in assessing power loss due to wake effects.

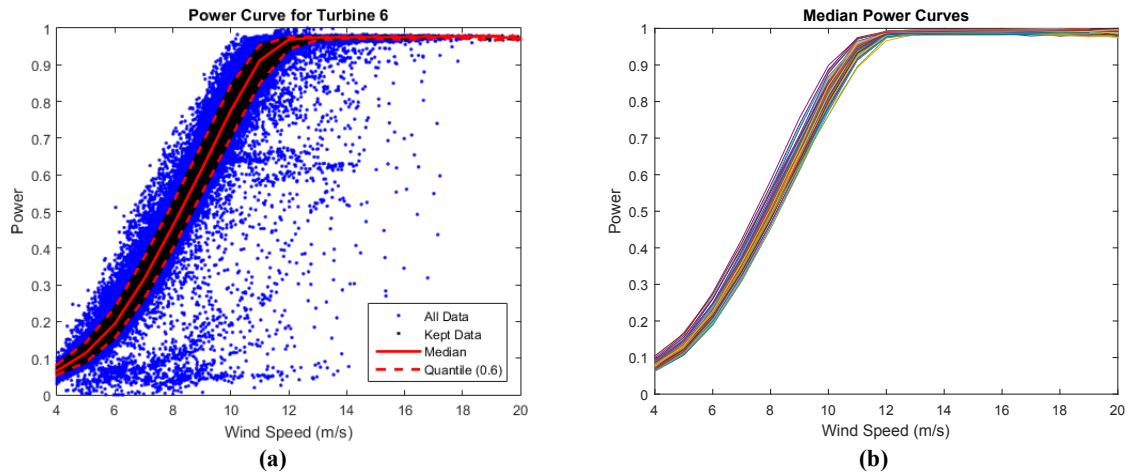


Figure 5. Power Corrections. On the left (a), the median power curve was computed using wind speed bins of 1 m/s, shown as a red line. Any data greater outside of the 60th percentile was then removed. This correction directly removes abnormal modes of turbine operation from consideration for further analysis. On the right (b), the full collection of median power curves (67) exhibit a large degree of uniformity over the wind farm, indicating that the power corrected data is suitable for estimating power losses due to wake effects on the farm.

2.2. Calculations

Turbine-Turbine Power Curve Pair Analysis

To motivate the method developed for visualizing wake effects across the entire wind farm, it is instructive to first consider the simpler case of two interacting turbines. For this exercise, consider the relatively isolated pair consisting of turbines 6 and 7, located near the northeast corner of the farm.

Turbines 6 and 7 are close enough (within 6 rotor diameters) that turbine 7 should experience wake effects given a westerly wind. However, wake effects were not observed by plotting the individual power curves for the two turbines, as shown in Figure 6(a). Fortunately, this observation is illusory, although the explanation is subtle.

It is important to realize that the wind speed observed by a nacelle is relative to that nacelle. In other words, a downwind nacelle will observe a wind speed slower than the true wind speed, precisely because it is in the wind shadow of the upwind turbine. Hence wake effects cannot be observed by using nacelle wind speed measurements. For the case of two turbines, this problem can be solved by using the upwind turbine wind speed measurement as the true wind speed, as shown in Figure 6(b).

Unfortunately, the solution used for turbines 6 and 7 will not scale to the entire wind farm, due to the numerous wind shadows and wind directions encountered over the full dataset. Imagine, for example, a row of turbines in line with the wind. The only turbine that is not in a wind shadow is the leading turbine. Thus the leading turbine will report the true wind speed, while every other turbine will report a wind speed slower than the true wind speed. The last turbine in the row will report the slowest wind speed. This situation is of course further complicated in an array of turbines, or as in the case of an actual farm, an irregular distribution of turbines.

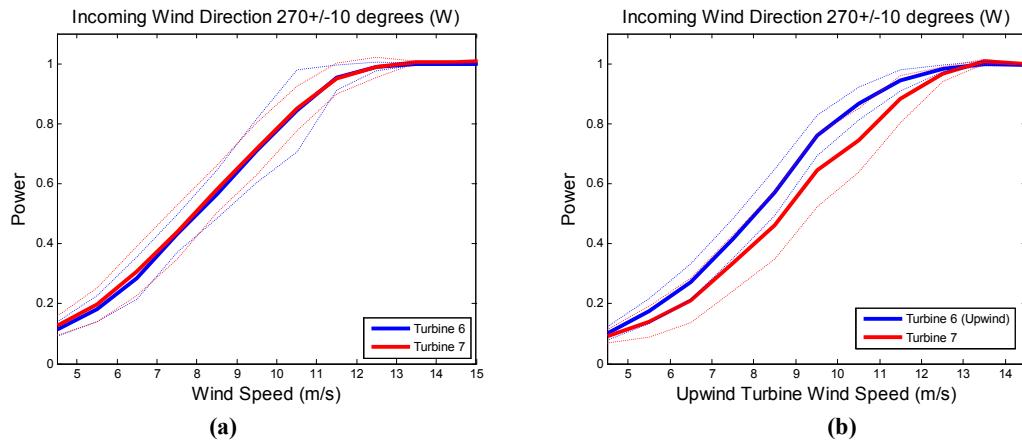


Figure 6. Wake Effect for Turbine 7. On the left (a), power curves are shown for the upwind turbine 6 and the downwind turbine 7 (given westerly winds). These curves show no wake effect because the individual nacelle wind speed measurements were used, which are relative. On the right (b), power curves are shown for the same two turbines, this time using the upwind nacelle sensor to measure wind speed. These curves show a wake effect, because the wind speed is now absolute for both turbines. Throughout this figure, the power curves were obtained by averaging the nacelle power measurements over bins with width of 1 m/s.

Directional Power Performance and Variance Plots

Since the nacelle wind speed measurements could not be used over the entire wind farm, calculations of relative power were substituted. In particular, a normalized measure of instantaneous power was employed. For turbine i , the *normalized instant power* is defined to be

$$P_N(t) = P_i(t)/\mu_{P(t)}, \quad (3)$$

where $P_i(t)$ is the power of turbine i over the ten-minute interval t , and $\mu_{P(t)} = \frac{1}{n} \sum_{i=1}^n P_i(t)$ is the average power over all turbines over the same interval (for our dataset $n = 67$). For example, if turbines 6 and 7 made up the entire wind farm, then the normalized instant power of turbine 6 would be $P_N(t) = 2P_6(t)/(P_6(t) + P_7(t))$. The normalized instant power avoids the use of wind speed measurements and can also be averaged over time and binned against wind direction to obtain polar plots showing the performance of a given turbine against the performance of the wind farm as a whole. For the normalized instant power plots, overlapping bins 1° apart covering $\pm 8^\circ$ sectors were used. The $\pm 8^\circ$ sectors correspond to anticipated wake effects for turbines separated by 7 rotor diameters. An example of normalized instant power plots, again assuming turbines 6 and 7 make up the entire wind farm, is shown in Figure 7.

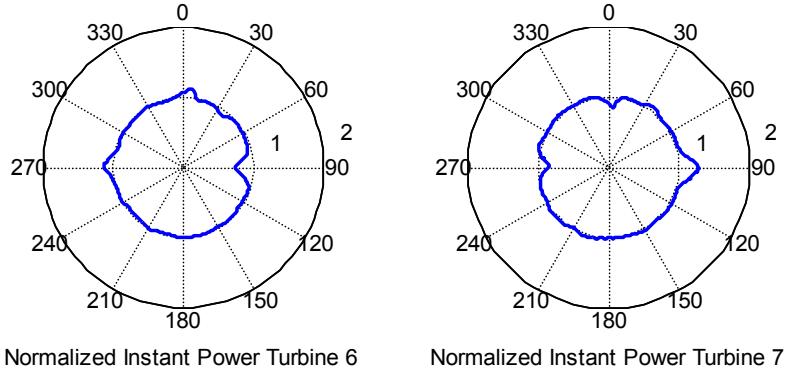


Figure 7. Normalized Instant Power Plots. These plots show normalized instant power averages over time versus wind direction for turbines 6 and 7. The plots are arranged to mirror the positions of the turbines in the wind farm, so that turbine 6 is west of turbine 7. Thus, given a westerly wind, turbine 6 is upwind and turbine 7 is downwind. In this case, turbine 6 is over-performing relative to turbine 7, so that a bump is observed in the plot for turbine 6 at the angle 270° . Similarly, turbine 7 is underperforming due to the wake effect, and a dip is observed in the plot for turbine 7 at the angle 270° . Both plots are obtained by averaging normalized instant power over time within wind direction bins of 1° .

In Figure 7, the mutual wakening of the two turbines at 90° and 270° can be clearly observed. For turbine 6, there is a hint of wake deficit from turbine 1 at 30° and a stronger deficit bearing 0° for turbine 7. The difference in the deficit is smaller at turbine 6 due to the larger distance to turbine 1. Turbine 7 also shows an apparent increase at 90° . This may be an artifact of the normalization, meaning that the farm is on average highly shadowed compared to turbine 7 in free wind speed at this bearing. However, it may also be an effect due to wake effects from another wind farm upstream of turbine 7 (to be discussed further in the Results section). Also note that these plots are slightly different from the plots generated for the entire wind farm (Figure ?), since this example is normalized against only turbines 6 and 7 (i.e. this example assumes the wind farm

consists only of turbines 6 and 7, even though the other turbines in the farm clearly affect the turbine 6 and 7 subset.)

A variation on the normalized instant power computation is power variability, as defined in Equation 2. Power variability averages over time can be computed and binned against wind direction (again using overlapping bins 1° apart covering $\pm 8^\circ$ sectors) as shown in Figure 8.

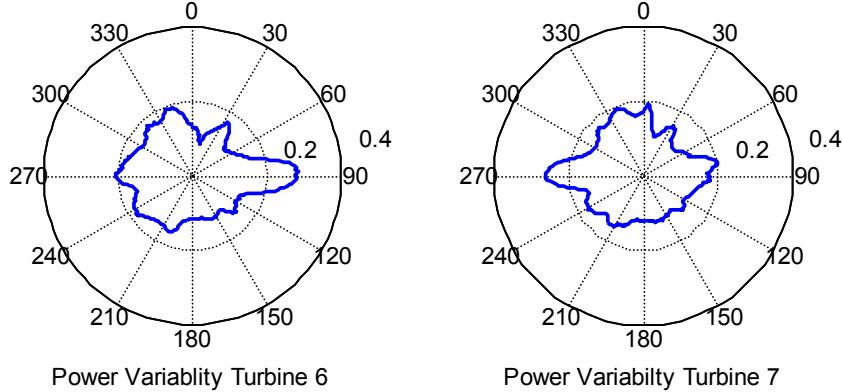


Figure 8. Power Variability Plots. These plots show power variability over time against wind direction for turbines 6 and 7. They are again arranged to mirror the positions of the turbines in the wind farm. Wake effects can be observed as increased variability, which manifest as bumps on the plots. Thus there are bumps at 90° (easterly wind) for turbine 6 and 270° (westerly wind) for turbine 7. Both plots are obtained by averaging power variability over time within wind direction bins of 1° .

Wind Shadow Calculation

To estimate power lost due to wake effects, an approximation of the wind speed in the absence of upwind turbines needs to be computed. This can be done if the upwind turbines throughout the farm are identified and used to extrapolate the true wind speed for the shadowed turbines. In other words, the first step is to determine the front of the wind farm in the case of different wind directions.

Suppose, for example, that a wind farm consists of only two nearby turbines, arranged in the east-west direction. If the wind blows from the east then the western turbine is in the shadow of the eastern turbine, so that the eastern turbine is at the front of the farm, and the wind speed measured by the eastern turbine is the true wind speed. The situation is reversed when the wind blows from the west, and in the case of a northerly wind, neither turbine is in the shadow of the other, and hence they are both at the front of the farm.

To identify the front of the farm, a simplified model of turbine wakes is used, known as the Park model [4, 18]. This model is also used in an effort to optimize wind farm performance using game theory [18]. In the Park model, the diameter of a wake increases linearly as a function of the distance behind the upwind turbine

$$D_w(x) = D_t + 2kx, \quad (4)$$

where $D_w(x)$ is the diameter of the wake as a function of the distance x behind the upwind turbine, D_t is the diameter of the turbine, and k is a roughness coefficient which depends on the

geography of the wind farm. In general, for farm land $k = 0.075$, and for offshore wind plants $k = 0.04$. Mathematically, k determines how rapidly the wake spreads behind the upwind turbine.

Using the Park model and a distance x_∞ beyond which we consider that the wake has dissipated, we can estimate which turbines are in wind shadows in a wind farm, depending on the wind direction. By performing this calculation for different values of k and x_∞ we can obtain different estimates of the wind shadows within the farm. Examples of this calculation for different wind directions and different values of k and x_∞ are shown in Figure 9.

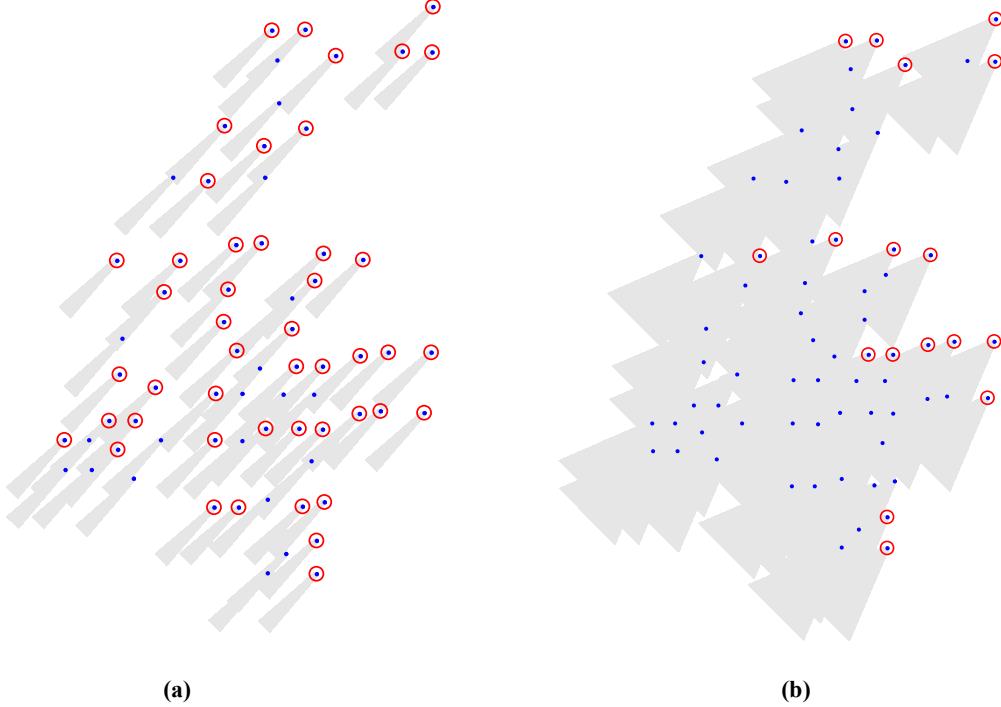


Figure 9. Wind Shadow Estimates. On the left (a), wind shadows are computed for wind blowing out of the north-east using $k = 0.075$ and $x_\infty = 15$ (rotor diameters). Turbines are marked with blue dots and wind shadows are shown using gray quadrilaterals. Turbines not in wind shadows are highlighted with red circles. On the right (b), wind shadows are computed using the same wind direction but with an extremely high $k = 0.4$ and $x_\infty = 20$. The wind front computed in the two examples shows how more or less conservative wind front estimates can be obtained by varying k and x_∞ .

Power Observer Calculation

In order to estimate power lost due to wake effects, it is necessary to estimate a wind observer for each turbine. The wind observer is defined as the free uninterrupted wind speed, as if the wind farm was not present. If a turbine is in an undisturbed state (not waked), the wind observer is just the wind speed measurement at the nacelle, assuming the nacelle anemometer has been adequately calibrated. If a turbine is in a wind shadow (waked), the wind observer should be the free wind speed, or the wind speed the turbine would experience in the absence of any upwind turbine. Fortunately, the wind observer for waked turbines can be easily estimated by using the wind speed measurement at the nearest non-waked turbine. The nearest non-waked turbine is determined using the previous wind shadow calculation.

Normally, the power curves would be used to estimate a power observer based on the wind observer, thus giving an estimate for power lost due to wake effect. However, a power observer can be calculated directly using exactly the same technique that was used to estimate the wind observer. Namely, the power observer for an un-waked turbine is provided by the power measurement at the nacelle, while the power observer for a waked turbine is estimated by using the nearest non-waked turbine power measurement. An example of the power observer calculation compared with the measured power against wind direction is shown in Figure 10. Using the power observer instead of the wind observer is more robust. No second order calibration is required (e.g. using the wind observer along with the power curve). Further, nacelle anemometers are generally less accurate than power sensors.

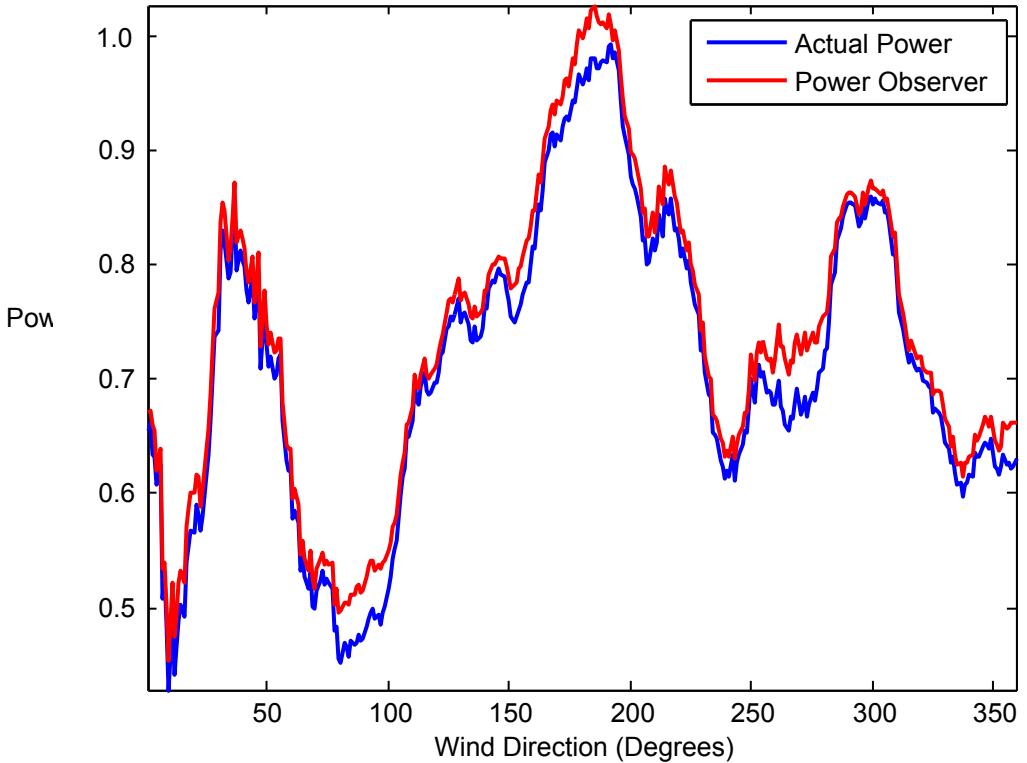


Figure 10. Power Observer versus Actual Power. Shown here is the power observer calculation (red) compared to the actual power measurements (blue). The curves are totaled over all turbines in the farm. For this estimate, $k = 0.075$ and $x_\infty = 15$ (rotor diameters).

Finally, the power observer calculations are used to estimate the total power lost over the wind farm due to wake effects. This loss is given as the potential percent improvement

$$G = \frac{\sum_i P_{obs,i} - P_{act,i}}{\sum_i P_{act,i}} \times 100, \quad (5)$$

where G is the potential percentage improvement in power (gain), $P_{obs,i}$ is the power observer for turbine i , and $P_{act,i}$ is the actual power measured for turbine i . The potential percentage gain for the calculation shown in Figure 10 is 3.0%.

Finally, we note that for this particular wind farm, the difference between the average observed wind power and the wind power observer (an estimate of the free wind power) is very small. Therefore, in our analysis, there is almost no difference in using normalized power from Equation (5), versus the wind power observer. In addition the observer is not calibrated for complex landscapes.

2.3. Visualization

Colored Polar Plots

The normalized instant power and power variability polar plots can be used to visualize wake effects across an entire wind farm. To facilitate viewing multiple plots simultaneously, a color scale can be added to the plots. For the instant power visualization these colors highlight over/under performing turbines, and for the power variability visualization the colors highlight high/low power variance per turbine. The color scales are computed to be comparable across the entire farm (i.e. the same scale is used for every plot in the entire visualization). Visualizations of the wind farm are given in Figure 11. Wake effects are easily seen using these images.

Note that the average power in the normalized power computation is taken over periods where most of the farm is operational, but not necessary the entire farm. Specifically, average power is computed when at least 62 turbines are operational, or 5 turbines are not operational. This measure avoids ignoring periods when the farm is operational except for a few down turbines (or turbines with previously discarded measurements).

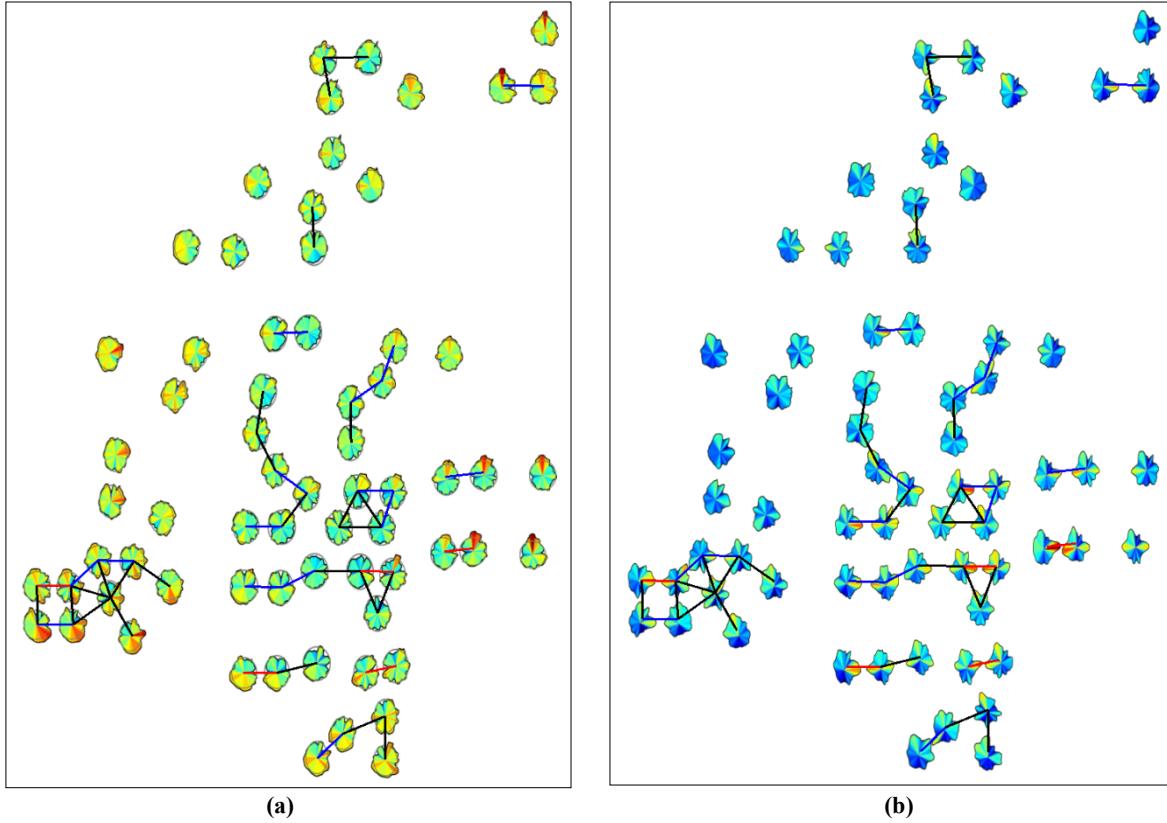


Figure 11. Wake Effect Visualizations. On the left (a), visualization is shown using normalized instant power, and on the right (b) visualization is shown using power variability. In both cases, the rose plots are positioned in place of the turbine icons seen in Figure 1. Further, the rose plots are colored according the radial magnitude. For the instant power plots, over-performing turbine directions are colored red, while under-performing turbine directions are colored blue. Grey circles show average performance (instant normal power value of 1). For the power variability plots, high variability directions are colored red, and low variability directions are colored blue.

The visualization can also be displayed using Google Earth (<http://www.google.com/earth>) for improved interactivity. An example of a fully interactive visualization is shown in Figure 12.

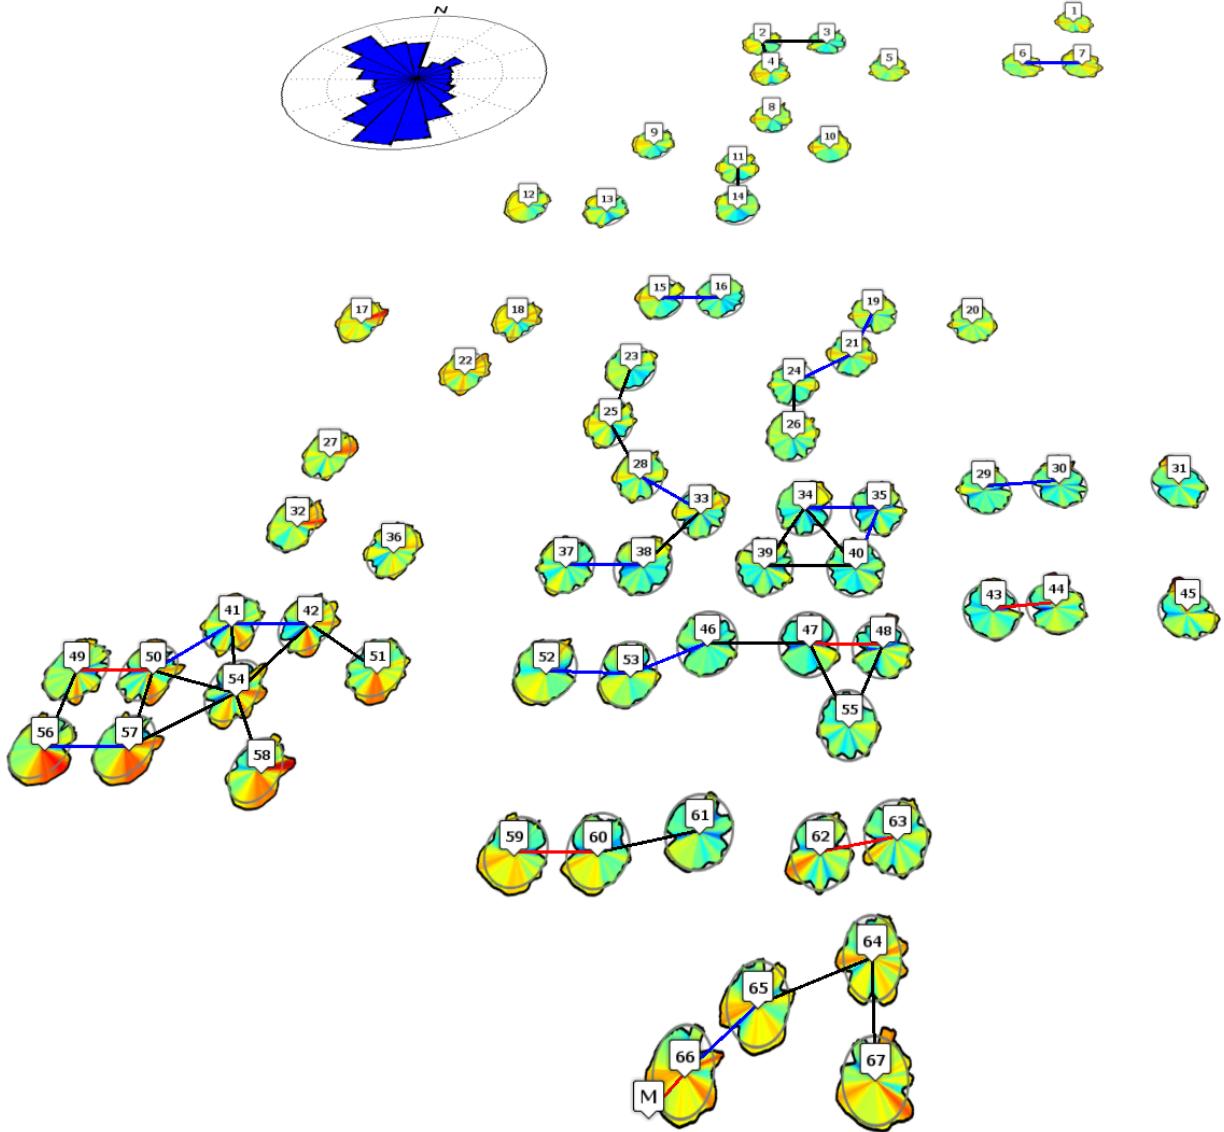


Figure 12. Google Earth Visualization. This visualization shows the wind farm wake effects using the normalized instant power plots, complete with the labels from Figure 1. Grey circles show average performance (instant normal power value of 1). The wind farm terrain imagery (not shown) can also be examined for correlations between performance and local topography.

2.4. Model

Fitted Parameters

Based on the data visualization described above, a wake model has been derived from an analysis using selected wake profiles. From 316 turbines with direct wakes, we identified 230 within a 12 rotor diameter (12D) range. Out of these, a total of 99 wake profiles were chosen based upon the upstream turbine being relatively undisturbed and un-waked. The wake profiles were sorted in 1.25D bins ($\pm 2/3D$) and averaged based on bearing to the upstream turbine, with about 16 wake profiles in each bin. A small subjective correction of the direction was applied when the individual profile appeared to be slightly off centerline. This correction was motivated

by the uncertainty of the directional data, and the fact that the upstream turbine inflow profiles appear to have a large effect on the downstream profiles.

The peak deficit for all of the profiles agree with the Park model [4], shown in Figure 13 plotted as an average of the nearest five points to the bearing against the Park model to the 3rd power, with $k = 0.075$ and $C_T = 0.7$. The Park model (sometimes known as Jensen's Park model) is given by

$$\frac{U_x}{U_0} = 1 - \frac{1 - \sqrt{1 - C_T}}{1 + 2k \frac{x}{D}}, \quad (6)$$

where U_0 is the upstream wind velocity, U_x is the wind downstream velocity at a distance x behind the turbine, C_T is the coefficient of thrust, and k is an empirical decay constant, given as 0.075 for flat land. Note that other combinations of k and C_T could have been used had the full details of the rotor been available.

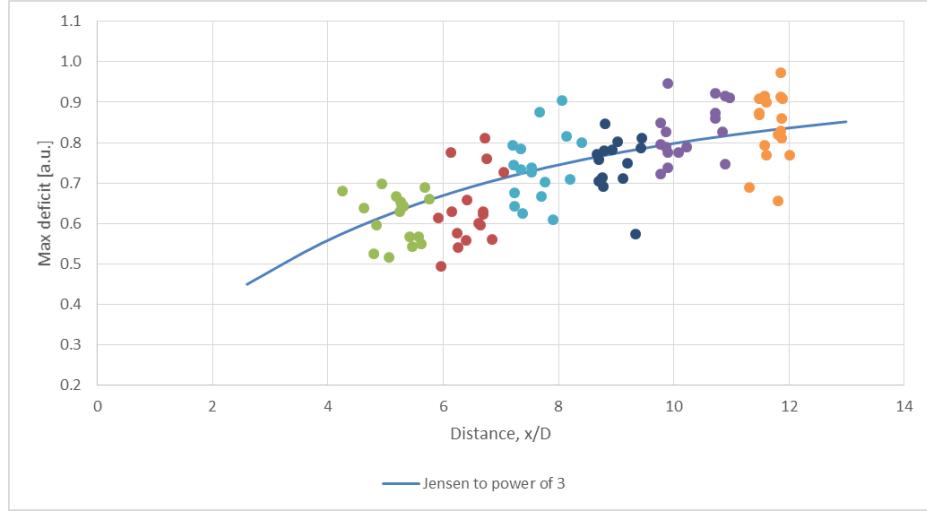


Figure 13. Peak Deficits. The peak deficit for all of the profiles plotted as an average of the nearest five points to the bearing against Jensen's Park model to the 3rd power, with $k = 0.075$ and $C_T = 0.7$. The profiles from each bin are colored coded.

When observing Figure 13, it is important to remember that each point represents a long period of measurement. However, even though these examples were selected to have relatively free upstream inflow, there are still obstacles from outside the wind farm.

The average normalized power wake profiles for each bin are shown in Figure 14(a). The standard deviations of the profiles are approximately the same across all ranges and directions and are on the order of 0.07. The possible range of the analysis is limited by this standard deviation to about 11D. The power variance is shown in Figure 14(b).

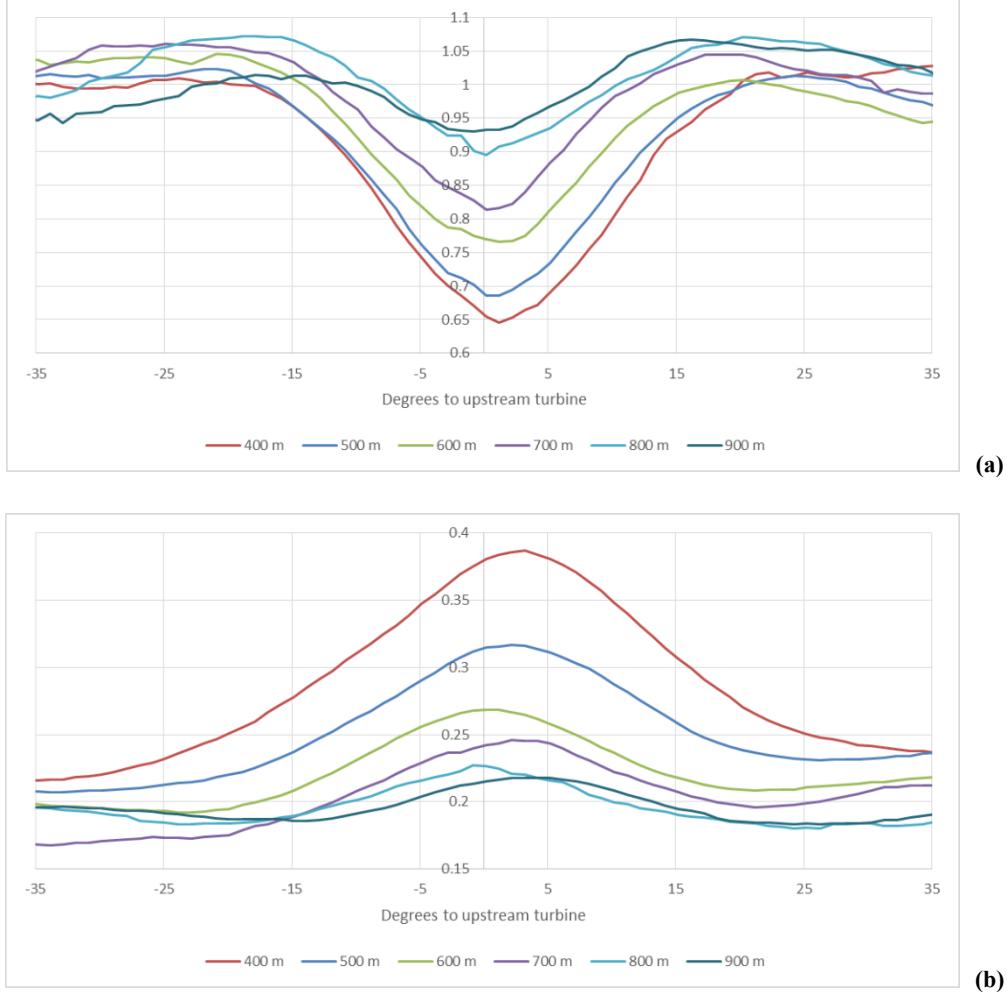


Figure 14. Wake Profiles from 3D to 12D. On the top (a), the average normalized power wake profiles for the 1.25D bins are shown. There are approximately 16 profiles in each bin. On the bottom (b), power variance of the wake profiles for the 1.25D bins is shown.

The average profiles shown in Figure 14 were fitted with a Gaussian bell with an offset to accommodate for the non-zero value outside the wake and a directional offset and summarized in Figure 15. These fitted curves are used in the model to follow.

The wake centerline is clearly offset within a few degrees, corresponding to approximately $\frac{x}{D} = 5$. The offset seems to diminish with distance, although detecting the offset at larger distances is difficult. Possible reasons for a wake failing to propagate from directly from the upstream to downstream turbine include: (a) all of the turbines investigated operate systematically with a small yaw error, (b) the Coriolis force produces a minor curving of the wakes, and (c) there is a systematic wind shear.

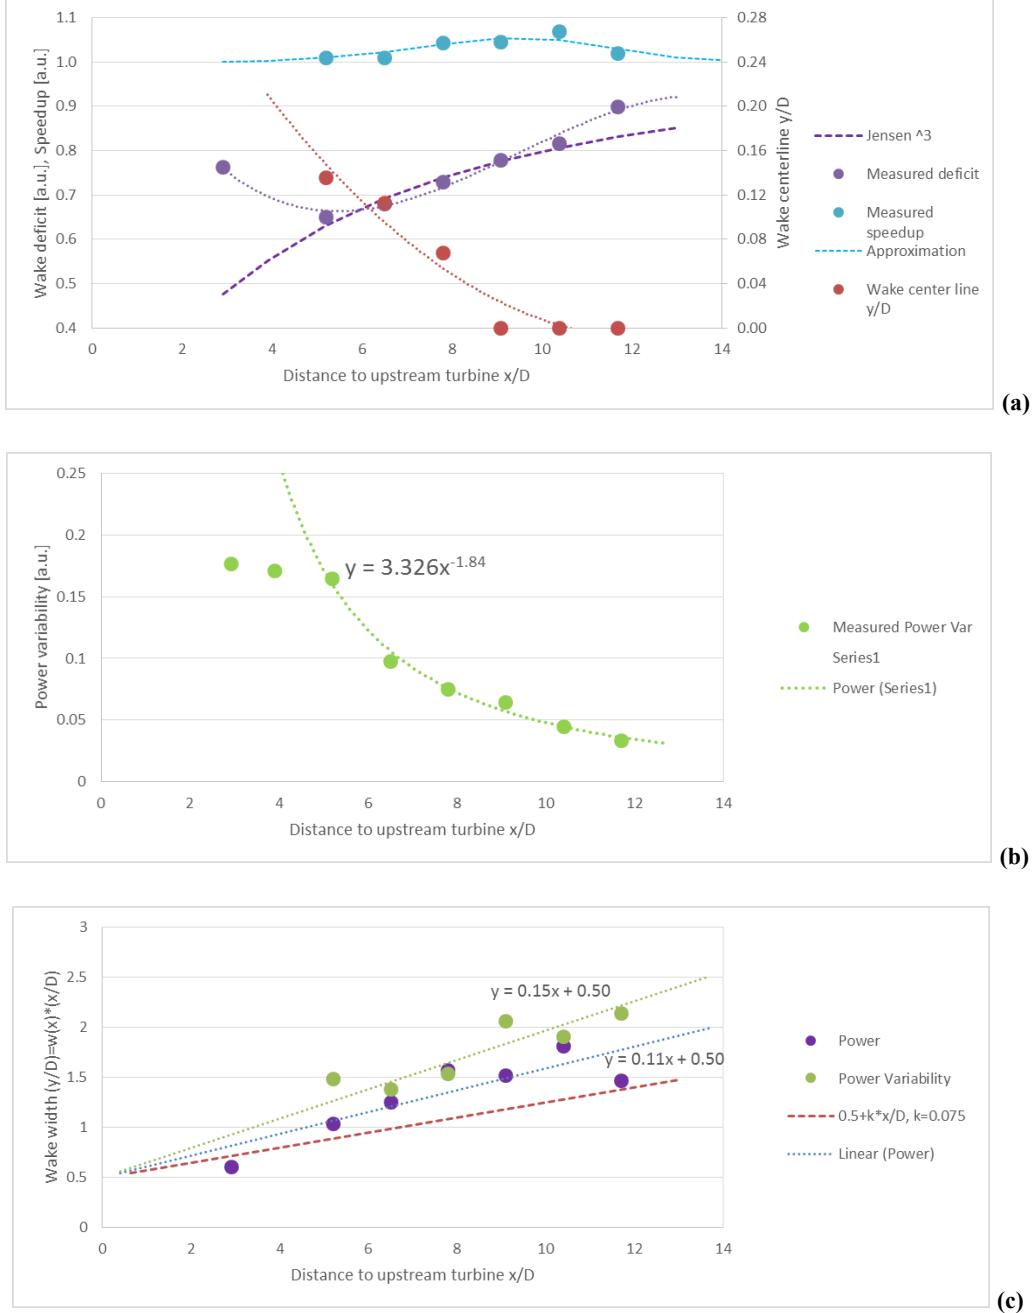


Figure 15. Fitted Parameters. On the top (a), we show wake deficit and centerline for the normalized power wake profiles. The fitted values outside the wake are given by $a(x/D) = 1 + 25/D \times \text{Weibull}(x/D, 10)$. In the middle (b), we show the measured peak power variability above the ambient level, with a fitted value. On the bottom (c), we show curves fitting wake width in linear coordinates for power and power variability. For reference, the Park model wake width is plotted for $k = 0.075$. The power performance fit is $w(x) = 0.11 \times (x/D) + 0.5$ and the variability fit is $b(x) = 0.15 \times (x/D) + 0.5$.

The average centerline deficit fits well with the Park model to the 3rd power using the assumed values of k and C_T , as shown in Figure 15(a). The power variability in the far wake is approximately proportional to $((x/D)^{0.5})^{-3}$, but flattens out for $x/D < 5$, as shown in Figure 15(c). The width of the wake in Figure 15(c) has been plotted in linear coordinates using $y/D = w \times (x/D)$ rather than in the observed coordinates to better compare with the Park model.

As can be seen, there is reasonable agreement, especially considering our arbitrary choice of $k = 0.075$. If the deficit curve in Figure 15(c) was computed using $k = 0.11$ (often suggested by experimental data), a C_T value of approximately 0.85 would provide a good fit.

Wake Model

For an individual turbine facing an upstream turbine, we have observed the centerline deficit be proportional to the Park model adjusted for upstream deficit. If we assume the wake is described by a Gaussian bell in polar coordinates, we can express the wake power deficit p_j at any given bearing by the product of all turbines in the wind farm:

$$p_j(\theta) = \prod_{i=0}^N p_i(\theta_{i,j}) \left(1 - (1 - d(C_T k_i \frac{x_{i,j}}{D})^3) e^{-\left(\frac{\theta - \theta_{i,j}}{w x_{i,j}}\right)^2} \right), \quad (7)$$

where $i \neq j$, $d(C_T k_i x_{i,j}/D)$ is the Park model (Equation 6) deficit of the upstream turbine with the width $w_{i,j}$ in the direction $\theta_{i,j}$ and distance $x_{i,j}$, the deficit of the upstream turbine is given by $p_i(\theta_{i,j})$, and the wake width is obtained from the curve fit shown in Figure 15(c).

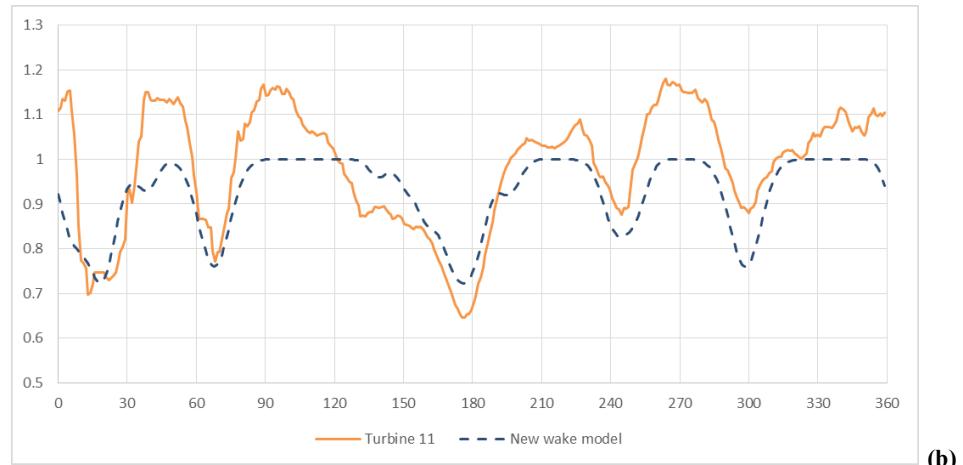
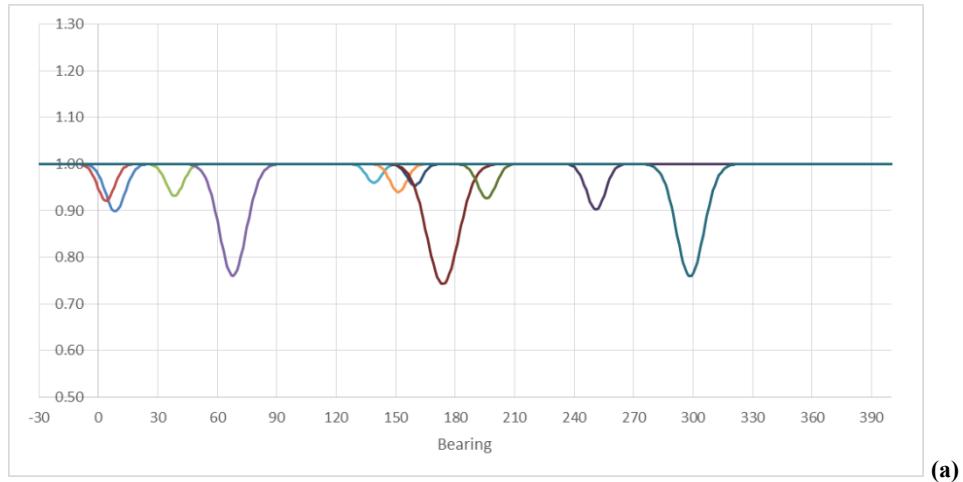


Figure 16. Initial Model. On the top (a), individual wake components for turbine 11 are modelled as Gaussian bells. On the bottom (b), all wakes for turbine 11 are modelled as the product of the individual Gaussian bells. Three of the wake directions have been adjusted a few degrees to align with the experimental data.

The individual components of the Gaussian bells are shown for turbine 11 in Figure 16(a) and the joint product in Figure 16(b) is compared to the experimentally observed values. The deficits align with the wakes, but it is clear the speedup effects are neglected. The model is therefore expanded with a cosine overlay:

$$p_j(\theta) = \prod_{i=0}^N p_i(\theta_{i,j}) \left(a(x_{i,j}) e^{-\left(\frac{\theta - \theta_{i,j}}{\alpha_{i,j}}\right)^2} \left(1 - \cos\left(\frac{\theta - \theta_{i,j}}{\beta_{i,j}}\right) \right) + \left(1 - (1 - d(C_T k_i x_{i,j})^3) e^{-\left(\frac{\theta - \theta_{i,j}}{wx_{i,j}}\right)^2} \right) \right), \quad (8)$$

where $i \neq j$, and we assume that $\beta = 1.5w$, $\alpha = 2\beta$ and $a(x/D) = 1 + 25/D \times \text{Weibull}(x/D, 10)$, using the fitted curve in Figure 15(a). With the cosine variation of the initial model, the results improve, as seen in Figure 17.

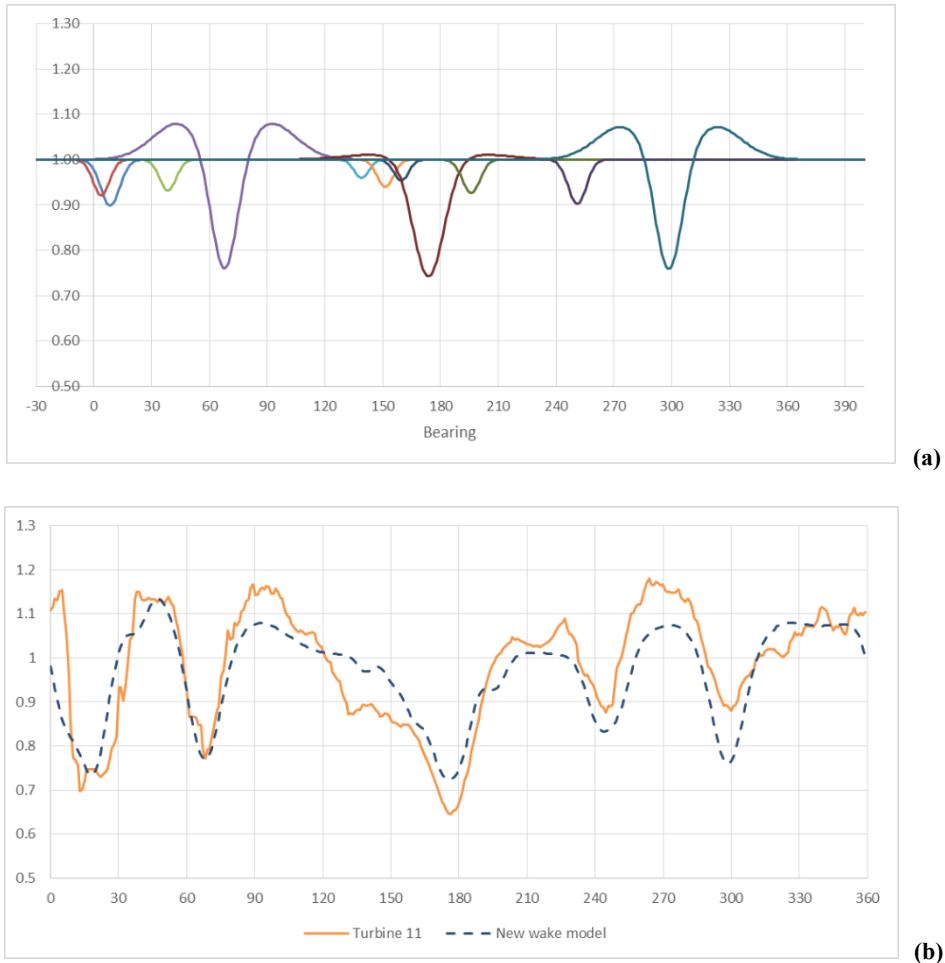


Figure 17. Cosine Model. On the top (a), individual wake components for turbine 11 are modelled using a Gaussian bell and a damped cosine overlay. On the bottom (b), all wakes for turbine 11 are modelled as the product of the individual Gaussian bells using the damped cosine overlay.

The individual components of the Gaussian bells are shown for turbine 11 in Figure 17(a) and the joint product in Figure 17 is compared to the experimentally observed values. As can be seen the agreement is much higher with these preliminary choices of parameters.

Finally, the wake power variability is modelled by:

$$pvar_j(\theta) = v_0 + \sum_{i=0}^N v(x_{i,j}) \cdot e^{-\left(\frac{\theta - \theta_{i,j}}{b(x_{i,j})}\right)^2}, \quad (9)$$

where $i \neq j$, v_0 is the background power variability (~ 0.15) and $v(x_{i,j})$, $b(x_{i,j})$ are the amplitude and width found from the fitted curve shown in Figure 15(b). The results are illustrated in Figure 18.



Figure 18. Power Variability Model. On the top (a), individual wake components for turbine 11 are modelled using Gaussian bells. On the bottom (b), all wakes for turbine 11 are modelled as the sum of the individual Gaussian bells.

3. RESULTS

3.1. Visualization

In this section, a more detailed analysis of the wind farm is provided, using the normalized instant power and power variability plots, as well as other simpler curves. To start, it is interesting to understand the overall performance of the farm in relation to wind direction. Although it is not useful to take the average of the normalized instant power plots (since they are normalized they average to 1), it is useful to consider the average turbine power relative to wind direction, as shown in Figure 19(a). As expected, this plot is highly correlated with the wind rose for the site, albeit with an unusual spike in production from the NE. The power variability plots, unlike the instant normalized power plots, can be averaged, since they are computed independently per turbine. The power variability averages over the wind farm is shown in Figure 19(b). There are three large lobes of variability at 90° , 270° , and 340° . The lobes at 90° and 270° are most likely associated with turbulence generated by closely positioned E-W turbine pairs, as there are a significant number of turbines in the $90^\circ - 270^\circ$ neighboring direction. This can be seen comparing the turbine to turbine bearing, as shown in Figure 19(c), to the variance distribution. The high variance at 340° is aligned with the main wind direction, and a large nearby wind farm to the NE. There is less variation along the other main wind direction to the SSE, but the neighboring wind farms in this direction are further away and shadow only the east part of the farm.

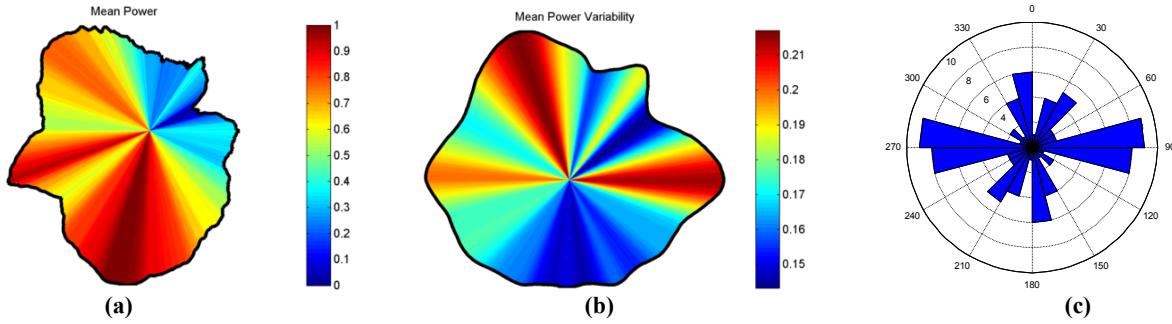


Figure 19. Overall Wind Farm Performance. On the left (a), the power averaged over all turbines versus wind direction is shown. As expected, this polar plot is highly correlated with prevailing wind directions (NW and S), although there is also an interesting spike in the NE direction. In the middle (b), the average power variability versus wind direction is shown. The power variability is aligned with prevailing winds as well as a large number of closely positioned E-W turbine pairs. On the right (c), a radial histogram of counts for turbine-turbine pairs within 7 rotor diameters is shown.

In Figure 11 it is apparent that individual turbine performance and variation is related to position within the wind farm, where turbines towards the center of the farm have lower performance and higher variability. This is more easily visualized by averaging the normalized instant power and power variability plots over all wind directions, as shown in Figure 20.

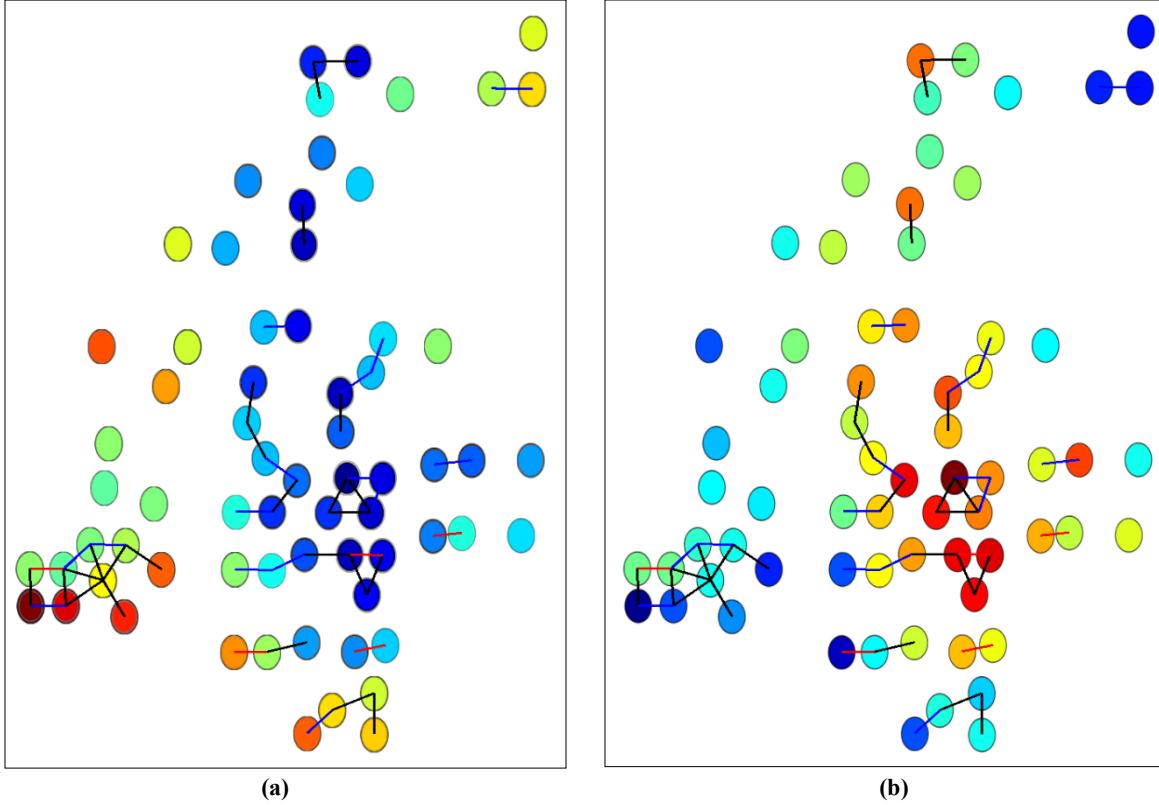


Figure 20. Average Turbine Performance over Wind Farm. On the left, the wind direction averaged normalized instant power is shown for each turbine, where red indicates high performance, and blue indicates low performance. On the right, the wind direction averaged power variability is shown for each turbine, where red indicates high variability, and blue indicates low variability.

Detailed Wake Analysis

In addition to providing an overview of the wind farm, the normalized power and power variability plots can be used to understand wake effects on a more detailed scale. Four distinct features have been observed in the data:

- wake deficit effects,
- speed up channels from two upstream turbines,
- shear point speedup from one upstream turbine, and
- shear point speedup from multiple upstream turbines or an upstream farm.

Wakes are characterized by a power deficit in the direction of a neighboring turbine and a distinct increase in power variance. In Figure 21, two distinct wake effects can be seen at turbine 11, originating from turbines 8 and 10. The first effect is the wake deficits, seen as dips in the normalized power plot of turbine 11 facing turbines 8 and 10, as well as peaks in the power variability.

The second observed effect is a speed up when a turbine is facing the midpoint of two upwind turbines. Here the power is higher than nominal, also seen in Figure 21 when turbine 11 is facing the midpoint between turbines 8 and 10. This effect is generally not included in the

standard wind farm modeling tools, even though it is apparently significant. Somewhat surprising, the power variance in this direction is significantly lower than average.

A third effect, which is a potentially new discovery, occurs when the upstream wake is barely touching the downstream turbine, causing an increase in power. This occurs in Figure 21, for example, when increased power is observed at turbine 11 just south of the downstream sector of turbine 10. Surprisingly, this is also associated with a low variance. We call this effect *shear point speedup*.

Both types of speedup effects seem to be amplified in cases of multiple upstream turbines and even a whole upstream wind farm. Finally, the three effects are also observed as products of the landscape, i.e. clusters of trees or buildings (data not shown).

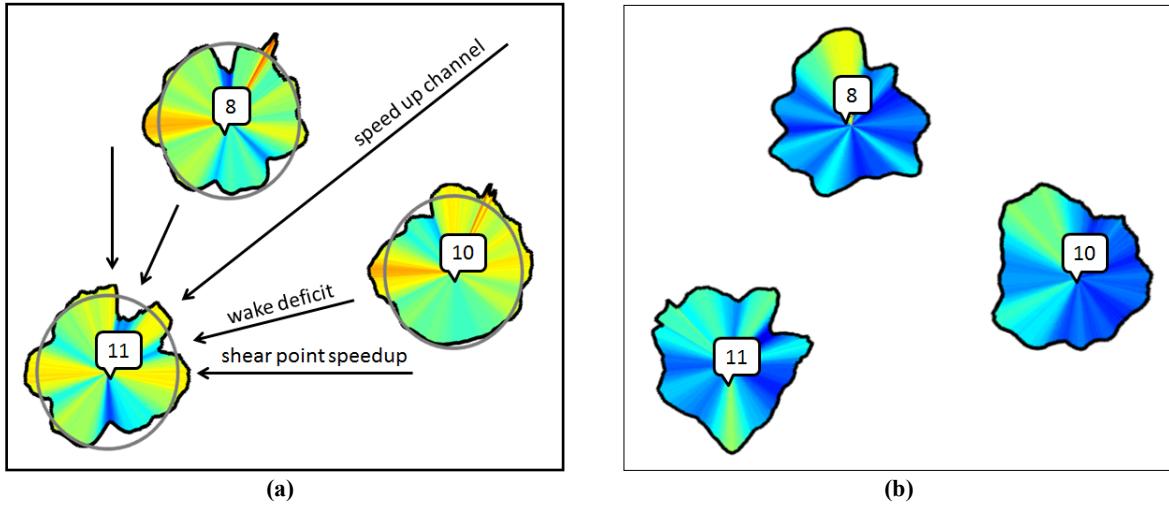


Figure 21. Wake Effects. On the left (a), three wake effects can be observed using instant normalized power plots for turbines 8, 10, and 11. Wake deficits can be seen as dips in the power production when turbine 11 is in the shadow of turbines 8 or 10; a speed up channel can be seen as a peak in the power production when turbine 11 is facing the midpoint of turbines 8 and 10; and shear point speedups can be seen when turbine 11 is tangent to the wake of turbine 10 or 8. On the right (b), the corresponding variability in power is shown.

Wake Deficit Profiles

To further illustrate the four wake effects observed using the instant power and power variability plots, profiles for various turbines are shown. The wake deficit, for example, is straightforward to detect when the wake is produced by a single upstream turbine which is relatively undisturbed. In this case, the power variance has a very clear peak in the direction of the upstream turbine, and the power produced has a corresponding minimum. Wake profiles showing wake deficits for 10 turbines pairs in terms of normalized instant power are shown in Figure 22, with the corresponding power variance profiles shown in Figure 23.

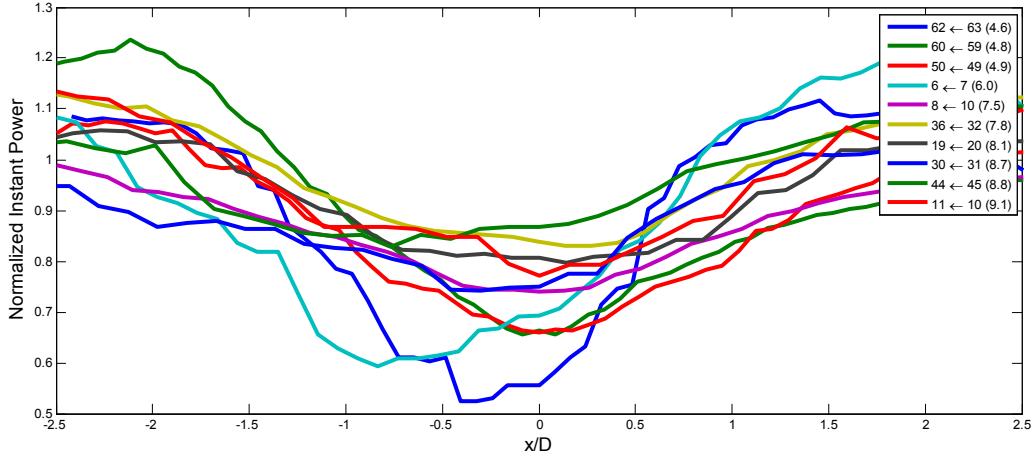


Figure 22. Normalized Instant Power for Waked Turbines. The x -axis gives the wind direction degree offset normalized by the $x/D = \theta^\circ \left(\frac{\pi}{180^\circ} \right) \left(\frac{x}{D} \right)$, where θ° is the offset in degrees, x is the distance between the two turbines, and D is the rotor diameter. For example, when the offset is 0° the wind is blowing straight from the upwind to the downwind turbine. The y -axis shows the normalized instant power for the downwind turbine. The turbine pairs selected are given in the legend, where the notation $d \leftarrow u (r)$ indicates the downwind turbine (d), the upwind turbine (u) and the distance (r) in rotor diameters. See also Figure 1 for turbine positions.

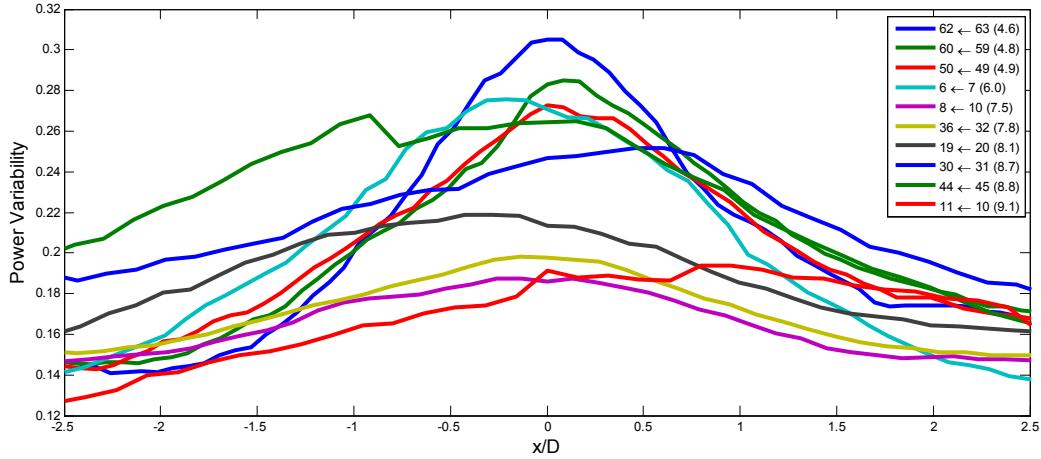


Figure 23. Power Variability for Waked Turbines. As in Figure 22, the x -axis gives the wind direction degree offset from the upwind turbine. The y -axis gives the power variability of the downwind turbine.

The detection of a wake deficit is straightforward when the wake is produced by a single undisturbed upstream turbine. In particular, the power variance signature is a very clear peak in the direction of the upstream turbine. However, in a multi-turbine wake situation, detecting wake effects is more difficult. For example, turbine 62 waked by turbine 63 shows a clear wake deficit and an increased variance, but has two distinct side lobes where the power increases over the nominal value outside the core of the wake. This is most likely due to wakes upstream of turbine 63 in an adjacent upstream wind farm. It is also interesting to observe that the peak variance is exactly pointing towards turbine 63, but the wake deficit is a few degrees off to the

left. Turbine 10 waked by turbine 11 exhibits the opposite behavior, where the wake is symmetric, but the variance profile is asymmetric. In general, however, the variance profiles are symmetric.

In addition to individual wake deficit profiles, a generic wake deficit effect can be observed across the entire wind farm. To see this effect, 854 turbine pairs were selected within 25 rotor diameters and an undistributed direct path between them (to observe potential wake effects). From these pairs, the minimum power deficit and maximum variance of the downwind turbine were recorded. These values are shown in Figure 24 versus distance between turbine pairs. For the normalized instant power in Figure 24(a), the semi-empirical expression for wind velocity deficit described by the Park model [4] is also shown (Equation 6). Reasoning that annual energy production is typically quasi-linear with wind speed averaged over time, the wind velocity deficit is compared with the normalized instant power in Figure 24(a) to surprisingly good effect. For the wind velocity deficit, k was taken to be 0.075 and C_T was set to 1.

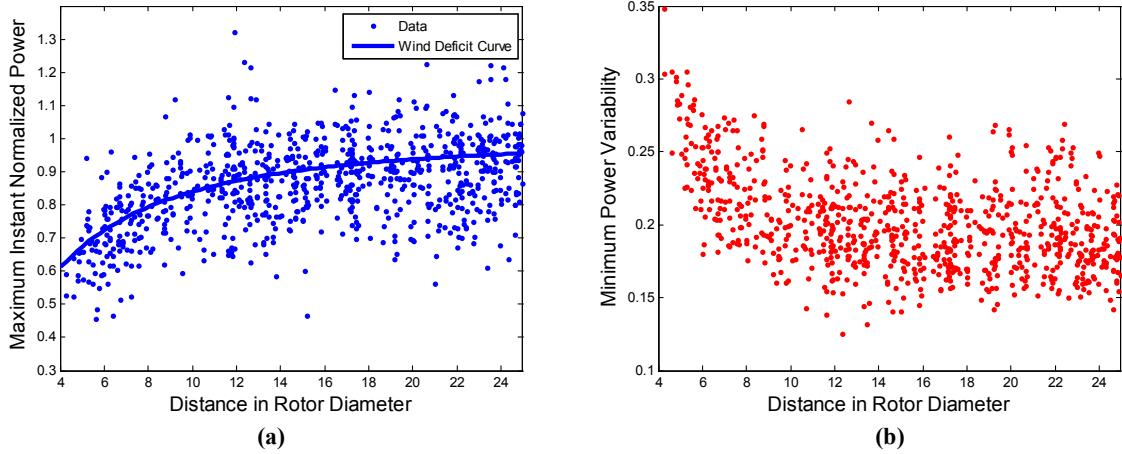


Figure 24. Power Deficit Versus Distance. Maximum instant normalized power and minimum power variance for a downstream turbine were collected for 854 turbine pairs. On the left (a), the maximum normalized instant power is plotted versus distance behind the upwind turbine. The semi-empirical wind velocity deficit from Equation 4 is shown as a solid curve. On the right (b), the minimum power variance for the downstream turbine versus distance is shown.

Channel Speedup Profiles

It has been observed from the normalized instant power plots that there is a significant performance improvement when a turbine is facing the midpoint between two upwind turbines. The physical explanation could be that the wakes of the two upstream turbines displace airflow, which accelerates as in an ordinary channel contraction. Even minor speedup effects would boost the performance of the downstream turbine. For the selected profiles in the Figure 25, the boost is between 1.1 and 1.3 times average. The profiles also display a wake deficit next to the peak, as the downstream turbine faces either one of the two upstream turbines.

Two turbines (17 and 58) show an excessive over performance (1.4 to 1.5 time average) facing what could be described as a duct or channel from the east, formed by rows of multiple neighboring turbines (these turbine profiles are not plotted). See Figure 12.

Channel speedups exhibit a lower variance than the average of the farm, as seen in Figure 25, a might be expected from multiple upstream turbines. Speculatively, this could be explained by the two upstream wakes displacing each other away from the downstream turbine, thus reducing wake turbulence. Further, if there is a speedup effect in a channel type arrangement, turbulence will generally be suppressed due to contraction.

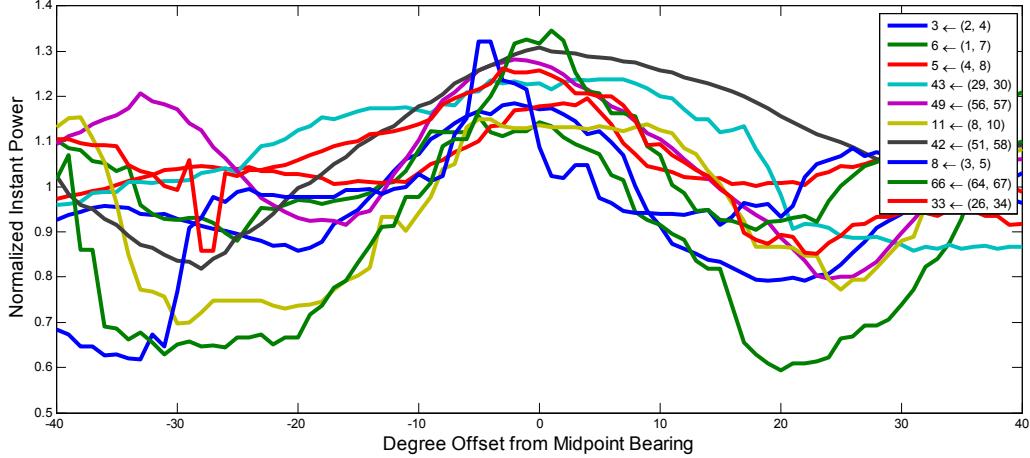


Figure 25. Normalized Instant Power for Channel Speedup. The x -axis gives the wind direction degree difference from the downstream turbine to the midpoint between the two upstream turbines. The y -axis gives the normalized instant power. The channel speed up effect can be seen as the peak at 0° , flanked by wake deficits on either side of the peak.

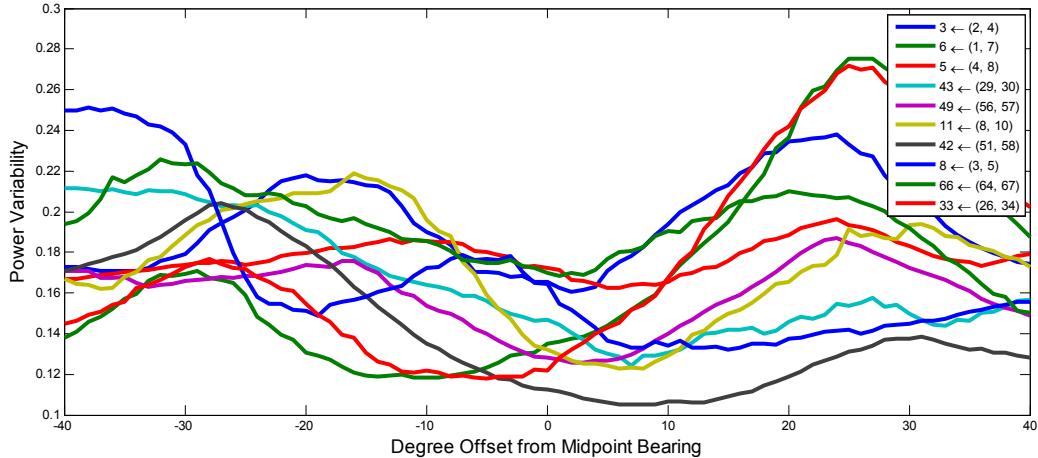


Figure 26. Power Variability for Channel Speedup. As in Figure 25, the x -axis gives the wind offset from the midpoint between the two upstream turbines. The y -axis gives the power variability. The channel speed up is associated with low variability.

Shear Point Speedup Profiles

For shear point speedup, the upstream obstacle, in this case a turbine wake, forces the incoming wind to go around and thus speedup along the side of the farm. This is seen as an increased power in the downstream turbine. The shear point speedup effect is investigated for seven turbines, shown in Figure 27 and Figure 28. Within the dataset, the increase occurs 15° to 20° from the bearing of the upstream turbine, with a magnitude of 1.1 to 1.22 times the nominal farm

level. This magnitude is less than the speedup channel effect, but still significant. Again, this effect is associated with a wake presence. Effects like these have been observed in connection with buildings upstream to wind turbines, (see, e.g., Corscadden *et al.*, [16]). Building speedup is also seen in this data set (for example turbine 7, bearing 80°), but is not investigated further.

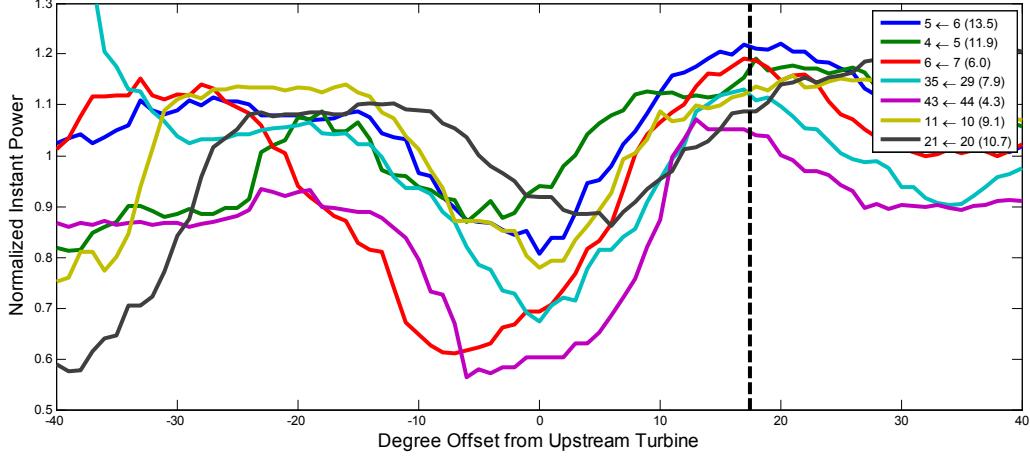


Figure 27. Normalized Instant Power for Shear Speedup. The x -axis gives the wind direction degree offset from the bearing between and upwind and downwind turbine pair, and the y -axis show the normalized instant power. Turbines were selected so that the region between 0° to 90° from the downwind turbine is undisturbed (no turbines nearby). The shear point speedup can be seen as improved power production between 15° and 20° , as indicated by the dotted line at 17.5° .

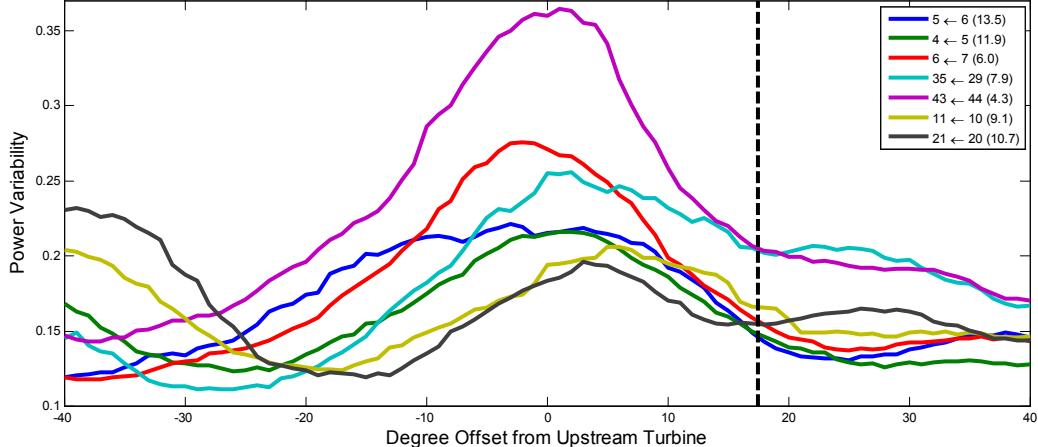


Figure 28. Power Variability for Shear Speedup. As in Figure 28, the x -axis gives the wind direction degree offset from the upwind turbine. The y -axis shows power variability.

Shear Speedup for Multiple Upstream Turbines

Turbines 1, 6, 30, 31, 44 and 45 exhibit a narrow and a very high increased performance peak close to due north. At first, this was misinterpreted as a normalization issue; however, it is due to a displacement from multiple upstream turbines. Turbines 1 and 6 face a long row of turbines in an upstream wind farm at approximately 15° . The 15° is comparable to what is observed for the shear point speedup from a single turbine. Several other turbines, for example 7, 20, 29 and

43, also show the speedup, but the wake profiles are masked by a combination of speedup and waking. The wake profiles for these turbines are shown in Figure 29 and Figure 30.

Turbines 44 and 45 are waked by a turbine and turbine 31 is facing multiple roughness elements. These three turbines show less performance increase than turbines 1, 6 and 30. All the turbines show a high variability towards the shear point to the left and slope down to an all low of 0.12 to the right. The increase in power variability seen in Figure 30 from 25° to 40° can be attributed to landscape roughness by farm buildings and clusters of tall trees (data not shown).

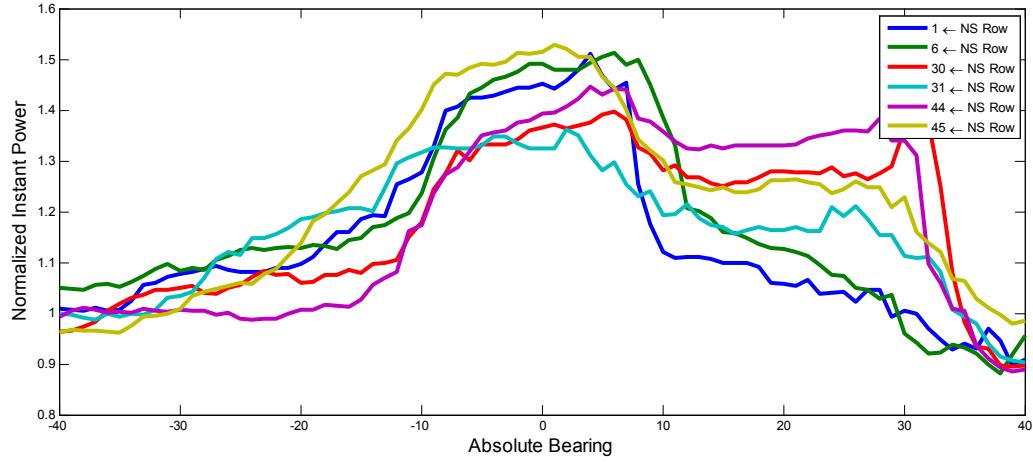


Figure 29. Normalized Instant Power Increase from Multiple Upstream Turbines. The normalized instant power profiles are shown for 6 turbines offset from an upstream row of turbines. The upstream turbine rows are all arranged from north to south.

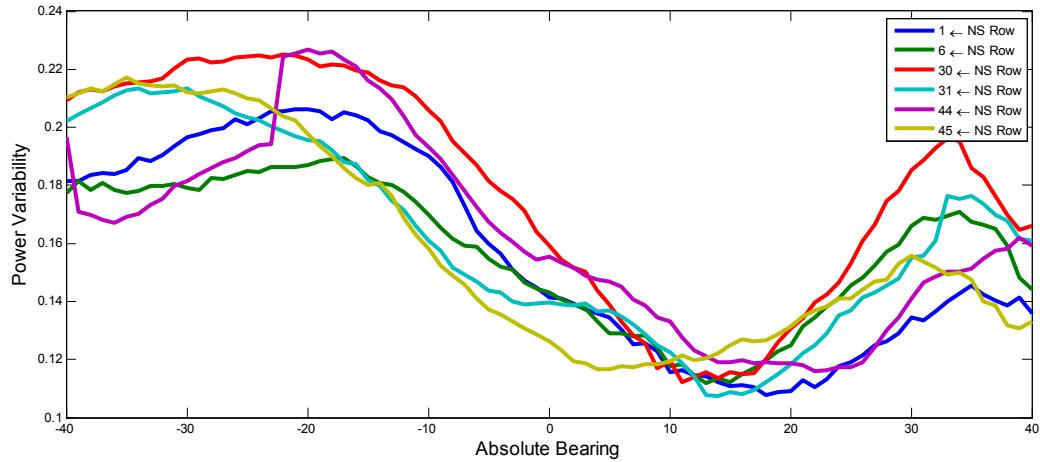


Figure 30. Power Variability from Multiple Upstream Turbines. The power variability profiles are shown for 6 turbines offset from an upstream row of turbines. The upstream turbine rows are arranged North to south.

3.2. Wake Loss Quantification

Using the power observer calculation, percent gain was estimated for parameters in the range $k = 0.01$ to $k = 0.5$ and $x_{\infty} = 5$ to $x_{\infty} = 30$ rotor diameters. This calculation was done over the full range of wind speeds, from 4-20 m/s, with results shown in Figure 31. The results indicate that the power observer calculation is stable with no gain beyond 6% expected, even using unrealistic parameters.

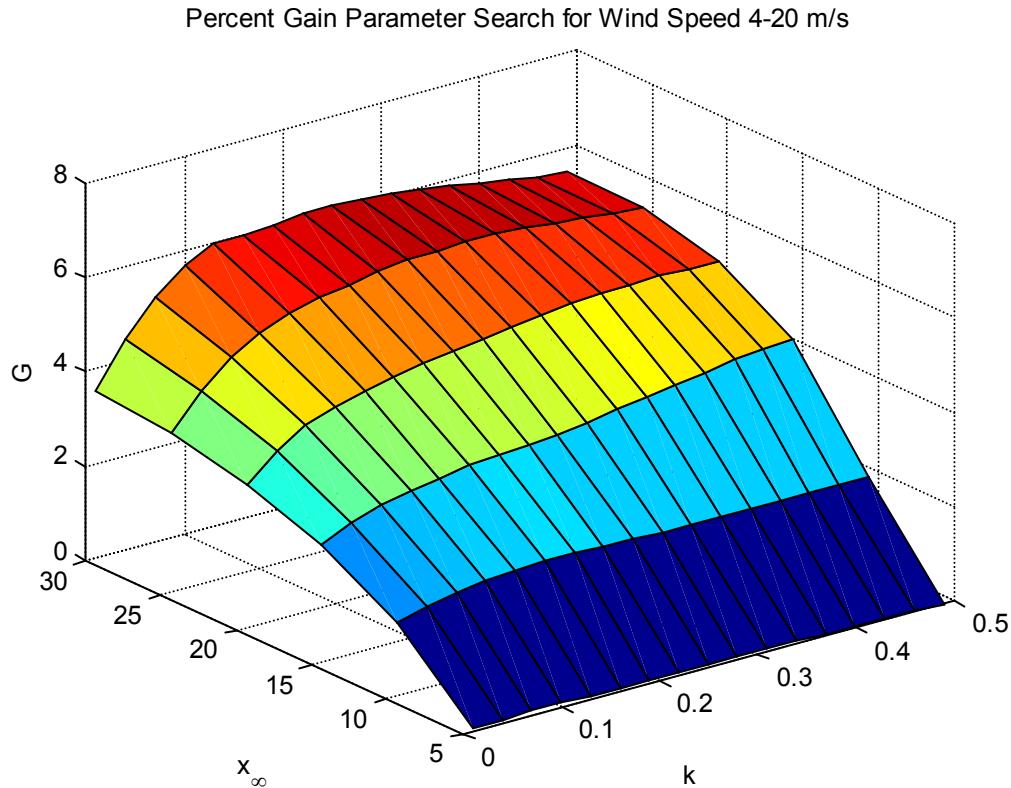


Figure 31. Percent Gain Parameter Search. Percent gain (G) is shown as a function of parameters in the range $k = 0.01$ to $k = 0.5$ and $x_{\infty} = 5$ to $x_{\infty} = 30$ rotor diameters.

Another interesting experiment can be performed if we examine data for wind speeds in the range from 6-10 m/s. In this wind speed range, known as Region II, the turbines are spinning freely without the application of pitch control to the blades. This is in contrast to higher wind speeds (Region III), where blades are pitched in order to limit potential damage. Since the blades are pitched, additional wind energy is available for downwind turbines (because the thrust on the rotor is reduced by pitching).

Thus, it might be expected that the percent gain of the wind farm would increase using data with wind speeds restricted from 6-10 m/s. In fact, this is the case, as can be seen in Figure 32 using data restricted to Region II wind speeds. It is worth noting, however, that the Region II restricted data yields a smaller data set, with approximately 10,000 time points per turbine (as opposed to the original data, which was restricted to wind speeds in the range 4-20 m/s).

Percent Gain Parameter Search for Wind Speed 6-10 m/s

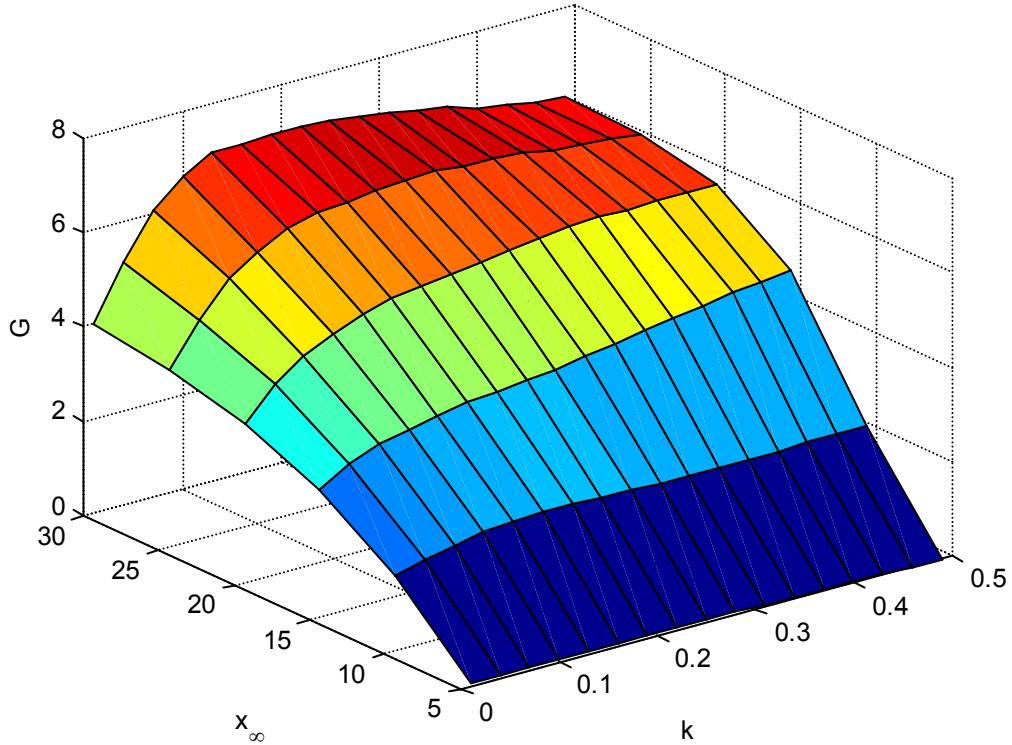


Figure 32. Region II Percent Gain Parameter Search. The percent gain of the wind farm is uniformly higher using wind restricted speeds from 6-10 m/s, as compared to wind speeds from 4-20 m/s (Figure 31). Parameters are again in the range $k = 0.01$ to $k = 0.5$ and $x_\infty = 5$ to $x_\infty = 30$ rotor diameters.

In using a large value of k , e.g. $k = 0.15$ and a large distance requirement, it is only the front turbines which are included in the power observer, thereby reflecting the maximum potential gain of about 7% in region II. By decreasing k , turbines deeper in the array are included in the power observer, and thereby the net gain is less than if only the front row is used. On the other hand, some of these interior turbines may include speedup effects, as discussed in the previous section.

Finally, the percent gain computed over region III is shown in Figure 33. As might be expected, gains are minimal in region III, due to the fact that the turbines are extracting less energy from the wind than is actually present, so that wake effects are minimal. Again, however, we note that the region III dataset is smaller, representing by only 3,000-4,000 time points per turbine.

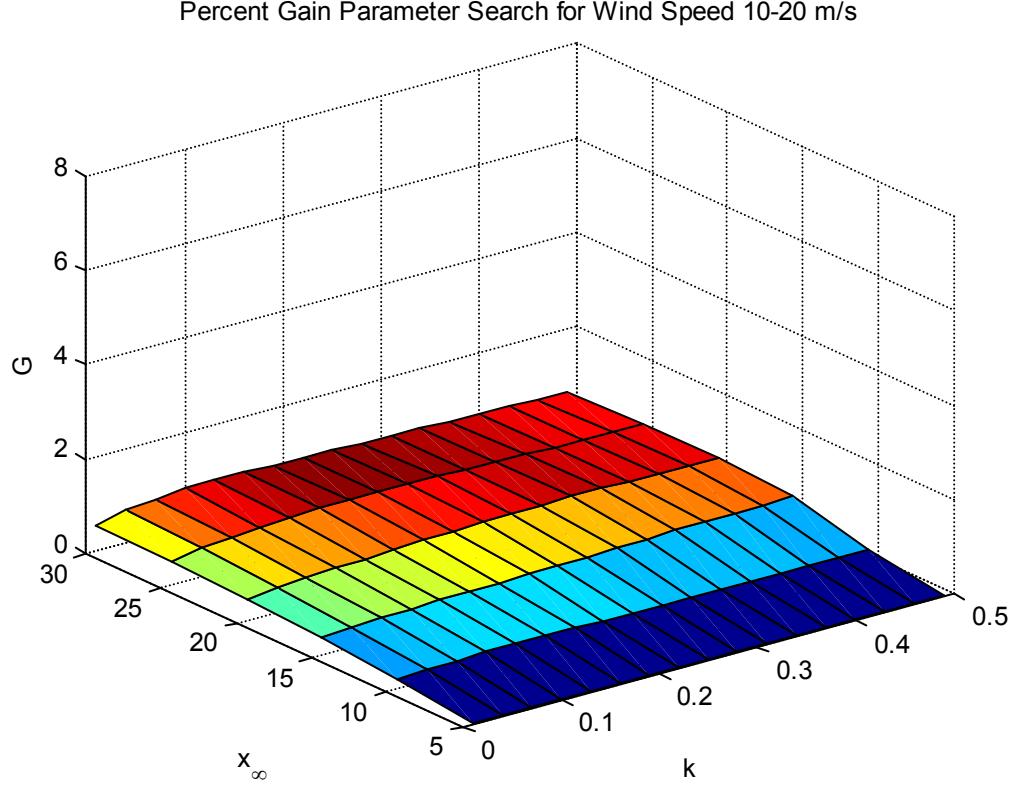


Figure 33. Region III Percent Gain Parameter Search. The percent gain of the wind farm is minimal using wind restricted speeds from 10-20 m/s, as compared to wind speeds from 4-10 m/s (Figure 31 and Figure 32). Parameters are again in the range $k = 0.01$ to $k = 0.5$ and $x_\infty = 5$ to $x_\infty = 30$ rotor diameters.

3.3. Modelling Wakes

In the examples below, we demonstrate the wake model developed in Section 2.4. For our demonstration, 16 wakes were used to make up the composite power rose. In Figure 34, turbines 11, 8, and 10 are shown. These turbines are relatively lightly waked, and show good agreement for turbines 11 and 8, while turbine 10 has some clear discrepancies. In Figure 35, turbines 39, 40, and 54 are shown. These are relatively heavily waked, and gain show good agreement for turbines 39 and 40, while turbine 54 has some clear discrepancies.

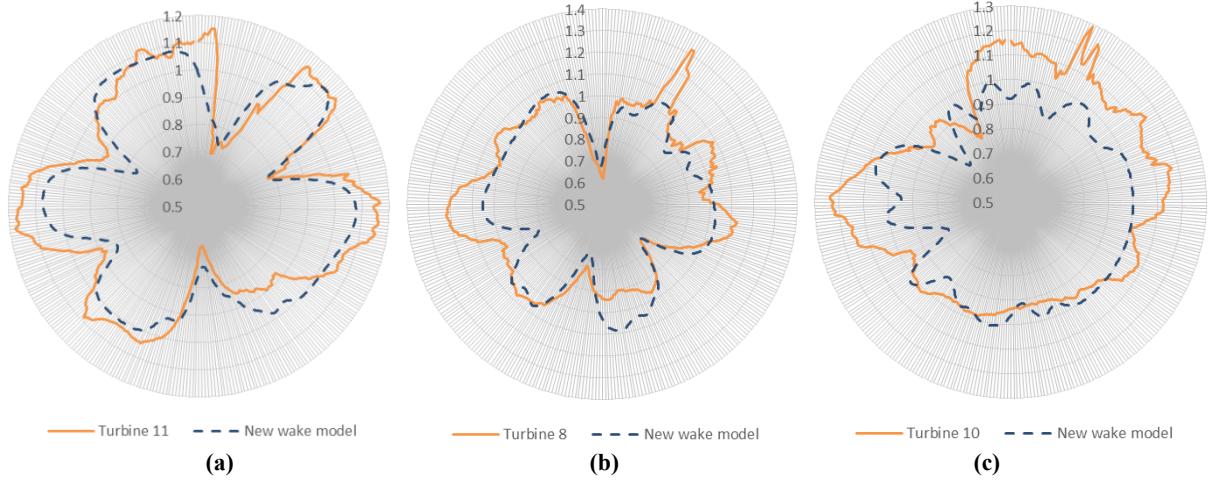


Figure 34. Power roses for turbines 11, 8 and 10. In this figure, we show the power roses as computed using the wake model from Section 2.4. On the left(a), we show the power rose for turbine 11; in the middle (b), the power rose for turbine 8; and on the right (c), the power rose for turbine 10.

Observing Figure 34, all three turbines show various degrees of increased productivity in the 270° direction, which is also seen in neighboring turbines. There are two straight rows of turbines from another wind farm at 285° and some additional turbines in the 270° direction. As previously discovered, rows of turbines typically coincide with a speedup at approximately 15° off center, so it seems reasonable to assume the increased productivity originates from the upwind turbines.

In the Northern direction, turbine 11 shows a strong peak which drops into a wake from turbine 8 at 18° . This is speed up is from the row of turbines 2, 4 and 8. The signature of reduced power variability is also observed (Figure 35). The effect is not seen on turbine 8 as it is facing 2 and is directly waked by 4, with increased power variability, seen in Figure 35. Turbine 10 is also affected by turbines 2, 4 and 8, but in addition there are potentially two channels between turbines 2 and 3 at bearing 350° and turbines 3 and 5 at bearing 10° .

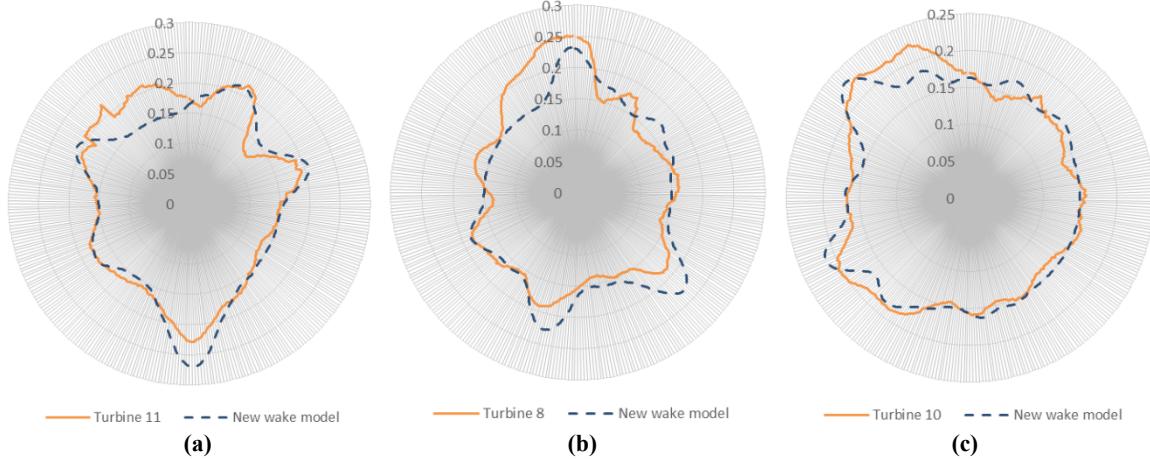


Figure 35. Power variability roses for turbines 11, 8 and 10. In this figure, we show the power variability roses as computed using the wake model from Section 2.4. On the left(a), we show the power variability rose for turbine 11; in the middle (b), the power variability rose for turbine 8; and on the right (c), the power variability rose for turbine 10.

Turbines 8 and 10 show a distinct spike at a 20° bearing. Five other turbines in the vicinity, and several others around the wind farm, also show this peak. Other than noise, there is no obvious explanation for this discrepancy. Finally, the discrepancies seen in the easterly direction seem to be associated with landscape and building features.

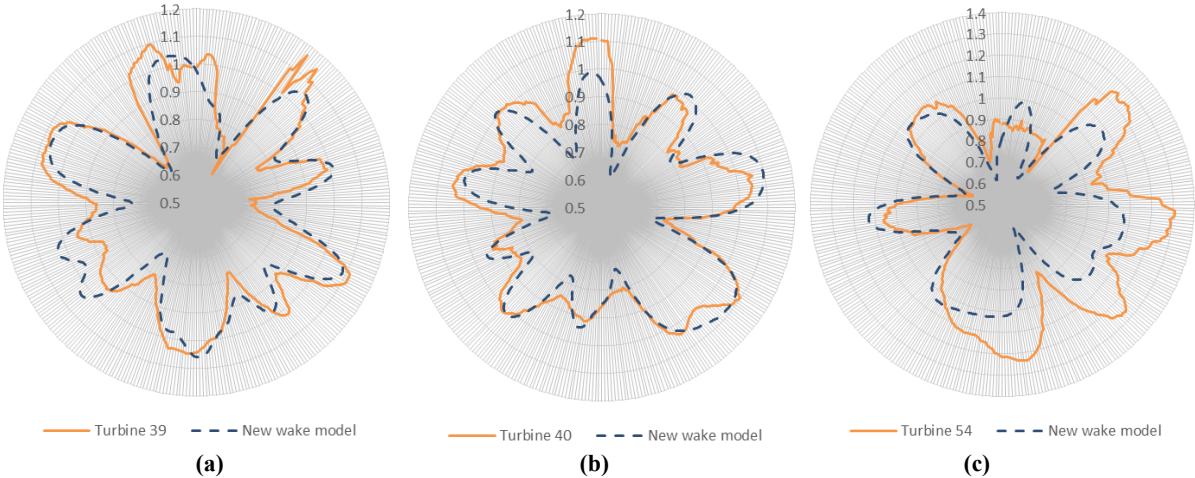


Figure 36. Power roses for turbines 39, 40, and 54. In this figure, we show the power roses as computed using the wake model from Section 2.4. On the left(a), we show the power rose for turbine 39; in the middle (b), the power rose for turbine 8; and on the right (c), the power rose for turbine 54.

Observing Figure 36, turbines 39 and 40 are in a very complex wake situations. The model emulates this well, but we notice the importance of directional accuracy. When turbine 39 is waked from turbines 53 and 52, at 218° and 240° , respectively, these two bearings had to be adjusted by -4° and 4° to show the two quasi symmetrical dimples in this direction. Although there are potentially many speedup channels for turbines 39 and 40, the only clear speed up effect is for turbine 40 is aimed in between turbines 34 and 35, which are very close. However, compared to the three previous turbines, the neighboring turbines are much closer and most of

them are in deeply waked situations themselves, so it is reasonable that the wakes are dominated by the nearest neighbor in this case.

Similar speedup effects are seen for turbine 54, which has two potentially strong channels at 5° and 260° , near turbine 41 and between turbines 57 and 50, respectively. Strong effects are also found at 45° and 95° . The model appears to under predict the performance of turbine 54 between 120° and 190° . However, all neighboring upstream turbines (56, 57, 58, 51 and 59), show a significant and nearly homogeneous over performance in this sector. It is difficult to pinpoint the cause as the nearest upstream wind farm is almost 5,000 meters away.

4. CONCLUSIONS AND FUTURE WORK

Based on an annual nominal directional power performance analysis, four wake effect mechanisms inside a 67 turbine wind farm have been identified. The analysis shows that directional decomposition of turbine performance within the farm just as powerful as classic power curve analysis, while providing a deeper insight into actual performance. The results have also been obtained without the use of wind speeds. However, it is clear that prior to data analysis, the data preparation and validation is critical. The scrubbing of the data for sensor calibration errors and abnormal operational situations are the two most important issues. Furthermore, automatic processing is needed.

In analysis, individual wake profiles and speed-up effects were clearly identified. These are not normally considered in general wake analysis. An individual turbine's average yearly operational situation is complex: even small wind direction changes can have large effects. Wakes were analyzed against nearest neighbor, but in the future, a multi-wake analysis could be extracted from the data as well. It is proposed that a superposition principle could be explored, similar to those applied in classic wind farm models [2-4]. It may also be desirable to revisit the data and explore different wake states based on power performance, i.e. restrict to region III rotor thrust of the power curve and analyze the wind farm when near nominal power, as was seen in Horns Rev data [17].

The data analysis has accurately corrected the turbine yaw position. The capture of the wakes and their characteristics may be sufficiently accurate to determine yaw errors in the upstream turbine based on a wake steering effect.

The wake losses identified are in reasonable agreement with classical models for wake losses. Comparing the observations to existing wake models, these only take into account the wake losses (i.e. recovery to nominal wind speed), but do not include the upside from speedup effects, due to channel or shearing points that has been observed. This new discovery seems to be absent from the existing knowledge base and is essential to understand, in order to improve wind farm performance.

Further, the suppressed variability of the speed up effect could be affecting turbine loading and reliability in a positive way, opposite the manner in which increased variability is known to have an adverse effect. Future correlation of turbine failures with directional observance could improve the understanding of wind farm operations with respect to reliability.

In addition to the strict statistical analysis of the SCADA data, we also developed an automatic method for quantifying wakes losses and applied it to the same 67 turbine wind farm. We found that wake losses are at most 7%. Further, most of the losses occur when the wind is blowing in the east or west directions. These losses coincide with the layout of the farm, since there are more turbines closely sited in the east-west direction than in the north-south direction. On the other hand, these losses do not contribute greatly to the overall losses of the wind farm because the wind does not often blow from the east or west. In other words, the wind farm is well sited, and wake losses are minimized.

This is not to say, of course, that other wind farms are not as well sited, or that wake losses couldn't be greater. In particular, the study Marden *et al.*[17] on the Horns Rev wind farm in Denmark claims potential improvements of 25% using a strategy for optimizing wind turbine control to minimize wake losses. To further benchmark the calculations used in this paper, a comparison with the results of Marden *et al.* [17] on the Horns Rev data should be undertaken.

Further, it would be interesting to perform these calculations on additional wind farms, thus both benchmarking the method and potentially pinpointing farms that might benefit from advanced control techniques. Finally, if our quantification method could be generalized to use only a wind rose collected over time, wake losses in the US fleet could be quantified. The result could potentially provide an absolute measure by compiling regional annual energy production data.

Lastly, we developed an ad-hoc model for wake losses and tested the model fit against the empirical data with reasonable success. The ad-hoc model is interesting because it only depends on geometry so could be applied without SCADA data. However, it needs significant further development and testing.

5. REFERENCES

1. Manwell JF, McGowan JG, Rogers AL. *Wind Energy Explained: Theory, Design and Application* (2nd edition). John Wiley & Sons, 2009.
2. Lissaman PBS, Bates ER. Energy Effectiveness of Arrays of Wind Energy Conversion Systems. Aerovirnoment Report AV FR 7050, Pasadena, CA, 1979.
3. Vermeulen PEJ. An Experimental Analysis of Wind Turbine Wakes. *Proc. Third International Symposium on Wind Energy Systems*, Lyngby, DK, 1980; 431-450.
4. Katic I, Hojstrup J, Jensen NO. A Simple Model for Cluster Efficiency. *Proc. of the 1987 European Wind Energy Conference*, Rome, 1987; 407-410.
5. Gonzalez-Longatt F, Wall P, Terzija V. Wake Effect in Wind Farm Performance: Steady-State and Dynamic Behaviour. *Renewable Energy* 2011; 39(1):329-338. DOI:10.1016/j.renene.2011.08.053
6. Sanderse B, van der Pijl SP, Koren B. Review of Computational Fluid Dynamics for Wind Turbine Wake Aerodynamics," *Wind Energy* 2011; 14:799-819. DOI: 10.1002/we.458
7. Barthelmie RJ, Pryor SC, Frandsen ST, Hansen KS, Schepers JG, Rados K, Schlez W, Neubert A, Jensen LE, Neckelmann S. Quantifying the Impact of Wind Turbine Wakes on Power Output at Offshore Wind Farms. *Journal of Atmospheric and Oceanic Technology* 2010; 27:1302-1317. DOI:10.1175/2010jtecha1398.1
8. Porte-Agel F, Wu Y-T, Chen C-H. A Numerical Study of the Effects of Wind Direction on Turbine Wakes and Power Losses in a Large Wind Farm. *Energies* 2013; 6:5297-5313. DOI:10.3390/en6105297
9. Zaher A, McArther SDJ, Infield DG, Patel Y. Online Wind Turbine Fault Detection through Automated SCADA Data Analysis. *Wind Energy* 2009; 12:574-593. DOI:10.1002/we.319
10. K. Kim K, G. Parthasarathy G, O. Uluyol O, W. Foslien W, S. Shuangwen S, Fleming P. Use of SCADA Data for Failure Detection in Wind Turbines. *Proc. of 2011 Energy Sustainability Conference and Fuel Cell Conference* 2011; NREL/CP-5000-51653.
11. Kusiak A, Li W. The Prediction and Diagnosis of Wind Turbine Faults," *Renewable Energy* 2011; 36(1):16-23. DOI:10.1016/j.renene.2010.05.014
12. Butler S, Ringwood J, O'Connor F. Exploiting SCADA System Data for Wind Turbine Performance Monitoring," *2013 Conference on Control and Fault-Tolerant Systems (SysTol)*, France, 2013; 389-394. DOI:10.1109/systol.2013.6693951
13. Schlechtingen M, Santos IF, Achiche S. Wind Turbine Condition Monitoring based on SCADA Data using Normal Behavior Models. Part 1: System Description. *Applied Soft Computing* 2013; 13(1):259-270. DOI:10.1016/j.asoc.2012.08.033
14. Johansen NA. *Verification of Simulated Fatigue Loads on Wind Turbines Operating in Wakes*, Master's Thesis, Technical University of Denmark, 2010.
15. Aggarwal C. *Outlier Analysis*, Springer, New York, 2013.
16. Corscadden KW, Lubitzb WD, Thomson A, McCabea J. Investigation of Wake Effects on the Energy Production of Small Wind Turbines. *Wind Engineering* 2013; 37(2):151-164. DOI:10.1260/0309-524x.37.2.151
17. Jensen LE. Wake Losses and Turbulence within the Horns Rev Off-Shore Wind Farm," IEA Annex XXIII: Offshore Wind Energy Technology and Deployment, Task 23, Subtask 1, 3.1 Workshop Programme on External Conditions, Layouts, and Design of Off-Shore Wind Farms, Dec. 12-13, Risoe, Denmark, 2005.
18. Marden JR, Ruben SD, Pao LY. A Model-Free Approach to Wind Farm Control Using Game Theoretic Methods, *IEEE Transactions on Control Systems Technology*; 21(4):1207-1214, 2012.

DISTRIBUTION

4 U.S. Department of Energy
Wind and Water Power Technologies Office
Attn: Mike Derby (Electronic Copy)
Patrick Gilman (Electronic Copy)
Nick Johnson (Electronic Copy)
Jose Zayas (Electronic Copy)
1000 Independence Avenue, SW
Washington, DC 20585

1	MS0932	Shawn Martin	9525 (electronic copy)
1	MS1124	Dave Minster	6121 (electronic copy)
1	MS1124	Benjamin Karlson	6121 (electronic copy)
1	MS1124	Brian Naughton	6121 (electronic copy)
1	MS1124	Carsten Westergaard	6121 (electronic copy)
1	MS1124	Jonathan White	6121 (electronic copy)
1	MS0899	Technical Library	9536 (electronic copy)



Sandia National Laboratories