https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

WIND ENERGY SCIENCE DISCUSSIONS

Fast Yaw Optimization for Wind Plant Wake Steering Using Boolean Yaw Angles

Andrew P. J. Stanley, Christopher Bay, Rafael Mudafort, and Paul Fleming

National Renewable Energy Laboratory, National Wind Technology Center, Boulder, CO 80303 USA

Correspondence: Andrew P. J. Stanley (PJ.Stanley@nrel.gov)

Abstract. In wind plants, turbines can be yawed into the wind to steer their wakes away from downstream turbines and achieve an overall increase in plant power. Mathematical optimization is typically used to determine the best yaw angles at which to operate the turbines in a plant. In this paper, we present a new method to rapidly determine the yaw angles in a wind plant. In this method, we define the turbine yaw angles as Boolean—either yawed at a predefined angle or nonyawed—as opposed to the typical methods of formulating yaw angles as continuous or with fine discretizations. We then optimize which turbines should be yawed with a greedy algorithm that sweeps through the turbines from the most upstream to the most downstream. We demonstrate that our new Boolean optimization method can find turbine yaw angles that perform well compared to a traditionally used gradient-based optimizer where the yaw angles are defined as continuous. There is less than 0.6% difference in the optimized power between the two optimization methods for randomly placed turbine layouts. Additionally, we show that our new method is much more computationally efficient than the traditional method. For plants with nonzero optimal yaw angles, our new method is generally able to solve for the turbine yaw angles 50–150 times faster, and in some extreme cases up to 500 times faster, than the traditional method.

1 Introduction

Wind energy capacity has grown rapidly in the United States in recent years (Administration, 2021b, a) and is projected to continue to grow as technology improves, costs decrease (Wiser et al., 2021), and public opinion and policy shift toward wind and renewable energy support (Stokes and Warshaw, 2017). One impactful improvement that has increased wind plant productivity is the use of active turbine yaw control for wake steering within a wind plant. When yawed, a pair of counter-rotating vortices is shed from a wind turbine, causing the downstream wake to deflect (Howland et al., 2016; Bastankhah and Porté-Agel, 2016). In a wind plant, where turbines are built close together to take advantage of high resources and logistical benefits, wake deflection can be actively exploited to steer wakes away from downstream turbines. Although yawed turbines experience a decrease in power production, many studies have shown that steering the wake away from other downstream turbines can result in a net gain for the power plant. This has been shown with models and simulations (Adaramola and Krogstad, 2011; Park et al., 2013; Gebraad et al., 2016; Lin and Porté-Agel, 2020) as well as with field tests (Fleming et al., 2016a, 2017, 2019).

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

To gain maximum performance from a wind power plant for a given wind condition, it is necessary to optimize the yaw angle at which each wind turbine should operate. This optimization often involves nonintuitive trade-offs because individual turbine performance is sometimes sacrificed to increase performance of the wind power plant as a whole. In addition to being nonintuitive, this optimization problem involves complex interactions because slightly adjusting the yaw angle of a single turbine can have effects that propagate to the rest of the wind turbines in the plant—both in their power production and in the wakes that they produce. To solve this optimization problem, the yaw angles of each wind turbine are either formulated as continuous between the upper and lower bounds (Gebraad et al., 2014; Fleming et al., 2016b; Gebraad et al., 2017) or with finely discretized yaw angle selections (Dar et al., 2016; Dou et al., 2020). The problem is then solved with a gradient-based (Fleming et al., 2016b; Gebraad et al., 2017) or gradient-free (Gebraad et al., 2014; Dar et al., 2016; Dou et al., 2020) optimization algorithm that determines the best combination of yaw angles in the wind power plant. While effective and relatively efficient for a one-off wind power plant analysis, there are some shortcomings to this problem formulation.

First, these formulations implicitly assume that real wind turbines are able to precisely achieve any yaw angle desired by the wind plant operator with respect to certain wind resources. In reality, there are significant uncertainties involved with wind measurements and estimations as well as with wind turbine yaw angle estimation (Quick et al., 2020). Thus, formulating the wind plant yaw control optimization problem with continuous or finely discretized yaw angles is unrealistic because real wind turbine uncertainties do not allow such precision.

Second, although the problem formulation and optimization methods currently present in the literature are effective and sufficiently efficient for the problems to which they have been applied, significant improvements in computational efficiency of performing the wind plant yaw angle optimization would enable further advancements in wind plant performance. Computationally efficient yaw optimizations facilitate the optimization of wind plant design, layout, and controls (Fleming et al., 2016b; Gebraad et al., 2017).

These fast yaw optimizations enable wind power plant design and layout optimization to be decoupled from the yaw optimization, which can be performed inside the plant evaluation. Because the yaw optimization is extremely fast, it could be performed within the optimization of the other design variables in a two-step process as part of the function evaluations. This would allow for additional gains over performing the control optimization after the turbine design and layout are fixed. This is impactful because it can reduce the number of design variables in the problem by an order of magnitude or more. Computational expense of optimization problems generally scales unfavorably with the number of design variables, meaning wind plant design, layout, and control codesign can quickly become infeasible, especially for large plants (Zingg et al., 2008; Rios and Sahinidis, 2013; Lyu et al., 2014; Ning and Petch, 2016; Thomas and Ning, 2018).

In this paper, we present a discrete, Boolean wind power plant yaw optimization formulation which, to this author's knowledge, has never been considered before. Each turbine is defined as either yawed or not yawed. This new method can quickly solve for the yaw angles of wind turbines in a power plant, which could enable the yaw angles and the rest of the design variables to be decoupled during optimization. Our new problem formulation and optimization are presented and discussed in comparison to a typical continuous yaw formulation optimized with a gradient-based optimizer.

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

We present the models we used in the paper as well as the optimization formulations. For single wind conditions, we demonstrate the performance of the Boolean problem formulation compared to a typical continuous, gradient-based yaw optimization, and show that there is not a significant sacrifice in performance associated with using the Boolean problem formulation and optimization. We also demonstrate an approximate savings of 50–150 times in computational expense when using our new method, with optimized power production for a random layout within 0.6% of the power from a traditionally used optimization method.

2 Modeling

In this section, we provide a brief overview of the models we used to evaluate wind power plant performance as well as the important wind turbine parameters.

We evaluated the wind plant performance using the open-source software FLOw Redirection and Induction in Steady State (FLORIS) (NREL, 2021). FLORIS is a wind power plant modeling software developed to be computationally inexpensive with optimization in mind. FLORIS is a steady-state, controls-oriented modeling tool that is commonly used in wind power plant control studies and wind plant layout optimization research. There are several modeling options available within the FLORIS framework. For this paper, we used the Gauss-Curl-Hybrid, or GCH, model (King et al., 2021). This is a Gaussian, controls-oriented wake model that captures some of the secondary effects of wake steering that are not captured by other wake models. In addition to the wake deficit and wake deflection captured in with the GCH model, we used the Crespo-Hernandez model to calculate the wake-added turbulence (Crespo et al., 1996) and the square root of the sum-of-squares method for multiple-wake combinations.

In making the important farm level calculations, FLORIS requires several wind turbine parameters. In this paper, we used wind turbine parameters for a 240-m-rotor-diameter, 15-MW turbine. Figure 1 shows the wind turbine dimensions that we used, along with the power and thrust coefficient curves. For the full set of input parameters and model settings used in this study, please refer to the model code input file that can be found within the code repository provided at the end of this paper.

3 Optimization Methods

In this section, we present the two optimization methods that we compare in this paper. We call these methods “continuous,” for the method that is typically used currently in wind plant yaw optimization, and “Boolean,” for our new method. For each, we were only interested in testing our simplified yaw control optimization. The wind turbine locations and design were fixed. We used only positive yaw angles and present results for scenarios in which we are interested in just one wind direction at a time, and for the entire wind rose. The objective of each optimization was to maximize the wind plant power production for the given wind condition or the annual energy production for the year. The design variables were the yaw angle of each wind turbine in the power plant for each wind condition being explored; these were bounded between 0 and 30 degrees. There were

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

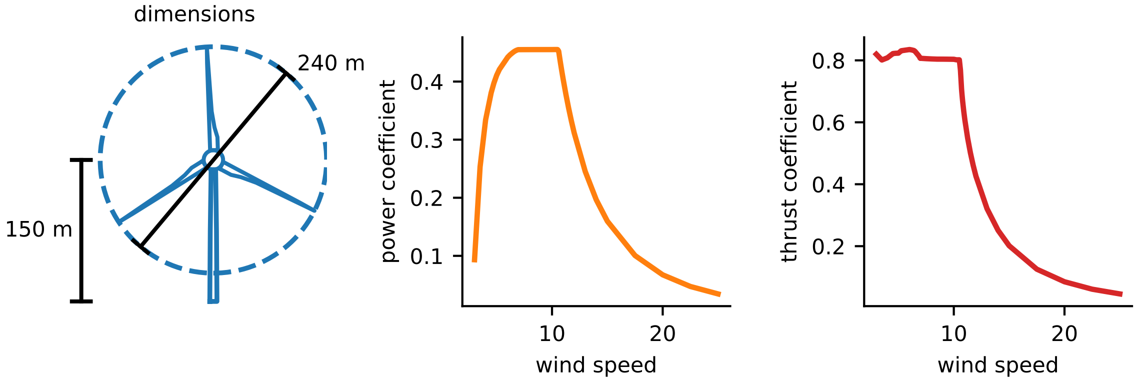

Figure 1. Parameters of the 15-MW wind turbine used in this study. From left to right, this figure provides the important wind turbine dimensions, the power coefficient curve, and the thrust coefficient curve.

no additional constraints beyond those bounding the design variables. The problem can simply be expressed as:

90

$$
\begin{aligned}
&\text{maximize} && \text{plant power or AEP} && \\
&\text{w.r.t.} && \gamma_{n,d} && n = 1 \ldots n_{\text{turbs}} \\
& && && d = 1 \ldots n_{\text{dirs}} \\
&\text{subject to} && 0 \leq \gamma \leq 30^\circ &&
\end{aligned}
$$

where $\gamma_{n,d}$ is the yaw angle of wind turbine n for wind resource d, n_{turbs} is the number of wind turbines in the plant, and n_{dirs} is the number of wind speed and wind direction combinations being considered in the optimization.

3.1 Continuous

For the continuous optimization method, we formulated the yaw angles as continuous variables, which represents a typical method to optimize yaw angles in a wind plant. The yaw angle of each wind turbine in the plant was optimized simultaneously with the commercial gradient-based optimizer SNOPT (Gill et al., 2005) within the pyOptSparse framework in Python (Wu et al., 2020). In this formulation, we normalized the objective function by the initial plant power with zero wind turbine yaw. We also scaled the turbine yaw angles by 0.1, meaning the turbine optimizer saw the design variables with bounds between 0 and 300. These design variables were multiplied by 0.1 within the objective function. We also used finite-difference gradients and started the optimization with each wind turbine at zero yaw. Each other setting was used as default. Refer to the optimization run scripts, found in the code repository linked at the end of this paper, to see details of this gradient-based optimization formulation.

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

3.2 Boolean

For the Boolean formulation, which is new to this paper, we assumed that each wind turbine could only be in one of two different states—yawed or nonyawed. The angle that should be used for the yawed wind turbines is explored in the following section, with the only requirement being that it must be between the upper and lower bounds of 30 and 0 degrees. To optimize the yaw angles in this formulation, we used the following greedy approach:

1. Sort the wind turbines from most upstream to most downstream.

2. Determine which turbines have downstream wind turbines in their wake.

3. From upstream to downstream, check one-by-one if yawing a wind turbine results in an increase in plant power. Fix wind turbine yaws that result in a power increase. Any wind turbines from Step 2 that do not have wind turbines in their wake are skipped and remain unyawed.

This greedy method is very computationally efficient, requiring at most one function call per wind turbine in the plant. Step 2 of the greedy optimization method requires checking to determine whether wind turbines have other downstream turbines in their wakes. To do this, we assumed the wake spread linearly behind each wind turbine using the equation of the Jensen wake model (Jensen, 1983) shown in Eq. 1.

$$r = \alpha x + r_0 \tag{1}$$

In this equation, r is the radius of the wake, α is the wake spread coefficient, x is the distance downstream of the waking wind turbine, and r_0 is the rotor radius of the waking wind turbine. For this paper, we used a large wake spread coefficient of 0.2. If any part of any downstream wind turbine was within this cone behind a wind turbine, the upstream turbine was designated as "waking" and the optimization algorithm above was checked to determine if this waking wind turbine should be yawed. If a wind turbine had no downstream turbines in its wake, the yaw angle was automatically assumed to be 0, and the algorithm did not check to determine if that wind turbine should be yawed.

4 Comparison of Boolean and Continuous Optimization Methods

In this section, we present and discuss the optimization results of our Boolean optimization method compared with the traditional continuous optimization. We compare the performance of each optimized wind power plant as well as the computational expense required for the optimization. We present three different scenarios: (1) wind turbines in a single row in-line with the incoming wind, (2) a regular grid of wind turbines with wind coming from several different directions, and (3) averaged results for wind turbines arranged randomly. For all of the results in this section, the freestream wind speed was set at 10 m/s, which is below the rated wind speed of the wind turbine we used. Past studies have shown wake steering to be most effective below the rated wind speed.

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

4.1 Turbines In-Line with Wind: Power Maximization

In this section, we present the results for wind plant optimizations for a plant with wind turbines that are in-line with the oncoming wind. Before comparing the performance of the different optimization problems, it was necessary to determine the Boolean yaw angle at which the wind turbines should be set. To determine this, we optimized turbine rows with 10–50 wind turbines and with spacings of 3, 5, and 8 rotor diameters between wind turbines using the Boolean optimization method with different Boolean yaw angles from 5–30 degrees at 5-degree increments. Figure 2 shows the optimized percent improvement over the nonyawed baseline case for these different Boolean yaw angles. From left to right, each subfigure shows results for the different wind turbine spacings; within each subfigure, the different lines represent the percent gain for different numbers of turbines. Notice the different y axes for each of the subfigures.

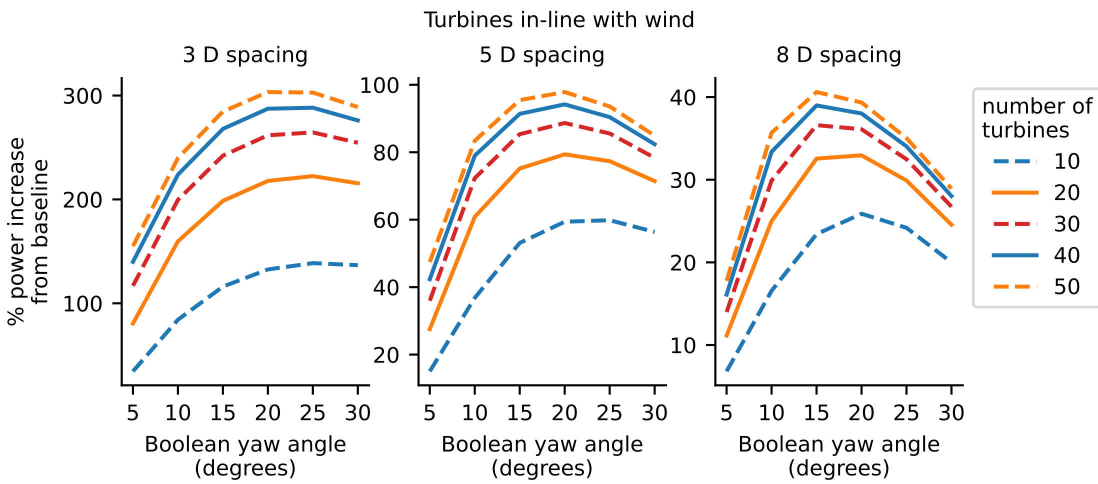

Figure 2. The absolute percent improvement over the nonyawed baseline for the Boolean problem formulation and optimization methods as a function of the Boolean yaw angle. Each subfigure shows the percent improvement as a function of the number of wind turbines in the power plant for Boolean yaw angles between 5 and 30 degrees. From left to right, the different subfigures present power plants with different turbine spacings of 3, 5, and 8 rotor diameters.

In Fig. 2, we see relatively poor performance at small Boolean yaw angles. The performance gains from wake steering increase with increasing yaw angle, reach a maximum, then begin to decrease again. Compared to the larger wind turbine spacings, the smaller wind turbine spacings benefit from larger yaw angles and also achieve a much higher percent improvement over the baseline power when using wake steering. For the power plant with 3 rotor diameter spacing between wind turbines, the optimal performance is almost identical for the yaw angles of 20 and 25 degrees, with a slight edge going to the 25-degree angle. At the 5 rotor diameter spacing, 20 degrees is clearly the best Boolean yaw angle. For the 8 rotor diameter spacing, the best performance is similar—between 15 and 20 degrees—with a small edge to 15 degrees. In each case, a 20-degree Boolean yaw angle is either the best or very close to the best, which led us to select 20 degrees for the remainder of the results in this section.

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

150 With 20 degrees determined as the Boolean yaw angle, we now compare the performance of the traditional, continuous formulation to our presented Boolean formulation. Figure 3 shows the performance of each optimization method as a function of the number of wind turbines in the plant. For this figure, the turbine spacings were held constant at 5 rotor diameters. The top

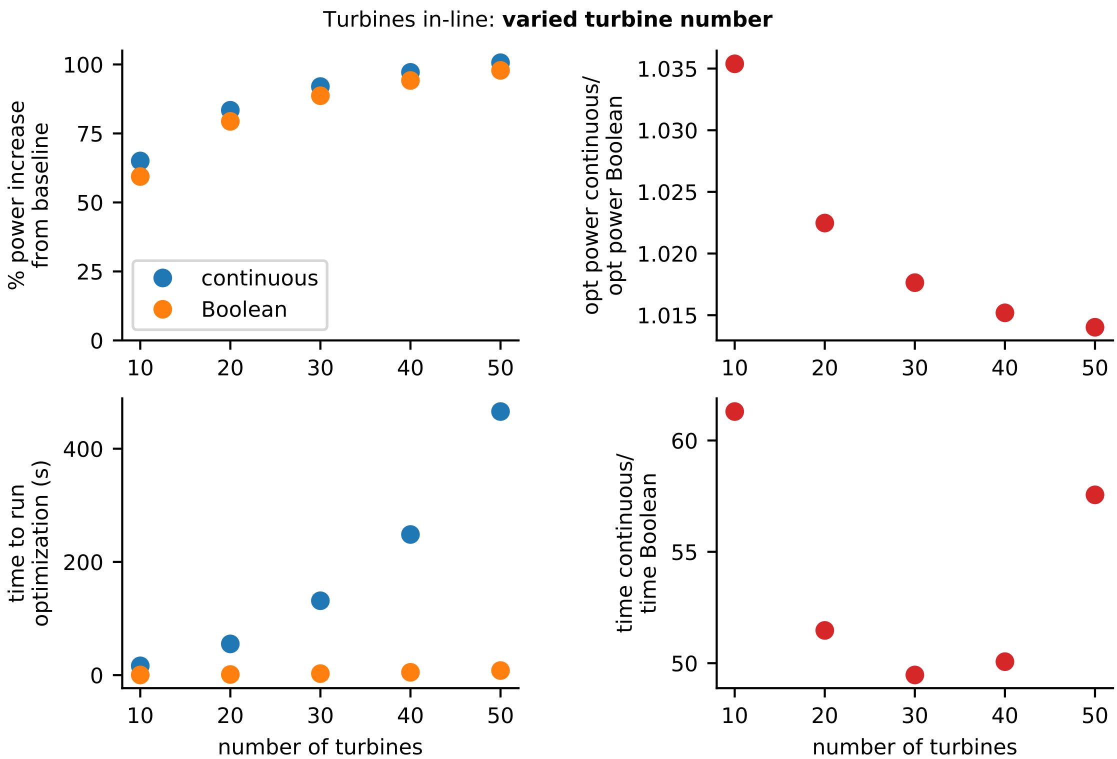

Figure 3. Comparison of a wind plant optimized with a traditional continuous formulation compared to our novel Boolean optimization. The results in this figure are for a single row of turbines in-line with the incoming wind. For the purposes of this figure, the spacing between wind turbines was constant at 5 rotor diameters, while each subfigure shows different metrics as a function of the total number of wind turbines in the plant. The top left subfigure shows the absolute percent increase in power over the nonyawed baseline for the continuous and Boolean formulations. The bottom left subfigure shows the absolute time required to run each optimization, again for the continuous and Boolean formulations. The top right subfigure shows the ratio of the optimized power with the continuous formulation to the optimized power with the Boolean formulation. The bottom right subfigure shows the ratio of the time required to optimize the continuous formulation to the time required to optimize the Boolean formulation.

two subfigures show the performance of each optimized plant. The top left subfigure shows the absolute percent improvement of each optimization method over the nonyawed baseline. The general trend and actual values for both the continuous and Boolean
155 optimizations are very similar in this subfigure. The percent improvement for using yaw-controlled wake steering increases with more wind turbines but begins to level out as a larger portion of the power plant operates under deep-array steady-state conditions. Although the trends are the same among each optimization method, the continuous formulation performs

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

slightly better. The top right subfigure helps us see how much better the continuous formulation performs compared with the Boolean optimization. With 10 wind turbines, the power production from the continuous optimization is about 3.5% higher than the Boolean optimization. This difference then decreases to less than 1.5% with 50 wind turbines. While 1.5%–3.5% is certainly a non-negligible improvement in the wind plant power production, the similarity in power production obtained using the continuous and Boolean formulation is sufficiently close for the purpose of control co-design. For actual operation, the continuous optimization can be used to determine the yaw angles for each turbine to capture the additional percentage points of improvement.

The bottom two subfigures of Fig. 3 show the difference in computational expense between the continuous and Boolean formulations. The bottom left subfigure shows the absolute time required to run each optimization. For the continuous formulation, the computation time is seen to increase exponentially with increasing design variables. As the number of wind turbines increases, the total number of function calls for optimization and the time for each function call increase, leading to poor computational scaling with increasing plant size. The computational expense for the Boolean optimization formulation also increases with power plant size, although the scale is much smaller, such that the computation time is minuscule and flat compared to the continuous computation time. The bottom right subfigure shows the ratio of time required for the continuous optimization to the Boolean optimization. As seen in the figure, the Boolean optimization was 50–60 times less computationally expensive than the continuous optimization.

While Fig. 3 shows the comparison of the different optimization methods as a function of the number of wind turbines, Fig. 4 shows the comparison of methods for a constant 50 wind turbines but for varied turbine spacing, from 3–8 rotor diameters. The subfigures in this figure represent the same information as that shown in Fig. 3, but for varied spacing. In the top left subfigure, we see that there are decreased gains from wake steering as the spacing of wind turbines increases. This is because, as wind turbine spacing increases, the wakes have more time to recover before reaching the downstream turbines. Thus, wake avoidance through wake steering is not as beneficial because the wind speed in the wakes is closer to the freestream. In the top right subfigure, we also see that the difference in the percent gain from the continuous optimization and the Boolean optimization is largest for the smaller wind turbine spacings. This indicates that the continuous formulation is more beneficial in scenarios of extreme waking, where small yaw adjustments can lead to a larger increase in plant power. In the top right subfigure, the results are similar to those in Fig. 3, where for 50 turbines the optimal power from the continuous optimization is between 1.4% and 2.2% greater than the Boolean power. In the bottom two subfigures, we see the difference between the computational expense for the various problem formulations and see that the Boolean formulation was 40–130 times faster than the continuous formulation.

Figures 3 and 4 show the comparison of optimized performance and computation time for a line of wind turbines in line with the incoming wind. From the scenarios optimized, there are a few key conclusions. First, the optimal Boolean yaw angle was found to be 20 degrees, which performed the best overall for different numbers of wind turbines and turbine spacings. This Boolean yaw angle appears to be sensitive to the wind turbine spacing and is likely a function of the wind speed as well. Second, the majority of the increase in power production from wake steering can be achieved with a Boolean problem formulation. Third, the continuous problem formulation still performs better than the Boolean formulation, between 1.5% and

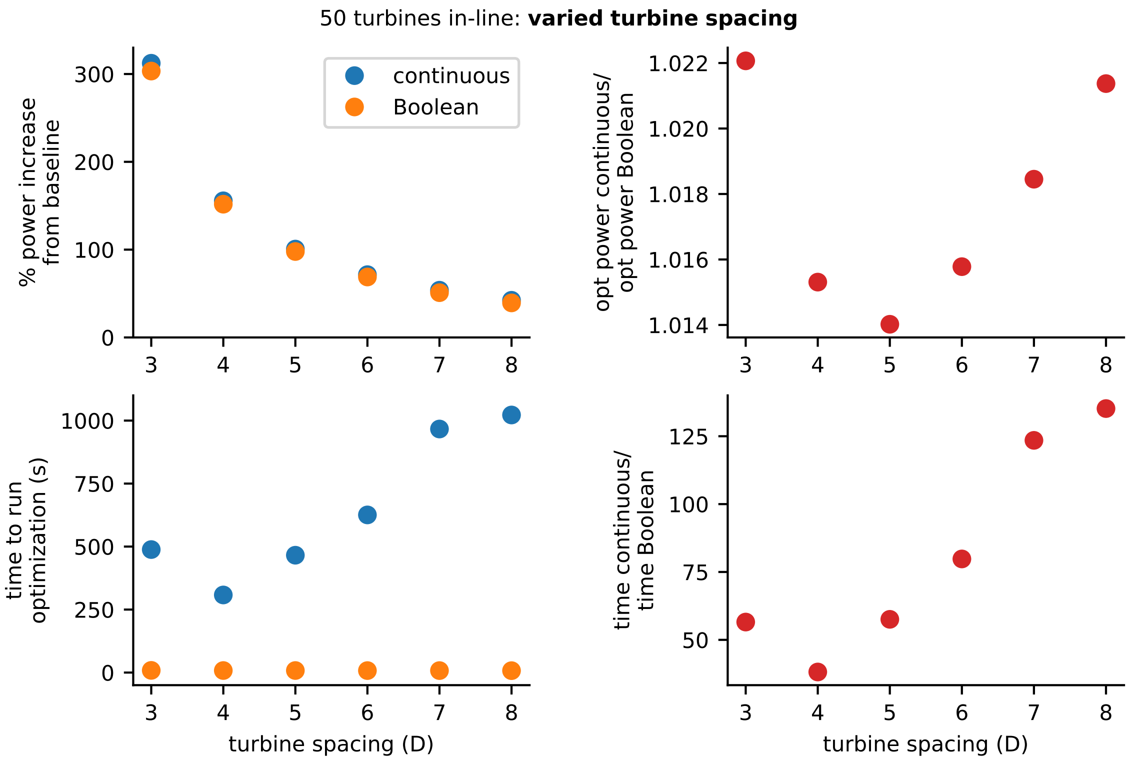

Figure 4. Comparison of a wind plant optimized with a traditional continuous formulation compared to our novel Boolean optimization. The results in this figure are for a single row of turbines in-line with the incoming wind. For the purposes of this figure, the number of wind turbines was constant at 50, while each subfigure shows different metrics as a function of the spacing between wind turbines. The top left subfigure shows the absolute percent increase in power over the nonyawed baseline for the continuous and Boolean formulations. The bottom left subfigure shows the absolute time required to run each optimization, again for the continuous and Boolean formulations. The top right subfigure shows the ratio of the optimized power with the continuous formulation to the optimized power with the Boolean formulation. The bottom right subfigure shows the ratio of the time required to optimize the continuous formulation to the time required to optimize the Boolean formulation.

3.5% better, depending on the number of turbines and the turbine spacing. Fourth, the Boolean problem formulation is able to solve the optimization much faster than the continuous formulation—between 40 times and 130 times faster. While the turbines in-line with the incoming wind provide a great example case, with the worst-case waking scenario, this case does not present the entire scenario. In reality, wind turbines are most often arranged in a grid or more random layout distributed over the landscape and are not usually directly in line with the incoming wind.

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

4.2 Turbines Arranged in a Grid: Power Maximization

In this section, we discuss a more realistic scenario where the wind turbines are placed in a regular grid. Although grid arrangements perform suboptimally compared with freely optimized layouts, grid layouts are easier to design and build, and there are often restrictions that require a grid layout. Also in this section, we compare the continuous and Boolean optimizations for different grid sizes and for different wind directions in the grid. For this section we assumed a constant grid spacing of five rotor diameters and a Boolean yaw angle of 20 degrees. Figure 5 shows the wakes for a nonyawed, 5-by-5 grid wind power plant for the 6 wind directions we considered between 270 degrees (due west) and 345 degrees in 15-degree increments. As seen in this figure, some wind directions result in high wake interactions between wind turbines, such as 270 degrees and 315 degrees; while others have minimal wake interactions, such as 300 degrees and 330 degrees. One can expect high gains from wake steering for the wind directions with the most waking, although a priori it is difficult to determine how the continuous and Boolean optimizations will compare.

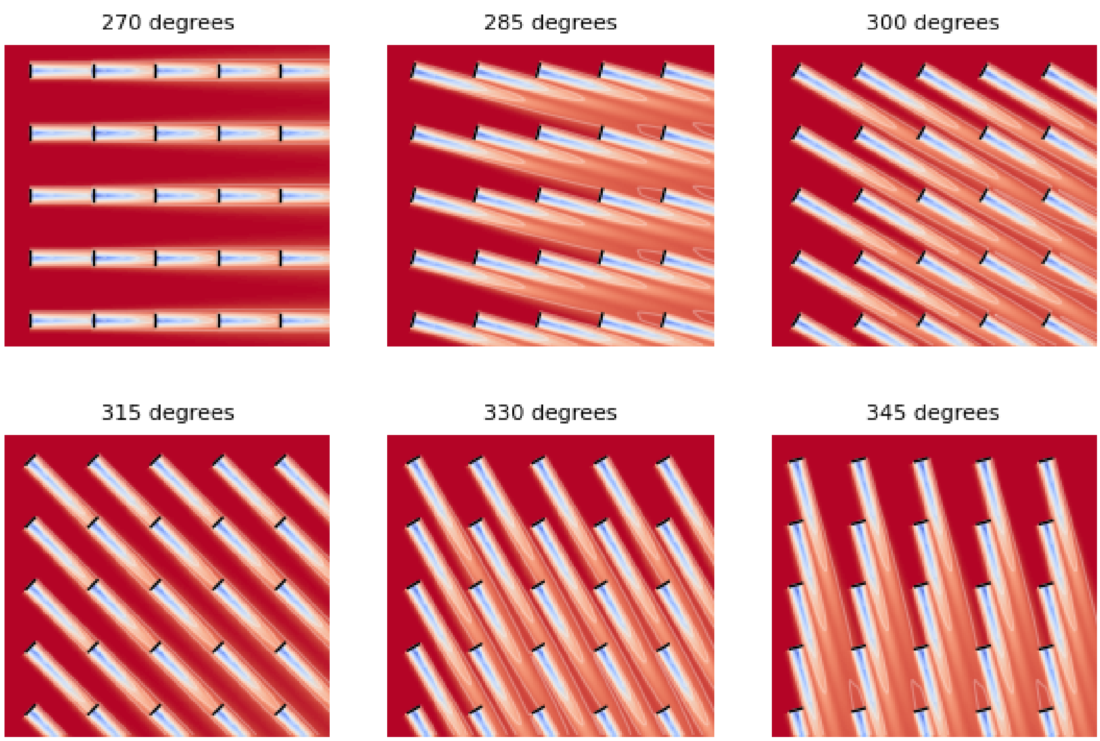

Figure 5. The flow field for a 5-by-5-square grid wind plant for different wind directions.

Figures 6, 7, and 8 show the results of the grid optimizations for different numbers of grid rows and for different wind directions. Figure 6 shows the percent increase in power that wake steering achieves compared to a nonyawed baseline for each of the optimization methods. Notice that for wind directions of 270 degrees and 315 degrees, the Boolean optimization

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

looks very similar to the continuous optimization—almost like the results for the wind turbines that were in-line with the wind direction. For these two wind directions, the turbines are behaving similar to the in-line wind plant. The interaction of the normal grid and the wind direction means that the power plant is just made of several rows in-line with the wind placed side by side. For the wind directions of 285 degrees and 330 degrees, there is very little or no performance improvement for either optimization method. For these wind directions, the grid is oriented such that there is very little wake interaction between wind turbines and, where there is wake interaction, it is very far downstream, such that the wake has already mostly recovered. Finally, for the wind directions of 300 degrees and 345 degrees, there is a more significant difference between the percent improvement achieved with the two different optimization methods. For these wind directions, the continuous optimization is able to realize about twice the percent improvement than the Boolean optimization, when compared to the baseline nonyawed case. Even though the percent improvement is small, this behavior is different than the other cases explored so far in which the Boolean optimization was able to provide most of the benefit that the continuous optimization provided.

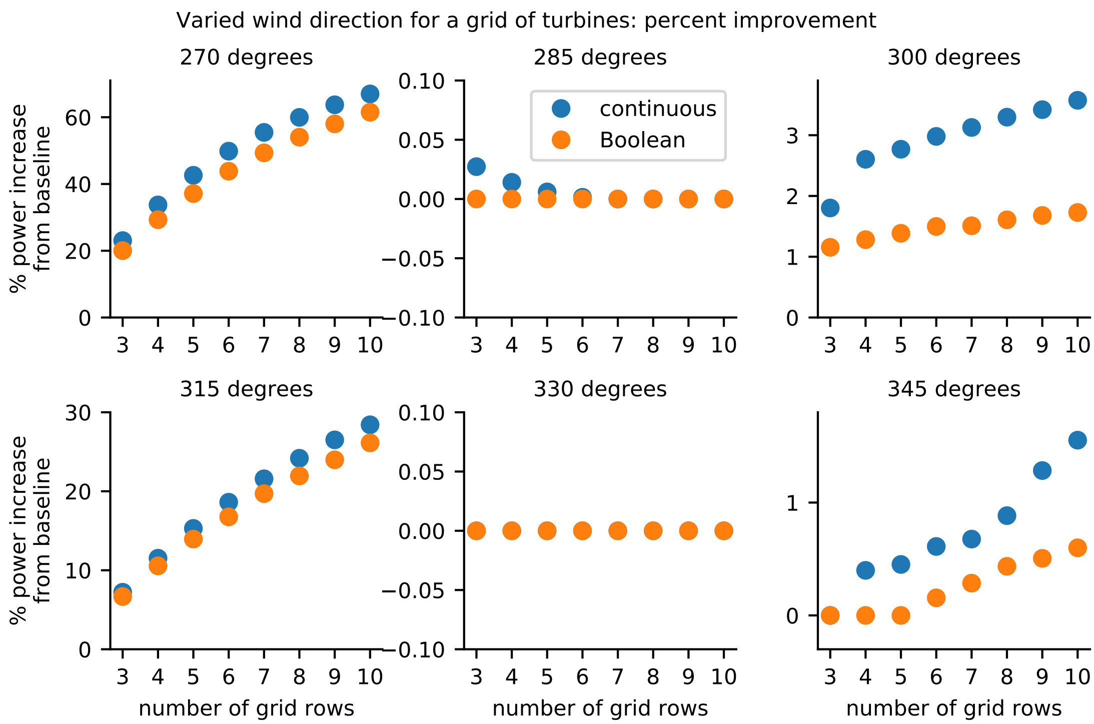

Figure 6. Comparison of the absolute percent increase in power over the nonyawed baseline wind plant optimized with a traditional continuous formulation compared to our novel Boolean optimization. These results are for a grid wind power plant where, within each subfigure, the x-axis indicates different numbers of rows in the plant. Each subfigure shows results for a different wind direction.

Figure 7 shows the ratio of the optimal power achieved with the continuous optimization to the Boolean optimization. For the wind direction of 270 degrees, the continuous optimization provided yaw angles that performed 2%–4% better than the

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

Boolean optimization, which is slightly higher than the percentage gain for the in-line wind turbines from the previous section. This additional benefit of the continuous optimization appears to be because there are fewer wind turbines in-line for the grid optimization, with only 3–10 rows. For the wind direction of 315 degrees, the continuous optimization only performs up to 2% better than the Boolean optimization. This is additional evidence for what we already saw in the previous section: When wind turbines are spaced further apart, there is less of an advantage to the continuous optimization. The results for wind directions of 285 degrees and 330 degrees are trivial—there is no power gain from yaw control with either optimization method, meaning that the ratio is 1. For wind directions of 300 degrees and 345 degrees, the continuous optimization method again performs up to 2% better than the Boolean optimization. Even though the relative percent gain between optimization methods was different for these wind directions, the absolute percent gain was very small.

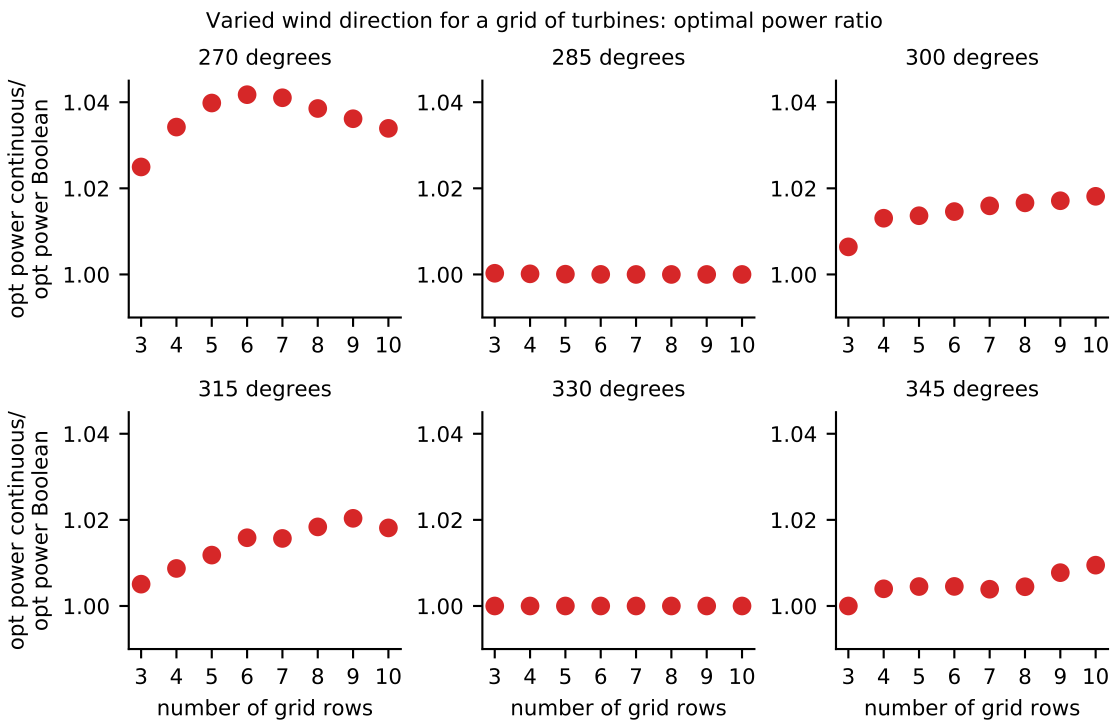

Figure 7. Comparison of the optimal power achieved with wake steering with yaw angles optimized with a traditional continuous formulation compared to our novel Boolean optimization. These results are for a grid wind plant where, within each subfigure, the x-axis indicates different numbers of rows in the power plant. Each subfigure shows results for a different wind direction.

Figure 8 shows the ratio of time required to optimize each plant with the two different optimization methods. First, let's examine the two right columns in this figure. For each of these optimizations, the difference in computational expense between the continuous and the Boolean methods is small compared to the 50–100 times multiplier we saw for the in-line power plant results. If we refer back to Fig. 6, we see that the percent gain from wake steering is nonexistent or very small for these wind

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

directions. This indicates that the optimized yaw angles were close to zero throughout the plant and there was relatively low sensitivity of plant power to the yaw angles of wind turbines in the plant. Thus, the continuous optimization converged quickly and was not notably superior to the Boolean optimization. Now, let's examine the left-hand column, which shows the results for wind directions of 270 degrees and 315 degrees. For these directions, the wind turbines in the grid are directly in-line with the incoming wind. For a wind direction of 270 degrees, the wind is in-line with the grid rows and, for 315 degrees, the wind is in-line with the grid diagonals. As we saw in previous results, when wind turbines are in-line with the incoming wind, the Boolean optimization solves much more quickly than the continuous optimization. For the grid, this affect appears to be exaggerated because it consists of several rows of wind turbines in-line with the wind. For these two wind directions, the Boolean optimization is about 150–500 times faster than the continuous optimization.

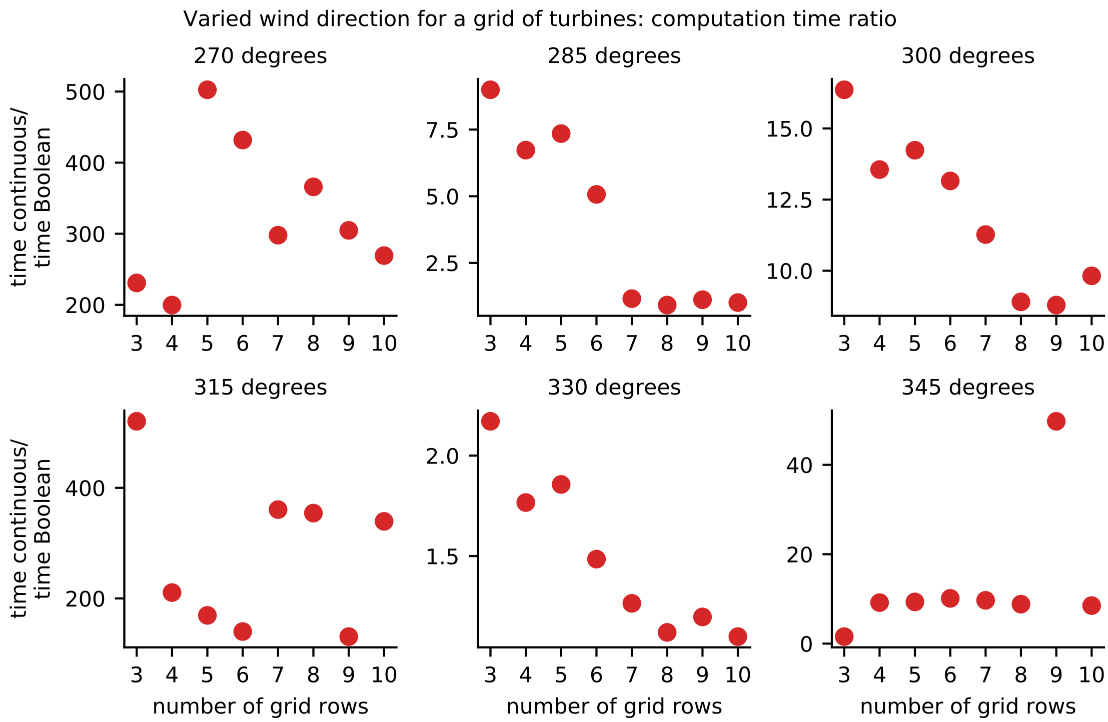

Figure 8. Comparison of the time required to solve the plant yaw angle optimization problem with yaw angles optimized with a traditional continuous formulation compared to our novel Boolean optimization. These results are for a grid wind power plant where, within each subfigure, the x-axis indicates different numbers of rows in the plant. Each subfigure shows results for a different wind direction.

4.3 Turbines Arranged Randomly

Sections 4.1 and 4.2 present and discuss the comparison of performance for each optimization method for regularly arranged wind plants, in a line and in a grid. In this section, we explore how the yaw optimization methods perform in plants with

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

the wind turbines arranged randomly. For sections 4.1 and 4.2, we used a Boolean yaw angle of 20 degrees for all of the performance comparison optimizations. While we showed that this was appropriate for the regularly arranged wind plants, it is possible that another angle is more appropriate for a random, irregular layout. Figure 9 shows the results of our test of which Boolean yaw angle is optimal. For this figure, we randomly generated seven wind plant layouts with the indicated number of wind turbines, with an average spacing of 5 rotor diameters. We assumed the wind came from due west for each of the optimizations, and we optimized the yaw angles in each of the 7 layouts using the Boolean optimization method. The results in the figure show the average performance of the 7 random layouts for each of the Boolean yaw angles that we tested and for each of the numbers of wind turbines. Because the layouts for Fig. 9 were randomly generated, there is little meaning to the

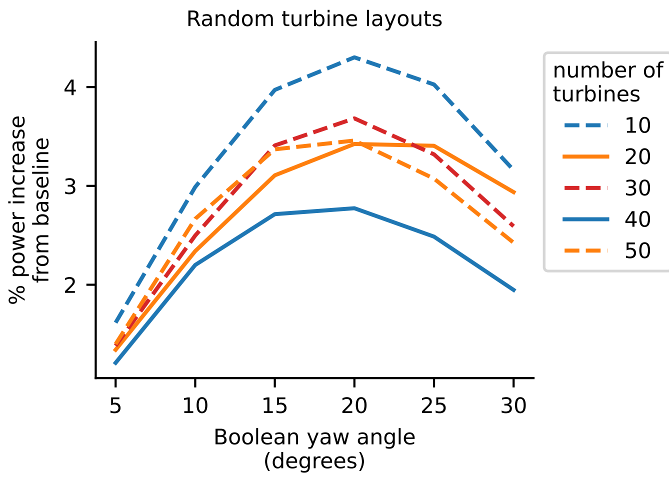

Figure 9. The absolute percent improvement over the nonyawed baseline for the Boolean problem formulation and optimization methods as a function of the Boolean yaw angle. The results shown are averaged for 7 randomly generated turbine layouts for the different numbers of turbines indicated on the x-axis. Boolean yaw angles between 5 and 30 degrees are shown.

trends of performance increase for the different numbers of wind turbines. However, in this figure we can see that 20 degrees is again the superior Boolean yaw angle, as we saw for the regular layouts.

Figure 10 compares the optimal performance of the continuous and Boolean optimization methods. As we did in discussing previous results, let's first examine the top row of subfigures, which compares the performance of the wind plants optimized with the different methods. As with Fig. 9, these results are the average of 7 randomly generated layouts with an average spacing of 5 rotor diameters, with different numbers of wind turbines indicated on the x axes. As was determined from Fig. 9, the Boolean yaw angle for these optimizations was 20 degrees. The top left subfigure shows the percent improvement achieved from wake steering compared to the nonyawed baseline. Notice that for these random layouts, the Boolean optimization performs very well compared to the continuous optimization, capturing the majority of the power gain from wake steering with the more simple optimization method. The top right subfigure shows the ratio of the optimized continuous power to the optimized Boolean power. In this figure, we see that for all numbers of wind turbines, the Boolean optimization performs within 1%

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

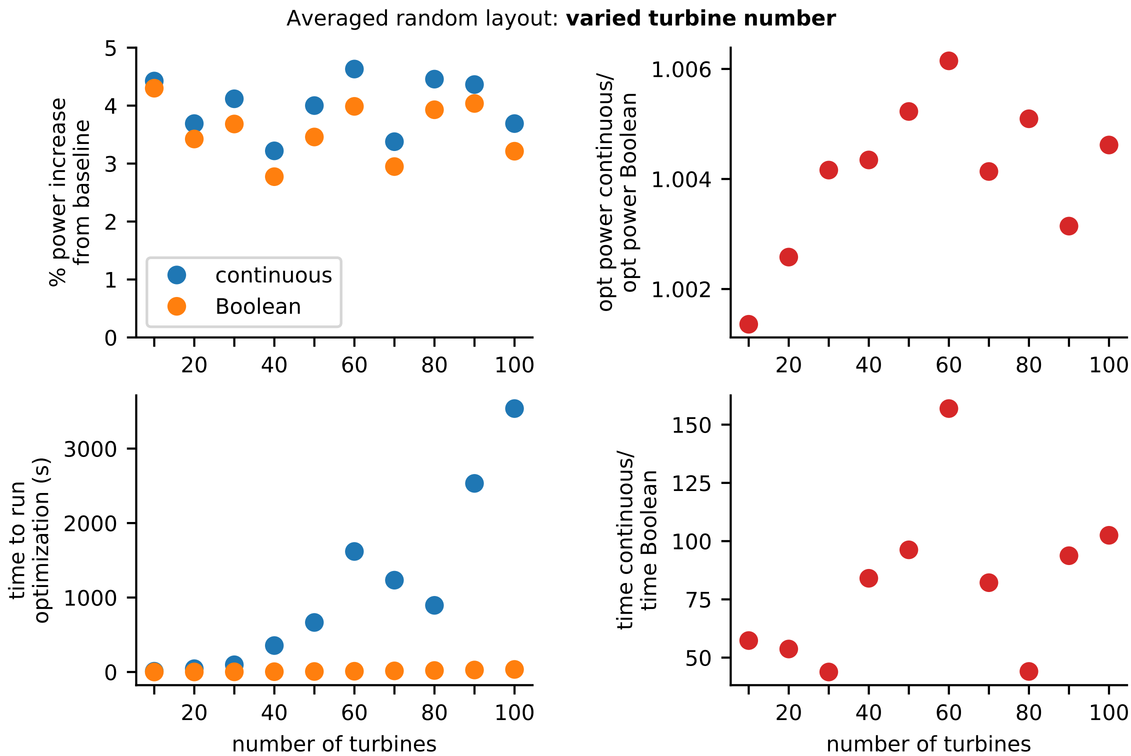

Figure 10. Comparison of a wind plant optimized with a traditional continuous formulation compared to our novel Boolean optimization. The results in this figure are the averaged results for 7 randomly generated wind turbine layouts for plants with different numbers of turbines. For the purposes of this figure, the average spacing was constant at 5 rotor diameters while each subfigure shows different metrics as a function of the number of wind turbines. The top left subfigure shows the absolute percent increase in power over the nonyawed baseline for the continuous and Boolean formulations. The bottom left subfigure shows the absolute time required to run each optimization, again for the continuous and Boolean formulations. The top right subfigure shows the ratio of the optimized power with the continuous formulation to the optimized power with the Boolean formulation. The bottom right subfigure shows the ratio of the time required to optimize the continuous formulation to the time required to optimize the Boolean formulation.

and, for most of the results, within 0.5% of the continuous optimization. This is much closer than the comparison of the two optimization methods for the previous regular layouts in which the difference with the optimized power was sometimes as high as 4% in some extreme cases.

The comparison for the computational expense of each optimization method is shown in the bottom row of Fig. 10. These timing results are similar to the results from the in-line power plant results, where the Boolean optimization is about 50–100 times faster than the continuous optimization, with one outlier about 150 times faster. In this random yaw optimization, there are always some wind turbines that are significantly waking. Because of this, the plant power is sensitive to some nonzero wake angles, which means the continuous formulation always takes much longer than the simple Boolean formulation.

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

5 General Discussion

In this section, we discuss the overarching performance of our new Boolean optimization method compared to the traditional continuous optimization method. We will discuss the optimal Boolean yaw angles that should be used, the performance of the optimized wind power plants with each method, the differences in computation time, and the potential applications of our Boolean method.

5.1 Optimal Boolean Yaw Angle

In Section 4, we determined that when using our new Boolean optimization method the best plant performance occurred with a Boolean yaw angle of 20 degrees. This was determined by comparing wind power plants that were optimized with different wind turbine spacings and numbers of wind turbines. For the cases that we compared, the yaw angle of 20 degrees provided either the best or close to the best plant performance for both the regular line and grid turbine layouts, as well as the irregular random turbine layouts. While it is clear from Fig. 2 that the optimal Boolean yaw angle has some sensitivity to the turbine spacing in the wind plant, we also expect that it is sensitive to the wind speed, which we did not test in this paper. It may be important to tune the Boolean yaw angle to the exact scenario at which the wind plant will operated, or even more finely adjust the Boolean yaw angle for additional gains when operating in the scenarios demonstrated in this paper. Because of the minimal computational expense required to run the Boolean optimization, this tuning of the yaw angle can be quickly achieved with very little effort.

5.2 Performance of Optimized Plants

In general, we see from Section 4 that the Boolean optimization method is able to achieve most of the gains from wake steering that the continuous optimization method can achieve. Additionally, in general, the optimized power from the Boolean optimization is very close to that of the continuous optimization. The Boolean optimized plants had the best comparison to the continuous optimized plants for the random turbine layouts. This indicates that the Boolean optimization method would perform particularly well for land-based wind power plants where the layout is not constrained to a regular grid. For land-based plants, wind turbine placement is often determined to a large extent by terrain features, land availability, and spatial constraints from local regulations, resulting in a more irregular layout where the Boolean method could perform well.

However, the Boolean optimized plants performed the worst compared to the continuously optimized plants for the regular turbine layouts with wind turbines that were directly in-line with the incoming wind. In these cases, the Boolean plants were between 1.5% and 4% worse than the continuous optimization plants. For the regular grid layouts where the wind direction resulted in turbines that were slightly offline with the wind direction, the Boolean method resulted in plants that performed about 0.5%–2% worse than the continuous method. At first glance, this seems to indicate that the Boolean optimization method may not be appropriate for wind power plants with a regular wind turbine layout, such as offshore wind plants in the United States where layouts are restricted to grids. However, even with these results, we expect the Boolean optimization method to be appropriate.

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

For the results shown in this paper, we only considered the power production comparison between the Boolean and the continuous optimization methods. In reality, we care about the overall energy production of the wind plant, not the instantaneous power production. The overall energy production takes into account all of the directions of incoming wind, as well as a distribution of wind speeds. Although for some orientations, the Boolean optimization performed relatively poorly, for most if compares very well to the more computationally expensive yaw optimization. Overall, we expect that the cases of poor comparison will be balanced out by the other wind conditions, and overall the Boolean yaw optimization will capture most of the gains that are possible from wake steering.

5.3 Computation Time

Except for the cases where the optimal yaw angle of all the turbines in the plant were zero or close to zero, the Boolean optimization was much more computationally efficient than the classic continuous optimization. In general, the Boolean optimization was more than 50 times faster than the continuous optimization and, in some cases, up to 500 times faster. As we mentioned in Section 3, we used the same scaling and convergence criteria for all of the continuous optimization runs. The computation time for any optimization problem is sensitive to these parameters. It is not only possible, but almost guaranteed, that there is some set of scaling and convergence criteria that would allow a specific optimization to find a similar solution faster than we achieved with our scaling. However, finding the best optimizer parameters for a specific optimization problem is often viewed as more of an art than a hard science. There are some general rules that provide an approximation of the appropriate values, but these almost always require several iterations to find a parameter set that works well. Accounting for the possible over-prediction of the Boolean method's computational advantage, it still vastly outperforms the continuous method. Even with the best scaling, we expect the Boolean method to be more than an order of magnitude faster than the continuous optimization. In addition, the Boolean method completely removes the time and experience required to find the appropriate gradient-based optimizer settings.

5.4 Application of Boolean Optimization

There currently is, and always will be, a need to improve the computational efficiency of analysis and optimization studies. For this specific process of quickly optimizing the yaw angles of wind turbines in a power plant for wake steering, we have identified the following reasons to improve computational efficiency as some of the most important.

First is the control optimization of very large wind plants. The computational expense of optimization problems scales poorly with increasing numbers of design variables. Very large wind power plants, or separate wind plants that are close together and have significant wake interaction, should benefit heavily from wake steering through yaw control. However, plants with hundreds or thousands of wind turbines may be limited by the poor scaling of the optimization problem.

Second is the rapid exploration of wind technology potential accounting for wake steering innovations. To determine suitable sites for wind development and help determine how public and private funds should be invested, it is necessary to project how wind energy will be deployed in coming years. Part of this requires understanding how wind energy will be improved through wake steering in existing plants and in plants that will be built while accounting for wake steering. With current computational

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

capabilities, researchers cannot simulate all of the scenarios that are possible and also take into account wake steering. The rapid Boolean optimization process is one step that enables further understanding.

Third is to perform control co-design of a wind power plant. The rapid optimization of yaw angles in a wind plant facilitates coupled wind turbine design, plant layout, and control optimization. As previously stated, the computational expense required to solve optimization problems scales poorly with increasing numbers of design variables without special treatment. When performing control co-design of wind plants, all of the design variables are coupled, resulting in huge numbers of design variables that can easily range up to tens of thousands or more for large wind plants. Problems of this size are infeasible to solve with most current optimization techniques. However, with a fast yaw optimization process, the turbine design and plant layout variables can be decoupled from the yaw angle optimization. If it is fast enough, the yaw optimization can be performed within the plant analysis step, dramatically reducing the number of design variables from tens or hundreds of thousands to fewer than 100.

6 Conclusions

In this paper, we present a novel optimization method to determine the yaw angles of turbines in a wind plant for optimal wake steering. In this method, turbine yaw is defined as Boolean, and the optimization is performed greedily from the most upstream wind turbine to the most downstream. At most, this optimization requires one function call per wind turbine in the plant. We show that with irregular wind turbine layouts, the Boolean optimization performs within about 0.6% of a more traditional, continuous yaw angle definition optimized with a gradient-based algorithm. For a regular grid of wind turbines, or a row of wind turbines in-line with the incoming wind, the Boolean method still achieves most of the power gain as the continuous optimization, with optimal power production within 1.5%–4% of the continuous optimization. The larger discrepancies between the two optimization methods occur in high waking scenarios that have a low probability of occurrence in plants where the layout has been optimized.

In addition to demonstrating the similarity in optimal wind power production achieved by the two different problem formulations, we also showed that the computational expense required to solve the Boolean optimization is much less than that required for the continuous optimization. For any case where the optimal yaw angles were nonzero, the Boolean optimization was around 50–150 times faster than the continuous optimization, with some extreme cases performing about 500 times faster. In addition to the faster computation, our presented Boolean optimization method does not require any scaling of the problem or consideration of the convergence criteria, which removes a large part of the setup time and experience required to solve these optimization problems.

This proposed method greatly simplifies the wind power plant yaw optimization process, achieves plant performance that compares well to more sophisticated methods, and does so at a greatly reduced computational expense. This can be applied to the yaw control optimization of large wind plants, gaining further understanding of wind energy potential and the impacts of wake steering on current and future wind plants, and for coupled turbine design, plant layout, and yaw control optimization. We expect this new method to have wide and immediate impacts in research and in improving wind plant performance.

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

7 Future Work

While there are a huge number of future studies and applications that could expand on this work, we identify and discuss three that we believe could be important.

First, perform further exploration and develop intuition of the best Boolean yaw angles to use in different wind plants. This could involve studying the sensitivity of the optimal Boolean yaw angle to parameters such as average turbine spacing, turbine design, and wind speed. It could also involve a more sophisticated yaw angle selection or optimization in which the Boolean yaw angle is determined by the relative spacing, offset between upstream and downstream wind turbines, or the number of downstream wind turbines that are waked by an upstream turbine. In addition to increasing the wind plant performance, this could further decrease the computational expense of the optimization.

Second, include uncertainty in evaluation of wind plant performance. Past studies have shown that when operating under realistic conditions, in which wind direction and wind speed have significant uncertainty, yaw control strategies should be more conservative, and power gains from wake steering are reduced (Quick et al., 2020; Simley et al., 2020). We expect that when considering uncertainty, the Boolean yaw angle would be affected, and the performance of the plant optimized with the Boolean method would be closer to the continuous method than it was in this paper, in which we assumed wind direction and speed were deterministic.

Third, perform control co-design of wind power plants, and compare the performance and required computational expense of plants that were optimized with a traditional method in which all of the design variables are coupled, to performing our Boolean yaw optimization within a function evaluation, decoupling the turbine design and plant layout variables from the yaw control. We expect minimal differences in optimized performance with significant reductions in required computational expense.

Fourth, refine the optimization methods by including more that one yaw angle. This would add to the computational expense of the algorithm, but could improve the wake steering while still keeping the computational expense relatively low. Multiple passes through the plant with refined yaw angles could further improve the performance.

Code availability. The code and optimized data for this specific paper can be found at:

https://github.com/pjstanle/stanley2021-yaw-optimization

Although any wake model could be used, the FLORIS framework that we used for the results in this paper can be found at:

https://github.com/NREL/floris

Although the optimizer we used in this paper, SNOPT, is commercial, the pyOptSparse framework is open source. It has options to use open-source optimizers as well. The pyOptSparse optimization framework can be found at:

https://github.com/mdolab/pyoptsparse

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

Author contributions. AS contributed with Conceptualization, Investigation, Methodology, Software, Visualization, and Writing – original draft preparation. CB contributed with Conceptualization, Software, and Writing – review & editing. RM contributed with Software and Writing – review & editing. PF contributed with Conceptualization, Funding acquisition, Supervision, Writing – review & editing.

Competing interests. The authors declare that they has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements. This work was authored by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by the U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Wind Energy Technologies Office. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

References

Adaramola, M. and Krogstad, P.-Å.: Experimental investigation of wake effects on wind turbine performance, Renewable Energy, 36, 2078–2086, 2011.

Administration, U. E. I.: Electric power monthly with data for April 2021, Table ES1.A, Tech. rep., U.S. Department of Energy, 2021a.

Administration, U. E. I.: Monthly energy review, Table 7.2a, Tech. Rep. DOE/EIA-0035 (2021/6), Office of Energy Statistics, U.S. Department of Energy, 2021b.

Bastankhah, M. and Porté-Agel, F.: Experimental and theoretical study of wind turbine wakes in yawed conditions, Journal of Fluid Mechanics, 806, 506–541, 2016.

Crespo, A., Herna, J., et al.: Turbulence characteristics in wind-turbine wakes, Journal of wind engineering and industrial aerodynamics, 61, 71–85, 1996.

Dar, Z., Kar, K., Sahni, O., and Chow, J. H.: Windfarm power optimization using yaw angle control, IEEE Transactions on Sustainable Energy, 8, 104–116, 2016.

Dou, B., Qu, T., Lei, L., and Zeng, P.: Optimization of wind turbine yaw angles in a wind farm using a three-dimensional yawed wake model, Energy, 209, 118 415, 2020.

Fleming, P., Churchfield, M., Scholbrock, A., Clifton, A., Schreck, S., Johnson, K., Wright, A., Gebraad, P., Annoni, J., Naughton, B., et al.: Detailed field test of yaw-based wake steering, in: Journal of Physics: Conference Series, vol. 753-5, p. 052003, IOP Publishing, 2016a.

Fleming, P., Annoni, J., Shah, J. J., Wang, L., Ananthan, S., Zhang, Z., Hutchings, K., Wang, P., Chen, W., and Chen, L.: Field test of wake steering at an offshore wind farm, Wind Energy Science, 2, 229–239, 2017.

Fleming, P., King, J., Dykes, K., Simley, E., Roadman, J., Scholbrock, A., Murphy, P., Lundquist, J. K., Moriarty, P., Fleming, K., et al.: Initial results from a field campaign of wake steering applied at a commercial wind farm–Part 1, Wind Energy Science, 4, 273–285, 2019.

Fleming, P. A., Ning, A., Gebraad, P. M., and Dykes, K.: Wind plant system engineering through optimization of layout and yaw control, Wind Energy, 19, 329–344, 2016b.

Gebraad, P., Thomas, J. J., Ning, A., Fleming, P., and Dykes, K.: Maximization of the annual energy production of wind power plants by optimization of layout and yaw-based wake control, Wind Energy, 20, 97–107, 2017.

Gebraad, P. M., Teeuwisse, F., van Wingerden, J.-W., Fleming, P. A., Ruben, S. D., Marden, J. R., and Pao, L. Y.: A data-driven model for wind plant power optimization by yaw control, in: 2014 American Control Conference, pp. 3128–3134, IEEE, 2014.

Gebraad, P. M., Teeuwisse, F., Van Wingerden, J., Fleming, P. A., Ruben, S., Marden, J., and Pao, L.: Wind plant power optimization through yaw control using a parametric model for wake effects—a CFD simulation study, Wind Energy, 19, 95–114, 2016.

Gill, P. E., Murray, W., and Saunders, M. A.: SNOPT: An SQP algorithm for large-scale constrained optimization, SIAM review, 47, 99–131, 2005.

Howland, M. F., Bossuyt, J., Martínez-Tossas, L. A., Meyers, J., and Meneveau, C.: Wake structure in actuator disk models of wind turbines in yaw under uniform inflow conditions, Journal of Renewable and Sustainable Energy, 8, 043 301, 2016.

Jensen, N. O.: A note on wind generator interaction, 1983.

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

Lin, M. and Porté-Agel, F.: Power Maximization and Fatigue-Load Mitigation in a Wind-turbine Array by Active Yaw Control: an LES Study, in: Journal of Physics: Conference Series, vol. 1618-4, p. 042036, IOP Publishing, 2020.

https://doi.org/10.5194/wes-2021-108
Preprint. Discussion started: 22 October 2021

Lyu, Z., Xu, Z., and Martins, J.: Benchmarking optimization algorithms for wing aerodynamic design optimization, in: Proceedings of the 8th International Conference on Computational Fluid Dynamics, Chengdu, Sichuan, China, vol. 11, 2014.

Ning, A. and Petch, D.: Integrated design of downwind land-based wind turbines using analytic gradients, Wind Energy, 19, 2137–2152, 2016.

NREL: FLORIS. Version 2.4, https://github.com/NREL/floris, 2021.

Park, J., Kwon, S., and Law, K. H.: Wind farm power maximization based on a cooperative static game approach, in: Active and Passive Smart Structures and Integrated Systems 2013, vol. 8688, p. 86880R, International Society for Optics and Photonics, 2013.

Quick, J., King, J., King, R. N., Hamlington, P. E., and Dykes, K.: Wake steering optimization under uncertainty, Wind Energy Science, 5, 413–426, 2020.

Rios, L. M. and Sahinidis, N. V.: Derivative-free optimization: a review of algorithms and comparison of software implementations, Journal of Global Optimization, 56, 1247–1293, 2013.

Simley, E., Fleming, P., and King, J.: Design and analysis of a wake steering controller with wind direction variability, Wind Energy Science, 5, 451–468, 2020.

Stokes, L. C. and Warshaw, C.: Renewable energy policy design and framing influence public support in the United States, Nature Energy, 2, 1–6, 2017.

Thomas, J. J. and Ning, A.: A method for reducing multi-modality in the wind farm layout optimization problem, in: Journal of Physics: Conference Series, vol. 1037, p. 042012, IOP Publishing, 2018.

Wiser, R., Rand, J., Seel, J., Beiter, P., Baker, E., Lantz, E., and Gilman, P.: Expert elicitation survey predicts 37% to 49% declines in wind energy costs by 2050, Nature Energy, 6, 555–565, 2021.

Wu, N., Kenway, G., Mader, C. A., Jasa, J., and Martins, J. R.: pyOptSparse: A Python framework for large-scale constrained nonlinear optimization of sparse systems, Journal of Open Source Software, 5, 2564, 2020.

Zingg, D. W., Nemec, M., and Pulliam, T. H.: A comparative evaluation of genetic and gradient-based algorithms applied to aerodynamic optimization, European Journal of Computational Mechanics/Revue Européenne de Mécanique Numérique, 17, 103–126, 2008.