https://doi.org/10.5194/wes-2023-1
Preprint. Discussion started: 9 January 2023
© Author(s) 2023. CC BY 4.0 License.

Open Access
WIND ENERGY SCIENCE DISCUSSIONS

Enabling Control Co-Design of the Next Generation of Wind Plants

Andrew P. J. Stanley, Christopher J. Bay, and Paul Fleming
National Wind Technology Center, National Renewable Energy Laboratory, Golden, CO, 80401, USA
Correspondence: Christopher J. Bay (Christopher.Bay@nrel.gov)

Abstract. Layout design and wake steering through wind plant control are each important and complex components in the design and operation of modern wind plants. They are currently optimized separately, but as more and more wind plants implement wake steering as their primary form of operation, there are increasing needs from industry and regulating bodies to combine the layout and control optimization in a co-design process. However, combining these two optimization problems is currently infeasible due to the excessive number of design variables and the very large solution space. In this paper we present a revolutionary method that enables the coupled optimization of wind plant layout and wake steering with no additional computational expense than a traditional layout optimization. This is accomplished through the development of a geometric relationship between turbines to find an approximate optimal yaw angle, bypassing the need for either a nested or coupled wind plant control optimization. The method we present in this paper provides a significant and immediate improvement to wind plant design by enabling the co-design of turbine layout and yaw control for wake steering. A small co-designed plant shown in this paper produces 0.8% more energy than its sequentially designed counterpart, and we expect larger comparative gains for larger plants with more turbines. This additional energy production comes with no additional infrastructure, turbine hardware, or control software; it is a free consequence of optimizing the turbine layout and yaw control together, resulting in millions of dollars of additional revenue for the wind plants of the future.

1 Introduction

The design and layout of wind turbines within a wind plant is a highly complex problem where the wind plant developer must weigh numerous competing goals and constraints against each other. An example objective is maximizing the expected energy production of the plant while minimizing the cost to build. The problem also includes constraints on layouts which could include specified boundaries, turbine spacing requirements, grids or other layout regularity, setback from shipping lanes or structures, and sea floor or terrain based constraints.

Predicting the impact of wind turbine wakes on total wind plant production plays a key role in wind plant design. Within a wind plant, wind turbines interact with each other through the wakes that they produce while extracting energy from the passing flow (Sanderse et al., 2011). These wakes have reduced wind speed which limits the energy that is available to downstream turbines in the plant; additionally, these wakes have higher turbulence than the ambient flow which increases loads and is detrimental for structural reliability. Negative wake impacts can be mitigated in a wind plant's design phase, as well as during plant operation.

https://doi.org/10.5194/wes-2023-1
Preprint. Discussion started: 9 January 2023
© Author(s) 2023. CC BY 4.0 License.

The primary way to minimize wake interactions during the plant design stage is through turbine layout optimization, often referred to as micrositing, which is an important step for both offshore and onshore wind plants (Hou et al., 2019; Balasubramanian et al., 2020). Through layout optimization, wake interactions can be minimized for the wind resource, atmospheric conditions, turbine design, and constraints unique to a specific site. As was already alluded to, wind plant layout optimization is a notoriously challenging problem because of the large number of interacting variables, complexity of the required models, and the complexity of the design space.

During operation, plant control can be used to reduce wake interactions. One plant-level control strategy is yaw misalignment for wake steering. A wind turbine whose yaw angle is misaligned with the incoming wind will produce a wake that is deflected compared to an unyawed turbine. This phenomenon can be exploited to intentionally steer wakes away from downstream turbines in the wind plant. Although a turbine with some yaw misalignment to the incoming wind will suffer reduced power production and increased loading, the benefits of wake steering can result in a net improvement for the entire plant. This has been demonstrated with several different fidelities of wind plant simulations (Jiménez et al., 2010; Gebraad et al., 2017; Martínez-Tossas et al., 2021) as well as with wind tunnel experiments (Campagnolo et al., 2020). Because of these promising simulations and experiments, wake steering is seeing increased momentum toward growing adoption at existing sites. There have now been several demonstrations of wake steering implemented at commercial wind plants (Fleming et al., 2017, 2019, 2020; Simley et al., 2021; Howland et al., 2022). There have also been several announcements of commercial implementations of wind plant control, either provided by the wind turbine original equipment manufacturers or consultants.

An enormous opportunity for improved wind plant performance presents itself by simultaneously optimizing wind plant layout and turbine yaw angles. Generally this process is called control co-design, which means to account for aspects of system control throughout the entire design process (Garcia-Sanz, 2019). Specifically, control co-design can be leveraged to maximize the capture of spatially varying wind resources, such as offshore sites with wind speed correlated to the distance from shore, or complex terrain where higher wind speeds can exist on higher elevation topologies. Control co-design would allow the potential for wake loss mitigation in the layout process. Control co-design can also optimally make use of the available space in lease areas where lease fees are significant, or reduce installation costs by condensing turbines into shallower offshore regions. Coupled with other design parameters and constraints, the possible benefits of control co-design are numerous. Currently, this possibility is severely limited in wind plants by the large number of design variables required to fully couple wind plant layout and yaw control optimization. In its most basic formulation, optimizing wind plant layout and yaw angles requires two design variables for every turbine (one for both the x and y coordinate), and one design variable per turbine *per wind condition* for the yaw angles. For an average sized wind plant ($\sim$ tens of turbines), the fully coupled problem can easily reach thousands or tens of thousands of coupled design variables. These optimizations scale very poorly as the number of wind turbines increases, as shown in Fig. 1.

In practice, this description of the scaling issue of the fully-coupled optimization understates the problem, as real-world wind plant design must include many further design constraints, some of which were already mentioned. Additionally, in practice, wind plant layout optimization is an optimization performed over an already-existing developer-specific tool chain. This makes

https://doi.org/10.5194/wes-2023-1
Preprint. Discussion started: 9 January 2023
© Author(s) 2023. CC BY 4.0 License.

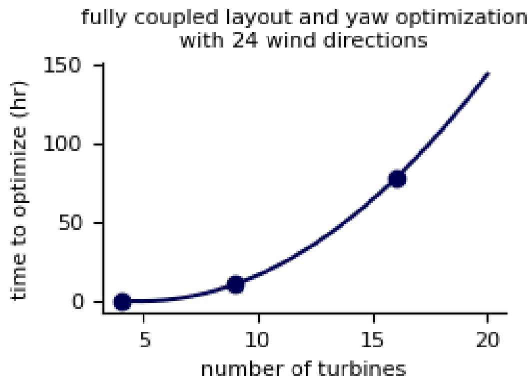

Figure 1. The time to solve the coupled wind turbine layout and yaw control optimization problem as a function of the number of turbines. These optimizations were performed with the gradient-based optimizer SNOPT, with the 24 wind direction bins.

co-design of the layout and yaw control even more prohibitive because practical implementation requires a control optimization nested within the existing plant optimization.

In this paper, we present a novel method to determine yaw angles for wake steering that will revolutionize wind plant optimization. This new method defines turbine yaw angles deterministically from the layout of a wind plant. Having the yaw angles implicitly defined by the turbine layout enables the fully coupled optimization of wind plant layout and turbine yaw angles for wake steering with no additional design variables than the isolated wind plant layout optimization problem. A small example plant presented later in this paper produced 0.8% more energy when the layout and yaw angles were optimized together compared to when the layout and yaw angles were optimized sequentially. To put this in perspective, 1 MW of wind capacity generates annual revenue on the order of $100,000. Therefore, a 0.8% increase in performance equates to an additional $800 per MW, or $800,000 per GW each year. This coupled, yet efficient optimization is frequently being requested by industry and is extremely relevant for the next generation of offshore and onshore wind plants for which it will be important to maximize the wind generation in a limited space. Expensive and limited lease areas (Friedman), increasingly strict siting regulations (Mai et al., 2021), and improved technology enabling larger turbines (Enevoldsen and Xydis, 2019) will drive the wind plants of the future to have turbines packed close together relative to the rotor diameter. This marked improvement in wind plant design can have significant implications in near-term renewable energy policy and wind deployment throughout the world, helping accelerate deployment and reach near- and mid-term wind energy goals.

2 Geometric Yaw Relationship

When optimizing the yaw offset angles in a wind plant, there are many different combinations of turbine yaw angles that result in almost identical plant performance. We determined that a sufficiently-optimal yaw angle for any individual wind turbine can be calculated as a function of the streamwise and cross-stream distance to its nearest downstream waked turbine, as shown in Figures 2A and 2B. Figure 2A shows a group of 5 turbines with the wind coming from the left. To determine the yaw angle

https://doi.org/10.5194/wes-2023-1
Preprint. Discussion started: 9 January 2023
© Author(s) 2023. CC BY 4.0 License.
eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

of the yellow turbine using our geometric yaw relationship, it is necessary to calculate the streamwise distance (dx) and the cross-stream distance (dy) to the nearest waked turbine shown in purple. The black circles represent the other turbines in this cluster. Notice that there are two turbines closer to the yellow turbine, but these are not waked and therefore do not affect the yaw angle of the yellow turbine. To determine if a turbine was waked, we assumed a wake radius $r_{\text{wake}} = 0.1x + r_{\text{turbine}}$, where r_{wake} is the radius of the wake, x is the streamwise distance downstream of the waking turbine, and r_{turbine} is the radius of the waking turbine. Figure 2B shows the same group of five turbines, but with the wind coming from the upper left corner. For this wind direction, the nearest waked turbine is different than in 2A.

To understand the relationship between optimal turbine yaw angles and their position relative to the nearest downstream waked turbine, we optimized the yaw for many different wind plants, including randomly generated layouts with different numbers of turbines, average turbine spacings, and wind speeds, as well as regular grid layouts with different numbers of rows and columns, turbine spacings, grid rotations, and wind speeds. These yaw angles were optimized with the gradient-based optimizer SNOPT (Gill et al., 2005, 2018), with the objective to maximize plant power modeled with FLORIS (National Renewable Energy Laboratory, 2022), a controls-focused wind plant simulation software incorporating steady-state engineering wake models with wake deflection modeling capabilities. The result was over 100,000 optimized turbine yaw angles. Figure 2C shows these optimal yaw angles as a function of position of the yawed turbine relative to its nearest waked downstream turbine, normalized by the turbine rotor diameter. A single point in this figure represents the yaw angle of a single turbine, represented by the color, as a function of the distance to the nearest downstream waked turbine, indicated by the point's position on the plot. A clear pattern emerges from Figure 2C. There is a divide between positive and negative yaw angles depending on whether the cross-stream distance to the nearest waked turbine is positive or negative. Additionally we can see that the turbine is only yawed if the cross-stream distance to the nearest downstream waked turbine is around one rotor diameter or less. Outside of that range, for the most part the upstream turbine has an optimized yaw angle near zero. Observing this pattern, we created the geometric yaw relationship shown in 2D, which can be used to instantly determine a near-optimal yaw angle for a wind turbine as a function of its location relative to the turbines around it. The specific relationship is a 1-dimensional gradient starting at the upstream turbine with a value of 30 degrees, and linearly decreasing to 0 degrees at 25 rotor diameters downstream. The sign of the yaw angle is determined by the lateral placement of the downstream turbine, as shown in 2D.

3 Results

In this section, we present two examples of implementations of geometric yaw in the wind plant layout optimization problem. Like many other scenarios, the scenarios presented in this paper perform better and look significantly different when the layout and yaw control are optimized together. Additionally, the use of our geometric yaw relationship during turbine layout optimization allowed these examples to be run on laptop machine on a single processor, which is infeasible with currently existing methods. Note that we don't compare computational time between the layout-only and layout-plus-control optimization. The time required to compute the necessary geometries is negligible compared to the full objective function, meaning that the difference in computation time between layouts optimized with and without geometric yaw is trivial. For each example

https://doi.org/10.5194/wes-2023-1
Preprint. Discussion started: 9 January 2023
© Author(s) 2023. CC BY 4.0 License.

Open Access
WIND ENERGY SCIENCE
DISCUSSIONS

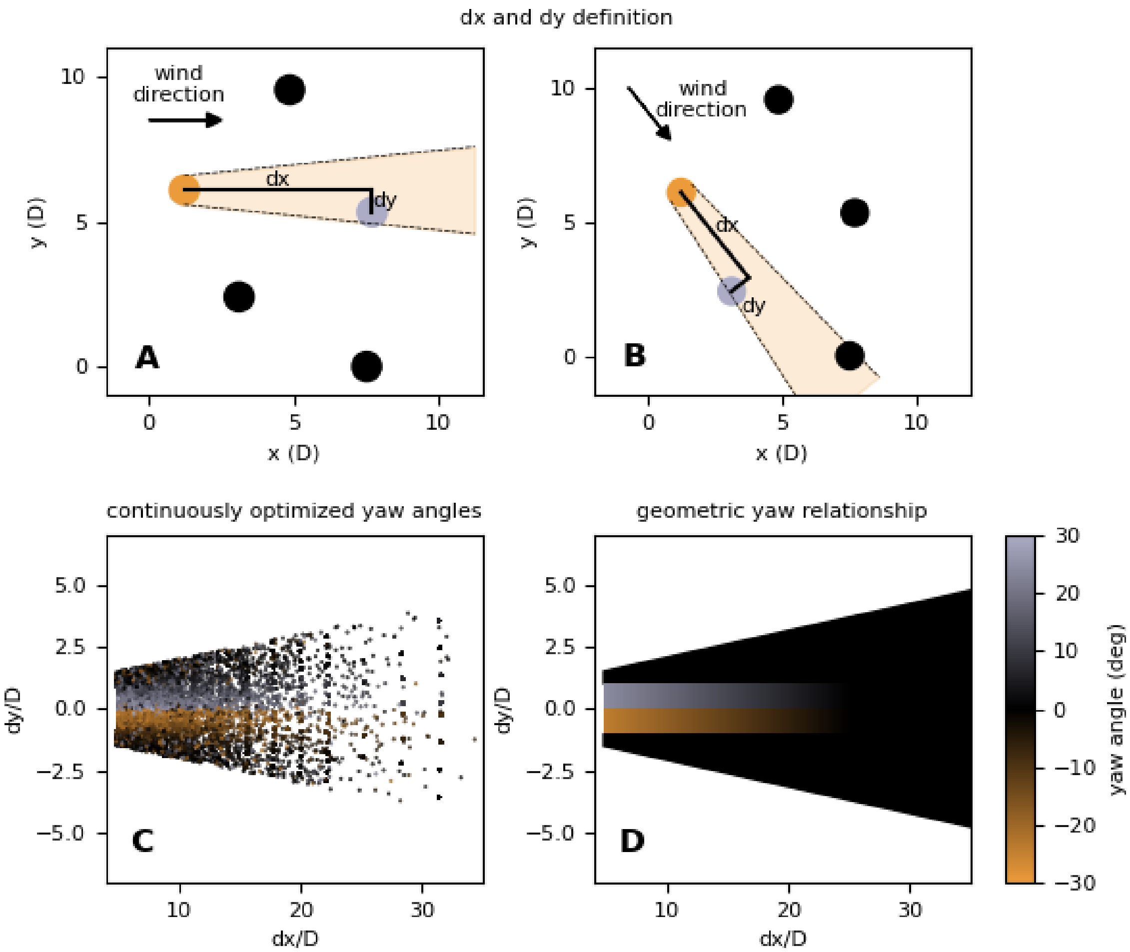

Figure 2. A description of the geometric yaw relationship presented in this paper. Figures 2A and 2B show how the streamwise (dx) and cross-stream (dy) distances to the nearest downstream waked turbine are defined for two different wind directions. Figure 2C shows the optimized yaw angles for over 100,000 individual turbines optimized continuously in a variety of wind plants with different numbers of turbines, layouts, turbine spacings, and wind speeds. These yaw angles are shown as a function of the streamwise and cross-stream distance to the nearest downstream waked turbine of the yawed turbine. Figure 2D shows the geometric yaw relationship that we defined through observation of the pattern that emerges in Figure 2C, which can be used to immediately determine a near-optimal yaw angle of any turbine in a wind plant.

https://doi.org/10.5194/wes-2023-1
Preprint. Discussion started: 9 January 2023
© Author(s) 2023. CC BY 4.0 License.

in this section we used the scipy (Virtanen et al., 2020) SLSQP gradient-based optimzer within the pyOptSparse optimization framework (Wu et al., 2020). To model the plant performance, we used FLORIS 3.1 as was discussed previously. For both of the examples we compare two different wind plant layouts and how they perform. The first is a layout that was optimized assuming no yaw control. After the layout was optimized, the turbine locations were fixed and the yaw angles were optimized continuously to determine the final plant performance. The second layout was optimized using our geometric yaw relationship to define the yaw angles during the layout optimization. After the layout was optimized and fixed, one final continuous yaw optimization was performed to determine the final yaw angles and plant performance.

3.1 One-Dimensional Plant

The first example we present is a simple one-dimensional wind plant. Although this problem would not occur in the real world, it is valuable to demonstrate the power of coupled layout and yaw optimization. In this example, 16 turbines were arranged in a straight line with constant wind in-line with the row of turbines. The objective was to maximize the power density of the array, which for these purposes was defined as the total power divided by the length of the row of turbines. The spacings between adjacent turbines were the design variables during the layout optimization, meaning that there were a total of 15 design variables (the number of turbines minus one). Figure 3 shows the results of this optimization. Figures 3A.1 and 3A.2 show the layout that was optimized assuming no yaw control, while Figures 3B.1 and 3B.2 show the layout that was optimized using the geometric yaw relationship. Figures 3A.1 and 3B.1 show the plant and performance for each optimized layout without yaw control. Figures 3A.2 and 3B.2 show these same layouts, but with the final optimized yaw angles.

This example clearly demonstrates two important principles. First, the layout optimized without wake steering, which has the turbines spaced relatively far apart (Figures 3A.1 and 3A.2) is very different than the layout optimized with wake steering, which has the turbines much closer together (Figures 3B.1 and 3B.2). Visually we can see that the difference between the layouts is significant, indicating that including wake steering during the layout optimization will lead to a different solution. Second, the layout optimized with geometric yaw outperforms the layout optimized without geometric yaw by over 7% when the plant is operated with wake steering (Figure 3B.2 compared to Figure 3A.2). However, when the plant is operated without yaw control, the layout optimized without geometric yaw outperforms the one optimized with geometric yaw by 2.6% (Figure 3A.1 compared to Figure 3B.1). From this observation, we can conclude that the layout should be optimized with the yaw control scheme that will be used during plant operation. Plants that will be operated with wake steering will benefit greatly from optimizing the turbine layout with geometric yaw.

3.2 Gaussian Hill Spatially Varying Inflow

The second example is a more realistic two-dimensional layout optimization with a full distribution of wind directions and spatially varying freestream wind speeds across the domain. In this example we optimized the layout of a wind plant with sixteen turbines, with the objective to maximize the annual energy production (AEP) of the plant. The turbines were constrained within a 2-by-2 kilometer square, and had a minimum spacing constraint of 2 turbine rotor diameters. We used a bimodal wind rose shown in Figure 4A divided into 72 discrete bins. From each wind direction we assumed a constant wind speed, indicated

https://doi.org/10.5194/wes-2023-1
Preprint. Discussion started: 9 January 2023
© Author(s) 2023. CC BY 4.0 License.

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

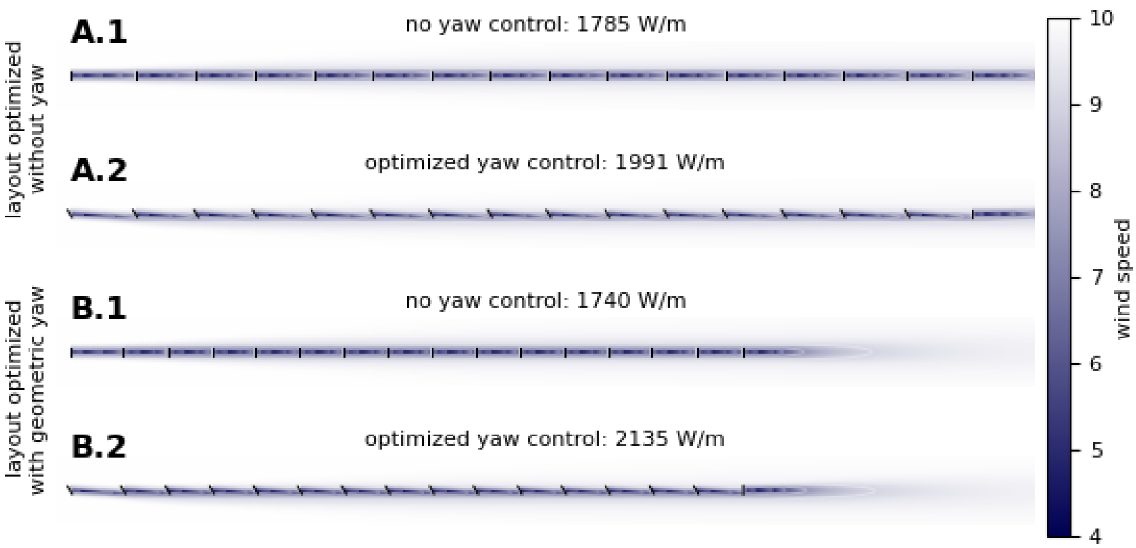

Figure 3. A simple one-dimensional optimization of a 16 turbine wind plant comparing the performance of the layout optimized without yaw control and the layout optimized with geometric yaw. Figures 3A.1 and 3A.2 show the layout that was optimized assuming no yaw control. Figures 3B.1 and 3B.2 show the layout that was optimized using the geometric yaw relationship. Figures 3A.1 and 3B.1 show the plant and performance for each optimized layout without yaw control. Figures 3A.2 and 3B.2 show these same layouts, but with the final optimized yaw angles.

by the color bar in Figure 4A. In addition to the full wind rose, we assumed there was a spatially varying wind speed throughout the domain for each wind direction. This was accomplished by applying a Gaussian wind speed multiplier to the domain with a standard deviation of 600 meters in each direction, which provided a maximum wind speed increase in the wind speed multiplier of 0.4 at the origin. This wind speed variation was meant to approximately simulate the spatial variation in wind speeds caused by a hill, including the speedup and wind shadow regions, so we also applied a penalty behind the hill to capture the wind shadow. This was done by applying a second Gaussian distribution 400 meters directly behind the origin in line with the wind direction. This negative Gaussian distribution again had a standard deviation of 600 meters and provided a maximum decrease in the wind speed multiplier of 0.2. The interaction of these two Gaussian distribution is a maximum wind speed multiplier of about 1.25 and a minimum slightly less than 1.0. The resultant wind speed multiplier distribution for wind coming directly from the left is shown in Figure 4B (note that this only visualizes the speedup/slowdown for one direction; the location of the highest speedup and the wind shadow change with the wind direction).

With the scenario fully defined, we optimized the plant layout both while assuming no yaw control and while using the geometric yaw relationship. Because of the large amount of local minima known to exist in the wind plant layout optimization problem, and because gradient-based optimizers are known to converge to local minima without full exploration of the design space, we randomly initialized 50 turbine layouts and performed each optimization (without yaw and with geometric yaw) from

https://doi.org/10.5194/wes-2023-1
Preprint. Discussion started: 9 January 2023
© Author(s) 2023. CC BY 4.0 License.

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

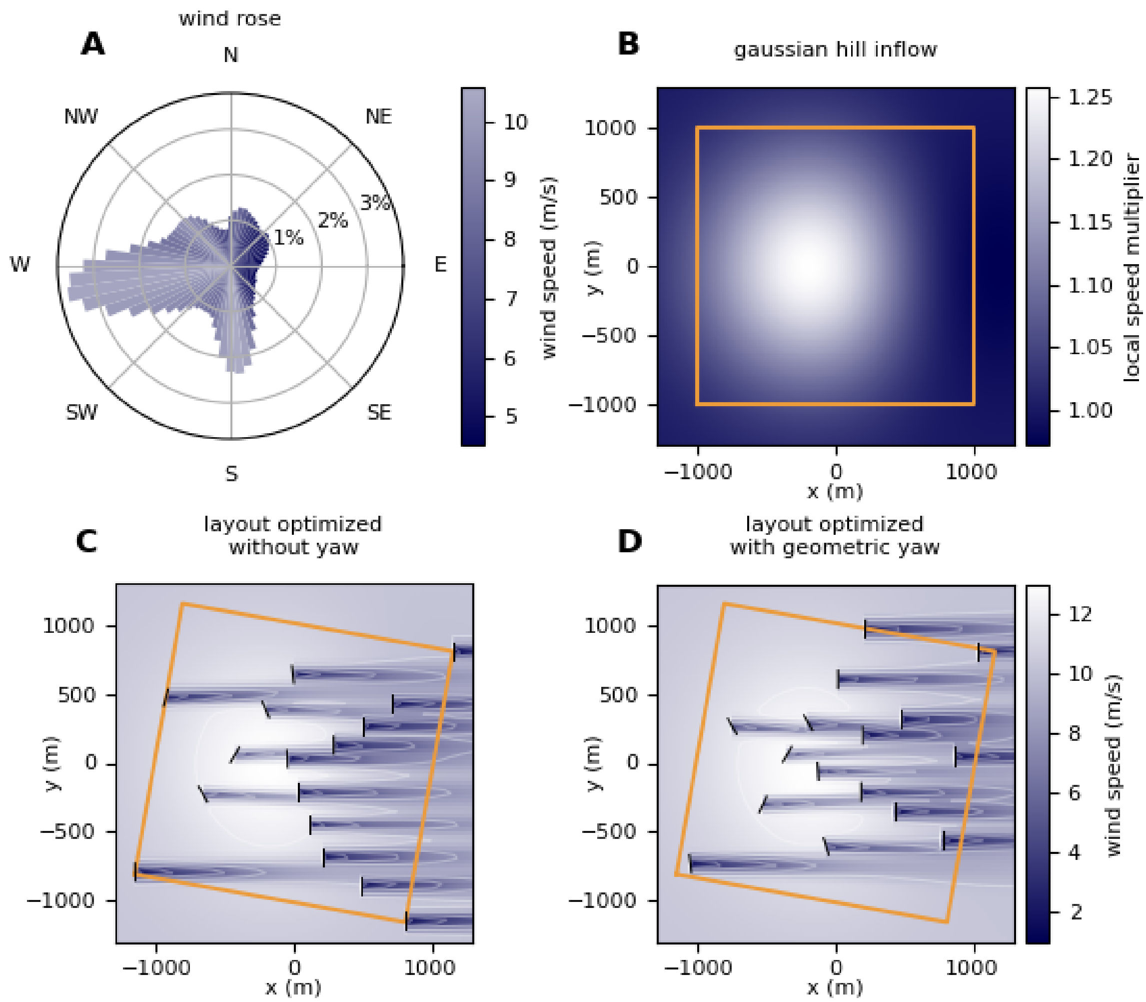

Figure 4. A turbine layout optimization of a 16 turbine wind plant comparing the performance of the layout optimized without yaw control and the layout optimized with geometric yaw. In this figure, there is a spatially varying wind speed multiplier applied over the domain. Figure 4A shows the wind probability rose used in this optimization. The wind directions were divided into 72 5-degree bins, with directionally averaged wind speeds indicated by the color bar. Figure 4B shows the Gaussian wind speed multiplier applied to the domain for one wind direction (from the left). In our simulation, the wind shadow rotated behind the Gaussian peak for each wind direction. Figures 4C and 4D show the optimal turbine layouts, and optimal turbine yaw angles and flow fields for the dominant wind direction. Figure 4C shows the layout optimized without yaw, and Figure 4D shows the layout optimized with geometric yaw.

https://doi.org/10.5194/wes-2023-1
Preprint. Discussion started: 9 January 2023
© Author(s) 2023. CC BY 4.0 License.

each of the 50 starting layouts. The best-performing layout from each method was selected as the final plant layout, to which one final yaw control optimization was performed to evaluate the final plant performance. With currently existing methods, this coupled turbine layout and yaw angle optimization problem would have needed 1,184 fully coupled design variables, two variables for each wind turbine to define the locations and one variable per turbine for each of the 72 wind directions to define the yaw angles. With our geometric yaw relationship we reduced that down to just 32 variables, the N×2 required for each turbine to define the layout, which allowed us to perform the optimization on our local machine with finite-difference gradients.

The layout optimized with geometric yaw produced **0.8% higher** AEP that the layout optimized without yaw. As previously discussed, because wind plants are enormous investments, a performance gain around 0.8% can easily equate to hundreds of thousands or millions of dollars annually depending on the plant capacity. This particular performance improvement is even more impressive in that it does not require any additional parts or technology, it simply involves building the turbines in better locations which were not found before this geometric yaw relationship. The optimized turbine locations, and the associated yaw angles and flow field for the dominant wind direction are shown in Figures 4C and 4D. Figure 4C shows the layout that was optimized without turbine yaw, and Figure 4 shows the layout that was optimized with the geometric yaw relationship. In these figures the wind plant boundary is represented by the orange squares, which appear rotated because the dominant wind direction is from 260 degrees, and the wind in this image is displayed coming from the left. Notice the extremely different layouts obtained with the two different optimization methods. Because the layout in Figure 4D was optimized with geometric yaw, the optimizer was able to place turbines closer together near the peak in the Gaussian wind speed multiplier, taking advantage of the wake steering to reduce wake interactions between nearby turbines. On the other hand, the layout in Figure 4C was optimized without yaw. This meant that the optimizer did not take as much advantage of the wind speed multiplier, and instead opted to spread turbines perpendicular to the dominant wind direction as displayed in the figure. This more complex example reiterates the conclusions found in the one-dimensional example, that optimizing the layout concurrently with wake steering leads to different optimal layouts and significant performance improvements.

4 Conclusions

In this paper, we presented a geometric yaw relationship that can be used to determine sufficiently-optimal turbine yaw angles for wake steering as a function of the layout of the wind wind plant. This method, or any improvement on the specific relationship presented in this paper, can be used to solve the coupled wind plant layout and yaw control optimization problem in a computationally efficient manner, and can find layouts that perform significantly better than layouts that are optimized without yaw. In Section 3.2, using geometric yaw we obtained a plant layout that performed 0.8% better than a layout optimized assuming no yaw, with no difference in the number of function calls or computation time required to optimize. In itself, 0.8% is already a non-negligible improvement in plant performance, however we expect much higher improvements in larger plants which have more wake interactions. The geometric yaw relationship presented in this paper enables fully coupled wind plant layout and yaw control optimization with no added expense compared to the regular wind plant layout optimization problem, and can revolutionize how wind plant layout optimization is approached by researchers and wind plant developers alike.

https://doi.org/10.5194/wes-2023-1
Preprint. Discussion started: 9 January 2023
© Author(s) 2023. CC BY 4.0 License.

Code availability. https://github.com/pjstanle/GeometricYaw/tree/paper/initial_submission

Author contributions. AS was responsible for Conceptualization, Investigation, Methodology, Software, Visualization, and Writing – original draft preparation. CB was responsible for Conceptualization, Investigation, Methodology, Software, and Writing – review and editing. PF was responsible for Funding acquisition, Project administration, Resources, Supervision, and Writing – review and editing.

Competing interests. Paul Fleming is a member of the editorial board of *Wind Energy Science*.

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 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-2023-1
Preprint. Discussion started: 9 January 2023
© Author(s) 2023. CC BY 4.0 License.

References

Balasubramanian, K., Thanikanti, S. B., Subramaniam, U., Sudhakar, N., and Sichilalu, S.: A novel review on optimization techniques used in wind farm modelling, Renewable Energy Focus, 35, 84–96, 2020.

Campagnolo, F., Weber, R., Schreiber, J., and Bottasso, C. L.: Wind tunnel testing of wake steering with dynamic wind direction changes, Wind Energy Science, 5, 1273–1295, 2020.

Enevoldsen, P. and Xydis, G.: Examining the trends of 35 years growth of key wind turbine components, Energy for sustainable development, 50, 18–26, 2019.

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., King, J., Simley, E., Roadman, J., Scholbrock, A., Murphy, P., Lundquist, J. K., Moriarty, P., Fleming, K., van Dam, J., et al.: Continued results from a field campaign of wake steering applied at a commercial wind farm–Part 2, Wind Energy Science, 5, 945–958, 2020.

Friedman, L.: Sale of Leases for Wind Farms Off New York Raises More Than $4 Billion, The New York Times, https://www.nytimes.com/2022/02/25/climate/new-york-offshore-wind-auction.html.

Garcia-Sanz, M.: Control Co-Design: an engineering game changer, Advanced Control for Applications: Engineering and Industrial Systems, 1, e18, 2019.

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.

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

Gill, P. E., Murray, W., Saunders, M. A., and Wong, E.: User's guide for SNOPT 7.7: Software for large-scale nonlinear programming, Center for Computational Mathematics Report CCoM, 15, 2018.

Hou, P., Zhu, J., Ma, K., Yang, G., Hu, W., and Chen, Z.: A review of offshore wind farm layout optimization and electrical system design methods, Journal of Modern Power Systems and Clean Energy, 7, 975–986, 2019.

Howland, M. F., Quesada, J. B., Martinez, J. J. P., Larrañaga, F. P., Yadav, N., Chawla, J. S., Sivaram, V., and Dabiri, J. O.: Collective wind farm operation based on a predictive model increases utility-scale energy production, arXiv preprint arXiv:2202.06683, 2022.

Jiménez, Á., Crespo, A., and Migoya, E.: Application of a LES technique to characterize the wake deflection of a wind turbine in yaw, Wind energy, 13, 559–572, 2010.

Mai, T., Lopez, A., Mowers, M., and Lantz, E.: Interactions of wind energy project siting, wind resource potential, and the evolution of the US power system, Energy, 223, 119 998, 2021.

Martínez-Tossas, L. A., King, J., Quon, E., Bay, C. J., Mudafort, R., Hamilton, N., Howland, M. F., and Fleming, P. A.: The curled wake model: a three-dimensional and extremely fast steady-state wake solver for wind plant flows, Wind Energy Science, 6, 555–570, 2021.

National Renewable Energy Laboratory: FLORIS Version 3.1, https://github.com/NREL/floris/releases/tag/v3.1, 2022.

Sanderse, B., Van der Pijl, S., and Koren, B.: Review of computational fluid dynamics for wind turbine wake aerodynamics, Wind energy, 14, 799–819, 2011.

https://doi.org/10.5194/wes-2023-1
Preprint. Discussion started: 9 January 2023
© Author(s) 2023. CC BY 4.0 License.

Simley, E., Fleming, P., Girard, N., Alloin, L., Godefroy, E., and Duc, T.: Results from a wake-steering experiment at a commercial wind plant: investigating the wind speed dependence of wake-steering performance, Wind Energy Science, 6, 1427–1453, 2021.

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors: SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python, Nature Methods, 17, 261–272, https://doi.org/10.1038/s41592-019-0686-2, 2020.

Wu, N., Kenway, G., Mader, C. A., Jasa, J., and Martins, J. R. R. A.: pyOptSparse: A Python framework for large-scale constrained nonlinear optimization of sparse systems, Journal of Open Source Software, 5, 2564, https://doi.org/10.21105/joss.02564, 2020.