



Gradient-Based Optimization of Wind Farms with Different Turbine Heights

Preprint

Andrew P. J. Stanley, Jared Thomas,
and Andrew Ning
Brigham Young University

Jennifer Annoni, Katherine Dykes,
and Paul Fleming
National Renewable Energy Laboratory

*Presented at the American Institute of Aeronautics and
Astronautics SciTech 2017
Dallas, Texas
January 16–20, 2017*

**NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC**

This report is available at no cost from the National Renewable Energy
Laboratory (NREL) at www.nrel.gov/publications.

Conference Paper
NREL/CP-5000-67661
April 2017

Contract No. DE-AC36-08GO28308

NOTICE

The submitted manuscript has been offered by an employee of the Alliance for Sustainable Energy, LLC (Alliance), a contractor of the US Government under Contract No. DE-AC36-08GO28308. Accordingly, the US Government and Alliance retain a nonexclusive royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for US Government purposes.

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

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Available electronically at SciTech Connect <http://www.osti.gov/scitech>

Available for a processing fee to U.S. Department of Energy and its contractors, in paper, from:

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
OSTI <http://www.osti.gov>
Phone: 865.576.8401
Fax: 865.576.5728
Email: reports@osti.gov

Available for sale to the public, in paper, from:

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Road
Alexandria, VA 22312
NTIS <http://www.ntis.gov>
Phone: 800.553.6847 or 703.605.6000
Fax: 703.605.6900
Email: orders@ntis.gov

Cover Photos by Dennis Schroeder: (left to right) NREL 26173, NREL 18302, NREL 19758, NREL 29642, NREL 19795.

NREL prints on paper that contains recycled content.

Gradient-Based Optimization of Wind Farms with Different Turbine Heights

Andrew P. J. Stanley*, Jared Thomas,[†] and Andrew Ning[‡]

Brigham Young University, Provo, Utah 84602

Jennifer Annoni[§], Katherine Dykes,[¶] and Paul Fleming[¶]

National Renewable Energy Laboratory, Golden, Colorado 80401

Turbine wakes reduce power production in a wind farm. Current wind farms are generally built with turbines that are all the same height, but if wind farms included turbines with different tower heights, the cost of energy (COE) may be reduced. We used gradient-based optimization to demonstrate a method to optimize wind farms with varied hub heights. Our study includes a modified version of the FLORIS wake model that accommodates three-dimensional wakes integrated with a tower structural model. Our purpose was to design a process to minimize the COE of a wind farm through layout optimization and varying turbine hub heights. Results indicate that when a farm is optimized for layout and height with two separate height groups, COE can be lowered by as much as 5%-9%, compared to a similar layout and height optimization where all the towers are the same. The COE has the best improvement in farms with high turbine density and a low wind shear exponent.

I. Introduction

As wind turbines extract energy from the air and convert it to power, an area of reduced wind speed is formed behind each wind turbine known as a wake. Because the air in a wake has less momentum, a wind turbine in a wake cannot extract as much energy and therefore produces less power. Several solutions have been developed to help remedy this problem, including layout optimization of the wind farm¹⁻³ and rotor yaw control.^{4,5} In general, wind farms are built with one turbine type and height, and layout optimization studies only analyze wind farms with identical turbines. Including more than one turbine height in the same wind farm could decrease wake interference even further and result in higher energy production.

Several studies have explored the use of different turbine heights in the same wind farm. Chen et al. used a genetic algorithm to optimize a wind farm layout of 25 turbines by changing the position and height of each turbine between two predefined heights. They found that the power increased by as much as 13.53% and the cost per unit of energy produced decreased 0.37%.⁶ Hazra et al. used a particle swarm method to optimize a wind farm, in which the turbine height and rotor radius are both design variables. In a 10-turbine wind farm, they found a 12.8% reduction in the cost of power.⁷ Studies such as these show promising results for multiple hub-height wind farms by increasing energy production and decreasing the cost of energy; however, they are limited to small problems. Both only analyze small wind farms and include few design variables: turbine position, tower height, and for the Hazra study, rotor size. For a farm of 25 turbines that has 3-4 variables per wind turbine, the optimization already has 75-100 design variables. Larger wind farms will have many more variables, making the optimization problem more computationally expensive. Because these studies use gradient-free optimization, having more design variables greatly increases the computational expense, making it difficult to impossible to use this method for large problems.

*Ph.D. Student, Brigham Young University Department of Mechanical Engineering

[†]M.S. Student, Brigham Young University Department of Mechanical Engineering

[‡]Assistant Professor, Brigham Young University Department of Mechanical Engineering, AIAA Senior Member

[§]Postdoctoral Researcher, National Wind Technology Center

[¶]Senior Engineer, National Wind Technology Center

Introducing varying hub height to layout and control optimization adds complexity and computational expense to the problem. The number of design variables increases by up to the number of turbines in the wind farm; one for each tower height. Additionally, a wake model must be developed or modified to operate in three dimensions, and a structural model for the tower must be added to account for potential failure as the height changes. The tower model adds more design variables for the diameter and thickness of the tower. Finally, the free stream wind speed from a given direction is no longer the same for all turbines in the farm as wind speed varies with height because of wind shear.

The previous studies addressed some of these problems and made optimistic conclusions, especially for small wind farms. They accounted for the change in wind speed at different heights and optimized for height and position. Hazra et al. also included rotor diameter as a design variable in their optimization. However, until now gradient-based algorithms have not been used in three-dimensional wind farm optimization. Gradient-based optimization is faster than gradient-free methods and is necessary for optimizing large wind farms including many design variables, such as yaw control coupled with the variables mentioned above. When yaw control is added to the optimization, thousands of design variables can be added, because each turbine must be optimized for each wind direction in consideration. A 50-turbine wind farm with 72 wind directions could have more than 4000 design variables. Optimization of this size is impractical with a gradient-free approach because of the huge computational expense. In this paper we do not couple yaw control into the layout/turbine design problem, but have built a gradient-based framework that will allow for subsequent coupled layout, turbine, and yaw design problems.

We will use gradient-based optimization to allow for the efficient optimization of large wind farms with many design variables. Specifically, we will optimize wind farms with different hub heights, and demonstrate the gains of wind farms with multiple hub heights compared to those with turbines at an identical height. Combining multiple hub heights in wind farms while continuing to optimize their layout may have significant impact on the cost of energy (COE) in wind farms.

II. Methodology

In this section, we describe the model used to predict the COE of a wind farm. First, the wake model is discussed, which is needed to calculate the wind speed at any point in the wind farm. Next, we discuss the annual energy production (AEP) and how it is calculated. We will then consider structural calculations made along the length of the tower that are important as constraints in our optimization. Finally, we introduce our cost model. Each of these components was used in our optimization.

A. Wake Model

To calculate the effective wind speed at each turbine, we used the FLORIS wake model presented by Gebraad et al.⁴ The FLORIS wake model is derived from the Jensen model,⁸ but rather than use one speed to describe the wind across the wake, three separate zones are defined, each with a different expansion and decay rate. A simple overlap ratio is used between zones to define the total effective wind speed at each turbine. Figure 1 shows the three separate wake zones, as well as their overlap on a rotor. Recent work has improved FLORIS to provide a smooth response and analytic gradients.⁹ These improvements enable solutions to be found for large optimization problems, and help to achieve reliable answers. Without analytic gradients, finite-difference gradients must be used, which often experience numerical difficulties, and do not scale well. Large problems are very computationally expensive with finite difference gradients. Because this wake model was designed to describe the wake in the horizontal plane, it was modified to calculate the effective wind speed at any point in three-dimensional (3-D) space. We assume that the wake is axisymmetric, such that any cross section is circular. Additionally, we continue the assumption from the original FLORIS model that the wake center neither ascends or descends, but remains at the same height from which it originated. FLORIS uses precomputed data, unique to the turbine model used, for the C_P and C_T curves that are used in the turbine power calculation.

A real wake may move in the vertical plane and may not maintain a perfectly circular cross section. To learn whether or not the assumptions we made were reasonable, we compared the model results to Simulator fOr Wind Farm Applications (SOWFA). SOWFA, a high-fidelity large eddy simulation tool that was developed at the National Renewable Energy Laboratory (NREL) for wind farm studies, is based on OpenFOAM and is coupled with NREL's FAST modeling tool.¹⁰⁻¹² SOWFA has been used extensively




Figure 1. The FLORIS wake model. The model has three zones with varying diameters, $D_{w,q}$, that depend on tuning parameters k_e and $m_{e,q}$. The effective hub velocity is computed using the overlap ratio, A_q^{OL} , of the total rotor-swept area to the part of the rotor-swept area overlapping each wake zone, respectively.

in previous wind farm control studies.^{4,13,14} The solver uses an actuator line model, or actuator disk model, coupled with FAST to study turbines in the atmospheric boundary layer. SOWFA solves the 3-D incompressible Navier-Stokes equations and transport of potential temperature equations, which take into account the thermal buoyancy and Earth rotation (Coriolis) effects in the atmosphere. The inflow conditions for these simulations are generated using a periodic atmospheric boundary layer precursor with no turbines.

SOWFA calculates the unsteady flow field to compute the time-varying power, velocity deficits, and loads at each turbine in a wind plant. This level of fidelity takes on the order of hours to days to run on a supercomputer using hundreds to thousands of processors, depending on the size of the wind plant. Our SOWFA simulations were run on NREL’s high-performance computer Peregrine.¹⁵

It is important to note that SOWFA has been validated and is a good representation of true atmospheric behavior in a wind farm. It has been compared with the 48-Lillgrund wind farm field data and shows good agreement through the first five turbines in a row aligned with the wind direction.¹⁶ In addition, SOWFA has been tested to verify that it captures the inertial range in the turbulent energy spectra and log layer in the mean flow, both of which characterize a real atmospheric boundary layer.¹² Further validation studies are ongoing.

To validate FLORIS-3D, actuator disk simulations of two-turbine scenarios were performed using SOWFA. The turbines were simulated using the NREL 5-MW reference turbine¹⁷ and were spaced 7 rotor diameters (7D) apart in the downstream direction. These scenarios were simulated under neutral atmospheric conditions with an 8 m/s mean wind speed and 10% turbulence intensity.

A baseline scenario was run in which both the upstream and downstream turbines were simulated at a hub height of 90 m. Next, the hub height of the downstream turbine was varied to verify that FLORIS-3D could capture the effects of varying hub heights. Specifically, the upstream turbine remained at a 90 m hub height and the downstream turbine was set at 65 m and 115 m hub heights. Figure 2 shows these results compared to the FLORIS-3D wake model.

When tuned for neutral atmospheric conditions and 10% turbulence, FLORIS and SOWFA predict very similar power production of each turbine for the turbines with different hub heights. There are only a few data points obtained from one atmospheric condition, but this indicates that even this simple wake model can be useful to predict wake losses in three dimensions.

B. Annual Energy Production Calculation

The instantaneous power production of a wind farm is highly dependent on the wind direction, due to the wakes created behind wind turbines. For this reason, AEP is a much better indicator of a productive farm than power. Wind farm AEP takes into account the power production for all wind speeds and directions as well as the associated frequencies. The wind direction frequency and wind speed data used in this study are from the Princess Amalia Wind Farm, an offshore farm in the Netherlands. The direction frequency data is binned into 5° increments and the wind speeds are averaged for each of the 72 bins.

To account for height differences for our inflow velocity, we adjusted the wind speed data for wind shear.




Figure 2. The results of the FLORIS model validation with SOWFA. As can be seen, the percent error between the power predicted by FLORIS and SOWFA is minimal (close to or less than 1%). Even more importantly, the FLORIS model correctly predicts the trends of power decrease as a result of wake interference in different 3-D situations.

We used a power law to estimate the wind speed at different heights (z):

$$U(z) = U_{ref} \left(\frac{z}{z_{ref}} \right)^\alpha \quad (1)$$

where the reference height, z_{ref} , of the reference turbine is 90 m, and the shear coefficient, α , was varied as will be discussed later.

C. Tower Model

Because the tower height was allowed to vary, it was necessary to include a model to calculate mass and perform structural analysis of the tower. The structural analysis was used to constrain the optimization, keeping the towers from growing unrealistically tall where failure from stress or buckling would be an issue. It was also necessary to provide gradients for all of our constraints, which included the von Mises stress, shell buckling, and global buckling at any point along the tower; the tower taper ratio; and the first natural frequency of the structure. NREL developed a finite element model called TowerSE¹⁸ that makes various calculations along the length of a tower. It is a powerful tool, but does not provide analytic gradients. We optimized several wind farms using TowerSE and finite difference gradients, and identified the shell buckling and first natural frequency as the only active constraints. We were then able to pull out the necessary calculations from TowerSE and find the associated gradients.

The tower mass was a simple calculation from the volume of the tower. The gradients were simple to solve by hand. We found shell buckling as a function of the tower geometry and the stresses at each location, following the method outlined in Eurocode.¹⁹ These calculations were made in Fortran 90 and exact gradients were obtained with the Tapenade automatic differentiation tool.²⁰ We simplified the frequency calculation by approximating the tower as a cantilever beam of constant cross section with an end mass. We used the method described by Erturk et al. to calculate the natural frequency.²¹ Because the turbine tower does not really have a constant mass density along the length and the mass from the rotor nacelle assembly is

slightly offset at the top, our calculation is slightly more conservative than that predicted by TowerSE by about 10%. For this reason we scaled our frequency calculation by 10% to more closely match the frequency calculated by TowerSE. We chose this simplified model so that we could find gradients, which were obtained using analytic sensitivity equations.

D. Cost Model

AEP is a standard objective in wind farm optimization problems because it is easy to calculate and is a valid measure when only power production is affected by the optimization. When the tower heights are included as design variables, this measure is no longer appropriate. Taller towers will result in higher AEP because of the higher wind speeds, but this increased energy production comes at the expense of higher turbine capital cost. Shorter turbines may also increase AEP from decreased wake interference. To accurately represent these intricacies, we evaluated our wind farm by its COE.

To find the COE, we defined the cost of the wind farm as:

$$\text{farm cost} = \text{FCR}[\text{TCC}(z_i, \vec{d}_i, \vec{t}_i) + \text{BOS}] + \text{O\&M}(x_i, y_i, z_i) \quad (2)$$

where FCR was the fixed charge rate, TCC was the turbine capital cost (sum of the tower, rotor, and nacelle costs), BOS were the balance-of-station costs, and O&M were the operation and maintenance costs. The variables z , \vec{d} , and \vec{t} represented the tower height, the vector describing the tapered tower diameter, and the vector describing the shell thickness, respectively. In our model, the rotor and nacelle were the same for all turbines and the tower cost was a function of the tower mass (m):

$$\text{Tower Cost} = \alpha m \quad (3)$$

where $\alpha = 3.08$ \$/kg. The balance of station cost was constant and is a function of wind farm capacity.²² Operation and maintenance costs scaled with AEP, and are therefore an indirect function of x , y , and z as well.²³

With the wind farm capital cost and AEP calculated, the cost of energy (COE) is found as:

$$\text{COE} = \frac{\text{FCR}[\text{TCC}(z_i, \vec{d}_i, \vec{t}_i) + \text{BOS}] + \text{O\&M}(x_i, y_i, z_i)}{\text{AEP}(x_i, y_i, z_i)} \quad (4)$$

where x and y represent the position of each turbine in the horizontal plane. Figure 3 shows a simplified flowchart of how COE is calculated.

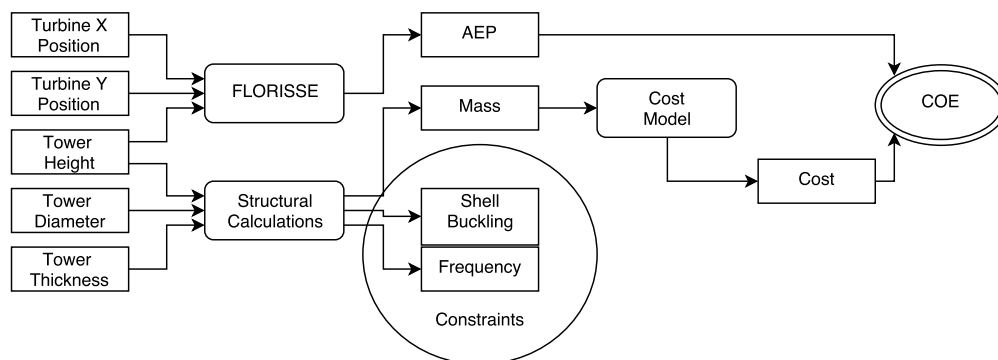


Figure 3. A simplified flow chart of COE calculation. Note that tower diameter and tower thickness are both vectors to describe the varying tower diameter and thickness along the height. The circled boxes on the bottom, shell buckling and frequency, are used as constraints in the optimization and are not used in the calculation of COE. The five boxes on the left are the design variables, and the five boxes in the middle (AEP, Mass, Shell Buckling, Frequency, and Cost) are intermediate variables.

E. Optimization

The purpose of this study was to optimize a wind farm for COE. To do so, we assigned each turbine to one of two groups, where all turbines in a group had the same tower height, diameter, and shell thickness.

Manufacturing each tower with custom dimensions would be very expensive and unrealistic. The cost and complexity both increase with the number of different turbine heights. We chose two groups because this is the smallest number that still allowed us to study the benefits of integrating turbine design. We parameterized the tower by specifying the diameter and shell thickness at the bottom, midpoint, and top of the tower and then linearly interpolating diameter and shell thickness at points in between.

It may be beneficial to do a binary optimization in which each turbine can change the height group to which it belongs, but this greatly increases the complexity of the optimization and makes it gradient-free. Binary variables, such as turbine group assignment, have no intermediate values. They are either one or the other. This means there is no way to use gradients in their optimization. Gradient-free optimization is more computationally expensive, which severely limits the number of design variables we can include in the problem. To maintain the gradient-based optimization, we assigned each turbine to one of the height groups before starting the optimization. Once assigned a turbine could not switch to the other group.

We ran several cases in which different design variables were included in the problem to allow comparison of their effects on COE. In all, the design variables we included were the position of each turbine (x_n, y_n) , the tower height of each group (H_1, H_2) , the tower diameter of each group $(d_{1,j}, d_{2,j})$, and the tower shell thickness of each group $(t_{1,j}, t_{2,j})$. Index j refers location on the tower ($j=1$ is at the bottom, $j=2$ at the midpoint, $j=3$ at the top). There are six total variables to define diameter (three for each height group), and six to define the tower thickness.

The position of each turbine was constrained so that it could not be within two rotor diameters of any other turbine in the wind farm. Also, each turbine was constrained so that it could not leave the convex hull of the original turbine layout at the beginning of the optimization. This constraint ensured that the turbines did not simply spread far apart to decrease COE. The tower heights were also constrained to be taller than the rotor radius plus the ground clearance, which we set as 10 m, which allowed us to separate the heights of different turbines while keeping a safe distance from the ground. The tower diameter was constrained to be less than 6.3 m for transportation, and greater than or equal to 3.6 at the top, to allow for the connection to the nacelle. Each tower was also structurally constrained by the shell buckling and natural frequency of the tower. The shell buckling constraint was applied to each height group for both the maximum thrust conditions and the survival load, with a safety factor of 1.35 for the loads and 1.1 for buckling resistance. The first natural frequency of the tower was constrained to be greater than the frequency at which the blades rotate and less than the blade passing frequency, with a factor of safety of 1.1. The diameter-to-thickness ratio was constrained to be greater than 120 at any point, to allow for welding. The optimization can be expressed:

$$\begin{aligned}
& \text{minimize} && \text{COE} \\
& \text{w.r.t.} && x_i, y_i, H_{1,2}, d_{(1,j)}, d_{(2,j)}, t_{(1,j)}, t_{(2,j)} \\
& && i = 1, \dots, n; j = 1, 2, 3 \\
& \text{subject to} && x_{\text{initial, min}} \leq x_i \leq x_{\text{initial, max}} \\
& && y_{\text{initial, min}} \leq y_i \leq y_{\text{initial, max}} \\
& && \sqrt{(x - x_i)^2 + (y - y_i)^2} \geq 2D_{\text{rotor}} \\
& && H_1, H_2 \geq r_{\text{turbine}} + 10 \text{ m} \\
& && d_{(1,j),(2,j)} \leq 6.3 \text{ m} \\
& && d_{(1,\text{top}),(2,\text{top})} \geq 3.6 \text{ m} \\
& && \frac{3\Omega}{1.1} \geq f_{1,2} \geq 1.1\Omega \\
& && \text{shell buckling margins: max thrust} \leq 1 \\
& && \text{shell buckling margins: survival load} \leq 1 \\
& && \frac{d_{(1,j)}}{t_{(1,j)}}, \frac{d_{(2,j)}}{t_{(2,j)}} \geq 120
\end{aligned} \tag{5}$$

Note that i is the index defining the wind turbine, and j is the index describing the location on the tower.

The results of gradient-based optimization, for problems with many local minima, are sensitive to the starting location. As in most optimization problems, there is no guarantee that the solution is the global solution. The best results can be achieved with a multiple-start approach, where several different starting

points are used for each condition, and the best solution is used. For our study, each optimization started from an equally spaced 5-by-5 turbine grid. The tower height groups were alternated so that the starting layout made a checkerboard pattern with 13 turbines in one height group and 12 in the other. These standardized starting points allowed us to better compare our solutions for each condition. The gradients for this optimization were all analytic. We calculated the partial derivatives of each small section of the model and included each part in a framework called OpenMDAO,²⁴ which calculates the gradients of the entire system. The analytic gradients are significant because they are more accurate and converge on a solution much faster than finite difference gradients. More importantly, they allow us to solve much larger optimization problems.

III. Results

Wind farms with multiple hub heights are more advantageous in certain conditions. Important factors that might affect this advantage include the wind farm boundary, wind shear exponent, rotor size, spacing constraints, and turbine type. We explored two factors: wind turbine density and wind shear exponent. We chose these factors because they are both site dependent and will be useful in determining if a site is a good candidate for a wind farm with different hub heights. To compare the results, we ran four different situations for each condition: the starting grid layout, an optimized layout in which the tower height was fixed, an optimized layout in which turbines could change height but must all be the same height, and an optimized layout in which turbines could change height within two different height groups.

A. Varied Turbine Density

The first variable studied was the turbine density in the wind farm, or the ratio of the area of the farm occupied by wind turbines to the total area of the wind farm:

$$\text{Turbine Density} = \frac{\pi R^2 N}{A} \quad (6)$$


where R is the rotor radius, N is the number of turbines, and A is the area of the wind farm. For this study, the shear exponent α from Eq. (1) was held constant at 0.1. Density was varied by changing the farm size between 64 square rotor diameters up to 400 square rotor diameters.

Figure 4(a) shows an optimized wind farm layout with low turbine density. The different colors correspond to the different height groups shown in Fig. 4(b). Figure 4(c) and Fig. 4(d) also show an optimized wind farm and the corresponding turbine heights, but for the case of high turbine density. In Fig. 4(a), for the case of low turbine density, the turbines are very far apart, and can easily move horizontally. Thus, we see in Fig. 4(b) that the optimized heights are the same. The high density case Fig. 4(c) is not able to move as much horizontally, so it has lower COE by separating the two height groups shown in Fig. 4(d).


Figure 5 shows the COE optimized for each wind farm under each condition previously discussed. The cyan points at the top represent the farms that have not been optimized and the black points are the farms that have been optimized for layout. The red points have been optimized for layout and height (there is only one height group), and the blue points have been optimized for layout and height, where two different height groups are allowed.

Low turbine density logically results in low COE because the turbines remain far apart and wake effects are not as high (See Fig. 4(a)). We can see from the blue points in Fig. 5 that optimizing with different hub heights significantly decreases COE for the cases with high turbine density. For the highest turbine density, slightly above 30%, there is a COE decrease of over 9% from the case with one height group to two height groups. There is a 6% and 4% decrease in COE for a turbine density of 20% and 15%, respectively. This is because at high density, the horizontal movement of turbines is severely limited by spacing constraints. Therefore, there is a greater benefit to moving vertically. At the highest turbine density, the black and cyan points are both the same. At this density, spacing constraints are so severe that there is no turbine movement without violating spacing constraints. The only decrease in COE that can be achieved is by moving up or down. Conversely, as the wind farm grows larger, the turbines can move completely or almost completely out of the wakes of other turbines only with horizontal movement.


From the data shown in Fig. 5, it appears that there is not a huge benefit to allow the turbines to change height together with a shear exponent of 0.1. The black points corresponding to layout optimization only have slightly higher COE than the red points, which show the farm with one height group. These wind




(a) Optimized turbine layout with a low turbine density of 1.92%.



(b) Optimized turbine heights with low turbine density.



(c) Optimized turbine layout with a high turbine density of 19.6%.



(d) Optimized turbine heights with high turbine density.

Figure 4. These figures show some of the optimization results with varied turbine density. An optimization with a low turbine density of about 2% is seen in (a), and has no separation between the height groups as seen in (b). A high turbine density of 20% is seen in (c), and results in large separation between the two height groups shown in (d).




Figure 5. This figure shows results for a 25-turbine wind farm that was optimized for cost of energy (COE) as a function of turbine density. The cyan points represent the original grid, or the starting COE. The black points have been optimized for location, but not for height. The red points are optimized for position and height but all the turbines are the same height. The blue points have been optimized for position and height, with two different height groups.

farms were all optimized with a low shear exponent (0.1). The wind speed does not vary quickly with height, meaning that the benefit of the slightly higher wind speeds from taller towers does not significantly outweigh the additional cost of larger towers.

Figure 6 shows the tower height for each of the height groups as a function of turbine density. This only applies to the case in which there were only two different height groups. The colors are the same as in Fig. 4. As shown, when the turbines are tightly packed (density higher than 5%), the optimizer varied the heights significantly to minimize COE. Any farm with a lower density does not benefit from different tower heights. Notice that when all the heights are the same, they are not at the maximum height. The shear exponent, 0.1, does not result in high enough wind speeds to make it worth the cost of building larger turbines.

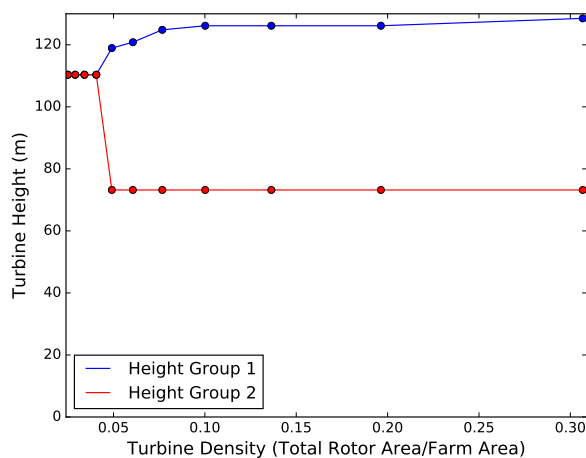




Figure 6. This figure shows the optimized heights of each height group as a function of turbine density when turbine heights are allowed to change with two height groups. See at around 5% turbine density, all the turbines are optimized to the same height around 110 m.

B. Varied Shear Exponent


Wind shear exponent determines how quickly wind speed increases with height (see Eq. (1)), and is determined by the terrain of a wind farm. Open water or a barren field will have a low wind shear exponent while lots of trees or buildings will have a high shear exponent. The sites with higher shear exponent are suited for taller turbines to take advantage of the much higher wind speeds, resulting in greater energy production. At lower wind shear, there is not as great of a benefit for the taller turbines. For these situations of lower wind shear, it is more beneficial for some of the turbines to be shorter, resulting in less wake losses in the wind farm. To observe the impact of shear exponent on the benefit of different hub heights, we kept the wind farm size constant at 144 square rotor diameters (turbine density of 13.6%) and varied the shear exponent from 0.08 to 0.26. Most sites have average wind shear exponents that fall in this range.




(a) Optimized turbine layout with a shear exponent of 0.08.



(b) Optimized turbine heights with a low shear exponent of 0.08.



(c) Optimized turbine layout with a high shear exponent of 0.22.



(d) Optimized turbine heights with a high shear exponent of 0.22.

Figure 7. These figures show some optimization results with varied wind shear exponent. The optimized layout with low shear is seen in (a), and results in large separation between the two height groups shown in (b). The high wind shear exponent optimized layout seen in (c) results in no separation between the height groups and results in both towers reaching their maximum height, shown in (d).

Figure 7 shows two optimized turbine layouts and heights for a low and high shear exponent. See in Fig. 7(b) that for a low shear exponent, the tower heights reach maximum separation, while in Fig. 7(d) for high wind shear, both towers reach maximum height to take advantage of the much higher wind speeds. Note that these optimized heights differ from the case of turbine density. In Fig. 4(b), we see that when the optimal turbines heights are the same, they do not reach the maximum limit. For this case, the wind shear is low enough that it is not worth the penalty in additional cost of building larger towers to reach the

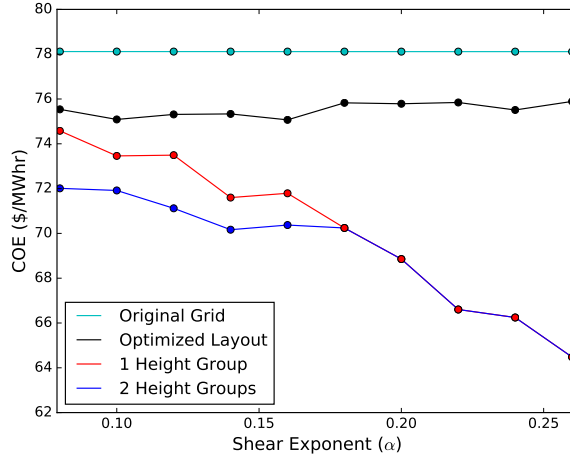


Figure 8. These are the results of a 25-turbine wind farm optimized for COE as a function of the wind shear exponent α . The cyan lines represent the original grid that has not been optimized. The black lines have been optimized for location but not for height. The red lines are optimized for position and height, but all the turbines are the same height. The blue lines have been optimized for position and height, with two height groups.

maximum height.

Figure 8 shows the optimized COE as a function of the wind shear exponent. As in Fig. 5, the cyan points at the top represent the farms that have not been optimized, and the black points are the farms that have been optimized for layout. The red points have been optimized for layout and height, but there is only one height group, and the blue points have been optimized for layout and height, in which two different height groups are allowed.

At low shear exponents (0.8-0.16), the COE of farms optimized with two height groups is lower than the farm optimized with one height group. The lowest shear exponent (0.08) resulted in a farm with nearly 5% lower COE from one height group to two, shown by the red and blue points in Fig. 8. The slightly higher shear exponents of 0.1 and 0.12 have similar benefits of 3% to 5% for the wind farms with different height groups. As the shear exponent increases, the COE from the one height group (red) and two height group (blue) converge, meaning the benefit of different height groups decreases as the shear exponent increases.

The cyan line at the very top represents the starting layout of turbines before any optimization. It does not vary with shear exponent because the turbine heights do not change. We chose the reference height (the starting turbine height) in our wind shear equation as 90 m, the point where our wind speed data was measured. When all the turbines are at the reference height, the shear exponent has no effect on the hub-height wind speed (see Eq. (1)). Because we are using a simplified power model that uses the hub speed to calculate power, AEP and COE remain unchanged.

Figure 9 shows the tower heights of each optimized height group when the tower height is allowed to change, as a function of shear exponent. These colors correspond to the same colors in Fig. 7 where the blue points are the taller height group, and the red points are the shorter height group. We see that at low shear values, the difference between the height groups is large and constant. After $\alpha = 0.18$, the smaller tower quickly approaches the maximum height until they are the same. This pattern also appears in Fig. 6, there is a large difference between the height groups until a certain turbine density, and where the height difference drops abruptly to zero. If similar behavior exists for other wind farms, these sharp transition values could help determine if a specific site is a good candidate for having different turbine heights.




Figure 9. These are the optimized heights of each height group as a function of wind shear exponent, when there are two different height groups. At a shear exponent of 0.22 and after, both height groups achieve maximum COE but both going to the maximum height.

IV. Conclusions and Ongoing Work

This paper demonstrated a method to optimize a wind farm that has turbines with different hub heights, with an eventual goal of performing coupled yaw control optimization. To do so, we modified the FLORIS wake model to work in three dimensions, and used this to predict the wind speed anywhere in a wind farm. This velocity information combined with wind frequency data from the Princess Amalia Wind Farm was used to calculate AEP. We also included a cost model, which combined with the wind farm AEP allowed us to calculate COE, which was used as the objective function during optimization. The cost model estimated the tower cost as a function of mass, which we calculated with a tower structural model. This model was also used to constrain the tower height, diameters, and thicknesses during optimization to not fail from shell buckling, or have a natural frequency below the blade rotation frequency or above the blade passing frequency.

The results indicate that wind farms with turbines of multiple hub heights can decrease the cost to produce energy for certain farms. Using the NREL 5-MW reference turbine, sites with high turbine density (higher than 15%) and 0.1 wind shear benefit greatly from different hub heights, with as much as a 5%-9% decrease in COE for very high turbine densities. This decrease is because layout optimization alone cannot move turbines out of wakes very successfully. Vertical movement provides an extra degree of freedom, which decreases wake interference. At lower turbine density, there is less benefit of having multiple hub heights in the same farm.

Farms with low wind shear can also benefit from allowing turbines to have different hub heights. At low wind shear, the wind speed closer to the ground is not much lower than the wind speed higher up. Therefore, the decreased wake interference from different hub heights outweighs the benefits of having all taller towers to capture the large wind speeds. Very low wind shears, 0.08-0.12, may decrease COE from 3 to 5% for a turbine density of around 14%. Greater wind shears do not provide as much of a COE decrease with different turbine heights, because the extra power production from the high wind speeds outweighs the benefit of decreased wake interference.

The most immediate continuation of this research will be to optimize a wind farm while including turbine yaw as a design variable. The current model has the capability to run this optimization and we expect the integrated yaw optimization to produce significantly lower COE than we have achieved so far. This integrated layout, turbine design, and yaw control optimization has many design variables and is only possible because of the analytic gradients included in this model.

This research will also be extended to include other aspects of turbine design. Specifically, we will expand the model to include rotor diameter as a design variable. This will potentially further reduce wake interference between turbines in a wind farm and decrease COE. We expect the benefits of multiple hub-

height farms to be greater when the relative size of the rotor to the tower height is smaller. The smaller relative size will allow different height groups to better avoid wake interference. When included with the other aspects of wind farm optimization addressed in this paper, the benefits could be magnified.

Acknowledgments

The BYU authors developed this journal article based on funding from the Alliance for Sustainable Energy, LLC, Managing and Operating Contractor for the National Renewable Energy Laboratory for the U.S. Department of Energy. The NREL authors were supported by the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308 with the National Renewable Energy Laboratory.

Funding for the work was provided by the DOE Office of Energy Efficiency and Renewable Energy, Wind Energy Technologies Office.

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.

References

- ¹Kusiak, A. and Song, Z., "Design of Wind Farm Layout for Maximum Wind Energy Capture," *Renewable Energy*, Vol. 35, No. 3, 2010, pp. 685–694.
- ²Şişbot, S., Turgut, Ö., Tunç, M., and Çamdah, Ü., "Optimal Positioning of Wind Turbines on Gökçeada Using Multi-Objective Genetic Algorithm," *Wind Energy*, Vol. 13, No. 4, 2010, pp. 297–306.
- ³Wagner, M., Veeramachaneni, K., Neumann, F., and O'Reilly, U.-M., "Optimizing the Layout of 1000 Wind Turbines," *European Wind Energy Association Annual Event*, 2011, pp. 205–209.
- ⁴Gebraad, P. M. O., Teeuwisse, F. W., van Wingerden, J. W., Fleming, P. A., Ruben, S. D., Marden, J. R., and Pao, L. Y., "Wind Plant Power Optimization Through Yaw Control Using a Parametric Model for Wake Effects—a CFD Simulation Study," *Wind Energy*, Vol. 19, No. 1, Jan. 2016, pp. 95–114.
- ⁵Fleming, P., Ning, A., Gebraad, P., and Dykes, K., "Wind Plant System Engineering through Optimization of Layout and Yaw Control," *Wind Energy*, Vol. 19, No. 2, Feb. 2016, pp. 329–344.
- ⁶Chen, Y., Li, H., Jin, K., and Song, Q., "Wind Farm Layout Optimization Using Genetic Algorithm with Different Hub Height Wind Turbines," *Energy Conversion and Management*, Vol. 70, June 2013, pp. 56–65.
- ⁷Hazra, J., Mitra, S., Mathew, S., and Zaini, F., "3D Layout Optimization for Large Wind Farms." *ISGT*, IEEE, 2015, pp. 1–5.
- ⁸Jensen, N., "A Note on Wind Generator Interaction," *Riso National Laboratory, Roskilde, Denmark, Riso-M-2411*, 1983.
- ⁹Thomas, J., Gebraad, P., and Ning, A., "Improving the FLORIS Wind Plant Model for Compatibility with Gradient-based Optimization," 2016, (in review).
- ¹⁰Churchfield, M. and Lee, S., "NWTC design codes-SOWFA," URL: <http://wind.nrel.gov/designcodes/simulators/SOWFA>, 2012.
- ¹¹Jonkman, J., "NWTC design codes (FAST)," *NWTC Design Codes (FAST)*, NREL, Boulder, CO, 2010.
- ¹²Churchfield, M. J., Lee, S., Michalakes, J., and Moriarty, P. J., "A Numerical Study of the Effects of Atmospheric and Wake Turbulence on Wind Turbine Dynamics," *Journal of turbulence*.
- ¹³Fleming, P. A., Gebraad, P. M., Lee, S., van Wingerden, J.-W., Johnson, K., Churchfield, M., Michalakes, J., Spalart, P., and Moriarty, P., "Evaluating Techniques for Redirecting Turbine Wakes Using SOWFA," *Renewable Energy*, Vol. 70, 2014, pp. 211–218.
- ¹⁴Fleming, P., Gebraad, P. M., Lee, S., Wingerden, J.-W., Johnson, K., Churchfield, M., Michalakes, J., Spalart, P., and Moriarty, P., "Simulation Comparison of Wake Mitigation Control Strategies for a Two-Turbine Case," *Wind Energy*, Vol. 18, No. 12, 2015, pp. 2135–2143.
- ¹⁵Regimbal, K., Carpenter, I., Chang, C., and Hammond, S., "Peregrine at the National Renewable Energy Laboratory," *Contemporary High Performance Computing: From Petascale toward Exascale, Volume Two*, Vol. 23, 2015, pp. 163.
- ¹⁶Churchfield, M. J., Lee, S., Moriarty, P. J., Martinez, L. A., Leonardi, S., Vijayakumar, G., and Brasseur, J. G., "A Large-Eddy Simulation of Wind-Plant Aerodynamics," *AIAA paper*, Vol. 537, 2012, pp. 2012.
- ¹⁷Jonkman, J., Butterfield, S., Musial, W., and Scott, G., "Definition of a 5-MW Reference Wind Turbine for Offshore System Development," Tech. rep., DOE, 02 2009.
- ¹⁸Ning, S. A., "TowerSE," Tech. rep., National Renewable Energy Laboratory (NREL), Golden, CO (United States), 2013.
- ¹⁹EN, C., "1-1, Eurocode 3: Design of Steel Structures," *Part 1-6: supplementary rules for the shell structures*, 1993.
- ²⁰Hascoet, L. and Pascual, V., "The Tapenade Automatic Differentiation Tool: Principles, Model, and Specification," *ACM Transactions on Mathematical Software (TOMS)*, Vol. 39, No. 3, 2013, pp. 20.
- ²¹Erturk, A. and Inman, D. J., "Appendix C: Modal Analysis of a Uniform Cantilever With a Tip Mass," *Piezoelectric Energy Harvesting*, 2011, pp. 353–366.
- ²²Mone, C., Maples, B., and M., H., "Land-Based Wind Plant Balance-of-System Cost Drivers and Sensitivities," 2014.
- ²³Moné, C., Smith, A., Maples, B., and Hand, M., "Cost of Wind Energy Review," Tech. rep., NREL/TP-5000-63267. Golden, Colorado: National Renewable Energy Laboratory, 2013.

²⁴Gray, J., Moore, K. T., and Naylor, B. A., “OpenMDAO: An Open Source Framework for Multidisciplinary Analysis and Optimization,” *AIAA/ISSMO Multidisciplinary Analysis Optimization Conference Proceedings*, Vol. 5, 2010.