

RESEARCH ARTICLE

WILEY

Blade planform design optimization to enhance turbine wake control

Jeffery Allen¹  | Ethan Young¹  | Pietro Bortolotti²  | Ryan King¹ | Garrett Barter² 

¹Computational Science Center, National Renewable Energy Laboratory, Golden, Colorado, USA

²National Wind Technology Center, National Renewable Energy Laboratory, Golden, Colorado, USA

Correspondence

Jeffery Allen, Computational Science Center, National Renewable Energy Laboratory, Golden, CO, USA.

Email: Jeffery.allen@nrel.gov

Funding information

Office of Energy Efficiency and Renewable Energy, Grant/Award Number: DE-AC36-08GO28308

Abstract

This study considers optimizing the planform of wind turbine blades to ultimately enhance wind plant controls, namely, wake steering strategies. Adjoint-enabled unsteady actuator line simulations are carried out to obtain gradients for optimization of several different performance objectives with respect to blade chord length at 10 locations along blade span. We demonstrate different blade design optimizations that can maximize time-averaged lateral wake deflection, entrainment of kinetic energy, or total power of multiple turbines. Our optimized designs can produce a 4+ % increase in wake deflection, a 4× increase in vertical kinetic energy entrainment, or a 3.6% increase in power when compared with the baseline case. While lateral wake deflection is only modestly sensitive to chord changes, we find that increasing the outboard chord length can dramatically increase kinetic energy entrainment, resulting in faster wake recovery and gains in net power. While this work develops only a few case studies emphasizing relative performance improvements and general trends, these results show the promise of a framework that combines mid-fidelity computation with adjoint-based optimization for control and design problems.

KEYWORDS

adjoint optimization, wind plant controls, wake steering, blade design

1 | INTRODUCTION

Modern wind plants are increasingly tasked with multiple performance objectives. In addition to designing plants that maximize power output and minimize the levelized cost of energy, the design and operation of wind plants is increasingly influenced by challenges regarding grid integration of variable generation renewables. This places a growing emphasis on making wind plants more controllable and predictable. Plant-level control strategies, such as wake steering or axial induction control, have emerged as a viable solution to some of wind energy's grand challenges¹ by offering a tool for plant operators to exert more control authority over an existing wind plant and also increase power production by reducing plant wake losses. However, these operating regimes are often not prioritized during the design phase, which results in turbines and plants that fail to maximize all of the available power capture at the plant level. Shifting to plant-focused design is further challenged because traditional turbine rotor design tools are based on blade element momentum theory. Models that can trace the impact of blade design choices through the complex turbulent flow physics of wake interactions to understand their effect on control strategies are typically too computationally intensive to be used in the core rotor design iteration loop. Therefore, turbine blades are designed in a *greedy* approach, where single turbine power production is the

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs](https://creativecommons.org/licenses/by-nc-nd/4.0/) License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2022 The Authors. *Wind Energy* published by John Wiley & Sons Ltd.

design objective. This paper's contribution is demonstrating wind turbine blade design using an *altruistic* approach, where plant-level power production and flow control are the ultimate design objectives.

1.1 | Plant-level flow control strategies

Plant flow control was originally proposed via adjustment of the axial induction (rotor thrust) of each turbine in the farm.²⁻⁴ However, consensus is building that greater control authority and loss recovery can be achieved through wake deflection via yaw misalignment. This has also been studied extensively in the literature, which includes not just model-based studies but also experiments in wind tunnels⁵⁻¹¹ and field trials.¹²⁻¹⁵ Wake deflection via steering with yaw offsets sacrifices some power production from upstream turbines but can significantly increase power production of downstream turbines. Therefore, even if the turbines themselves are designed with a greedy objective, they can be operated in an altruistic sense. Additionally, other studies have shown benefits of wake control steering beyond recovery of wake losses. The wake structure behind a turbine misaligned with the wind sets up a pair of counter-rotating vortices that creates a kidney-bean-shaped wake.¹⁶⁻²⁰ This increases the vertical mixing and entrainment of higher velocity fluid from the layers above the wind plant, re-energizing the lost momentum.^{21,22} The question remains, how much more benefit is available if the turbines were more tailored for this type of operation?

Wake steering requires a plant-level super controller that can override an individual turbine's default yaw settings. It also forces a turbine engineer to expand their design focus from a single turbine to the whole plant. If wake steering will be an accepted plant control strategy, then it is worth exploring whether there are blade designs that cater to the performance benefits of this strategy better than standard approaches. This could entail calibrating different spanwise induction profiles along the blade that deviate from classical actuator disk models, depending on a turbine's placement in the array and the local wind rose. This could have numerous advantages including

- increased power production by requiring less yaw misalignment to achieve the same wake deflections and/or entrainment of higher momentum flow;
- accelerate wake recovery by enhanced kinetic energy entrainment;
- load reductions by reducing partial waking alignments (via greater deflection) on downstream turbines and less yaw misalignment on upstream turbines;
- enhanced wind plant dispatchability, essentially the ability to track to a desired power production profile over time that might be specified by utility operators to support overall grid stability and reserves as the load and other generation sources vary; and
- better understanding of flow physics.

1.2 | Background on rotor design

The design of wind turbines has historically been optimization of the single turbine. One possible merit figure of a conventional optimization process is the levelized cost of energy,²³ and although some of the inputs belong to the wind farm, the models do not usually investigate how different turbine designs impact the interactions among the turbines. One notable exception consists of the recent paper from Stanley and Ning, 2019,²⁴ in which the turbine configuration and the plant layout are optimized simultaneously. The design approach is based on blade element momentum rotor aerodynamics and on a continuously differentiable wake model. The design variables parameterize hub height, rotor diameter, rated power, tower diameter, and tower shell thickness. Blade chord and twist distributions are also part of the optimization, but they are treated as implicit variables mapping a surrogate model for minimum blade mass.

While customizing the blade design for every turbine in the plant at every project site is obviously infeasible due to the need for industrialization and efficient manufacturability, there are less extreme and more practical options. For instance, General Electric announced in 2018 its Cypress turbine product line with two-piece blades allowing for increased blade length and better site customization.²⁵ For plant flow controllability, it is therefore conceivable that front row turbines could have a slightly different blade design than inner array turbines, or perhaps like the Cypress approach, a few different possible blade building blocks that can be mixed and matched together. This study only strives to highlight the design space possibilities when considering plant flow controllability, but the authors certainly acknowledge that the design space will be far more constrained for real-world applications.

One of the obstacles to moving from a turbine-focused blade design paradigm to a plant-focused one is that it requires a significant jump in computational resources and design tool sophistication. Turbine-focused design can rely on design tools built around traditional blade element momentum models, which are widely used, well understood, and computationally efficient. Plant-focused design requires a significant growth in the simulation control volume, from just the rotor plane to at least the full wake extent, and also in the fidelity of the numerical scheme, likely to a computational fluid dynamics (CFD) model that can connect the subtleties of blade design parameters to wake behavior. This work uses recent improvements in the WindSE software package,²⁶ developed at the National Renewable Energy Laboratory (NREL). WindSE is a Reynolds-

averaged Navier–Stokes model designed around analytical gradient and adjoint methods, with the ability to capture terrain-induced effects. The recent addition of an unsteady solver with an actuator line method gives it a unique niche to explore plant-level controls and design problems.

1.3 | Original contributions

Compared to the work from Stanley and Ning, 2019,²⁴ the work described in this paper elevates the fidelity level of the underlying wind turbine model and optimization framework. The focus of the design optimization is also different. While the work from Stanley and Ning, 2019, parameterizes the overall turbine configuration, this work investigates in more detail the solution in terms of flow and rotor aerodynamics. Three optimization studies are conducted to consider different flow control and plant performance objectives:

1. Wake deflection—The goal of the first optimization is to investigate the potential for improved wake steering abilities by varying the blade outer chord distribution to maximize the lateral deflection of wind turbine wakes.
2. Kinetic energy entrainment—The goal of the second optimization is to determine a chord distribution that enhances the entrainment of kinetic energy to re-energize the flow in the plant.
3. Power optimization—The goal of the third optimization is to find a blade design that optimizes the power output of multiple turbines by balancing wake deflection, entrainment, and individual power output.

The paper is organized as follows: Section 2 contains detailed descriptions of the methods employed in the unsteady solver, the actuator line model, the adjoint gradient computation, and the baseline rotor configuration. The details of the three studies and their results are reported in Sections 3.1, 3.2, and 3.3, respectively. Finally, a discussion of the results and conclusions for further rotor design studies is contained in Section 4.

2 | METHODS

2.1 | WindSE unsteady solver

Altering flow characteristics, such as combined wake interactions and vertical entrainment, to optimize plant power output requires a complete time-varying solution of the flow field. The physics governing this flow are set by the filtered Navier–Stokes and continuity equations

$$\rho \left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} \right) = -\nabla P + \mu \nabla^2 \mathbf{u} - \nabla \cdot \boldsymbol{\tau} + \mathbf{F} \quad (1)$$

$$\nabla \cdot \mathbf{u} = 0 \quad (2)$$

where \mathbf{u} is the velocity field implicitly filtered by the computational grid, P is the pressure field, $\boldsymbol{\tau}$ is the subgrid stress tensor, \mathbf{F} is an external forcing term attributable to rotor induction, μ is the dynamic viscosity, and ρ is the fluid density. These equations are discretized spatially using first-order continuous piecewise finite elements. First-order elements are employed for both the velocity and pressure function spaces rather than invoking a combination of higher-order elements in order to reduce the total degrees of freedom. An additional stabilization term is added as well. Temporally, the equations are solved using an implementation of the fractional-step method, which allows the fluid velocity to be updated while enforcing the incompressibility constraint.^{27,28} The first stage of this multi-step process is to discretize Equation (1) in time and then solve for the predicted fluid velocity at the next timestep

$$\frac{\mathbf{u}^* - \mathbf{u}^k}{\Delta t} = \nu \nabla^2 \mathbf{B} - \mathbf{A} \cdot \nabla \mathbf{B} - \nabla P^k + \mathbf{F}^k \quad (3)$$

where \mathbf{u}^* is the predicted velocity at timestep, $k + 1$, and \mathbf{u}^k is the velocity at the current timestep, k . Terms \mathbf{A} and \mathbf{B} are expressions of the filtered velocity field chosen to improve numerical stability: \mathbf{A} is an explicit Adams–Bashforth projection, $\mathbf{A} = 1.5\mathbf{u}^k - 0.5\mathbf{u}^{k-1}$, and \mathbf{B} is an implicit Crank–Nicolson interpolation, $\mathbf{B} = 0.5(\mathbf{u}^* + \mathbf{u}^k)$. The discrete timestep size, Δt , is proportional to the ratio of the minimum mesh length scale to the maximum velocity, $\Delta t = \text{CFL} \times (\Delta_{\min}/|\mathbf{u}|_{\max})$, where the CFL term is a scaling quantity derived from the Courant–Friedrichs–Lewy condition. In the studies presented here, the maximum velocity, \mathbf{u}_{\max} , is effectively a constant value associated with the blade tip speed, \mathbf{u}_{tip} ; however, the relationship is presented here in the more general form. In general, the predicted velocity returned using this scheme is not divergence free. Correcting this to enforce governing Equation (2) is done by solving for the change in the pressure field, ϕ , required to bring the flow to a divergence-free state.

$$\nabla \cdot \mathbf{u}^* = \Delta t \nabla^2 \phi \quad (4)$$

This corrective pressure change is now used to update the predicted fluid velocity, \mathbf{u}^* , to a divergence-free state. This result of this operation is the true, divergence-free velocity associated with timestep $k + 1$.

$$\mathbf{u}^{k+1} = \mathbf{u}^* - \Delta t \nabla \phi \quad (5)$$

$$p^{k+1} = p^k + \phi \quad (6)$$

The total viscosity, ν , in Equation (3) is the sum of the background kinematic viscosity, ν_0 , and a subgrid scale eddy viscosity, ν_T ,

$$\nu_T = (C\Delta)^2 |\mathbf{S}| \quad (7)$$

$$\mathbf{S} = \frac{1}{2} (\nabla \mathbf{u} + (\nabla \mathbf{u})^T) \quad (8)$$

employed to model the dissipative effects of the smallest turbulent length scales without the prohibitive mesh refinement that would be necessary to simulate them, where C is the Smagorinsky constant, currently 0.17, Δ is the characteristic length associated with each cell of the computational grid, and \mathbf{S} is the resolved strain rate tensor shown in Equation (8).

For boundary conditions, the inflow is a log-law profile, $u_x(z) = u^* / \kappa \ln(z/z_0)$, where u^* is the reference velocity in m/s and κ is the von Karman constant. A no-slip condition, $\mathbf{u} = 0$, is enforced on the ground. The spanwise and top boundaries have a free-slip, $\mathbf{u} \cdot \vec{n} = 0$, condition. Finally, the outflow boundary is set to no-stress.

2.2 | Actuator line model implementation

The effects of blade geometry and rotor induction are imposed via the external forcing term, \mathbf{F} , in Equation (1). These time-dependent blade forces are distributed spatially using an implementation of the actuator line method, in which a summation of point forces projected onto small spherical volumes defined by Gaussian distributions are positioned to reproduce the effect of discrete rotor blades.^{29,30} The forcing term is given by

$$\mathbf{F}(x, y, z, t) = - \sum_{j=1}^N \mathbf{f}_j(x_j, y_j, z_j, t) \frac{1}{\epsilon^3 \pi^{3/2}} \exp\left(-\frac{|\mathbf{d}_j|^2}{\epsilon^2}\right) \quad (9)$$

where N is the number of blade discretization segments, defined by the ratio of the rotor radius to the minimum mesh length scale, $\lceil R/\Delta_{\min} \rceil = 10$ in the current work, and ϵ controls the characteristic width of the Gaussian, $2\Delta_{\min}$. \mathbf{d}_j is the distance vector between any point in the fluid with position (x, y, z) and the j th actuator point located at (x_j, y_j, z_j) . The force being distributed from each actuator point, \mathbf{f}_j , is the sum of the lift and drag vectors, \mathbf{f}_l and \mathbf{f}_d , which are generated by projecting the magnitudes of lift and drag, \mathcal{L} and \mathcal{D} , onto directions perpendicular and parallel, respectively, to the relative fluid velocity. The magnitudes of lift and drag are given by

$$\mathcal{L}(x, y, z, t) = \beta \frac{1}{2} C_l(\alpha) \rho c w |\mathbf{u}_{rel}(x, y, z, t)|^2 \quad (10)$$

$$\mathcal{D}(x, y, z, t) = \beta \frac{1}{2} C_d(\alpha) \rho c w |\mathbf{u}_{rel}(x, y, z, t)|^2 \quad (11)$$

where c is the chord length specified at the actuator point, w is the width defined by the spacing of the actuator points, $R/(N - 1)$, and \mathbf{u}_{rel} is the relative fluid velocity arising from the prescribed rotation of blade segment and the local flow. The coefficients of lift and drag, C_l and C_d , are interpolated from lookup tables of airfoil polars as a function of local angle of attack, α , which itself is a function of the relative velocity and the twist of the blade. The tip-loss factor, β , diminishes the force near the tip of the blade

$$\beta = \frac{2}{\pi} \cos^{-1} \left[\exp\left(\frac{3}{2} \frac{(R-r)}{r \sin(\alpha_{rp})}\right) \right] \quad (12)$$

where r is a scalar in the range $[0, R]$ that expresses radial position along the blade and α_p is the angle formed by the relative velocity and the rotor plane. The rotor power, W , generated by a turbine represented in this manner is

$$W = \omega \sum_{j=1}^N r_j (\mathbf{f}_j \cdot \hat{\mathbf{n}}_j) \tag{13}$$

where ω is the angular velocity of the rotor, r_j is the radial distance to actuator point j and $\hat{\mathbf{n}}_j$ is a unit vector constructed in the rotor plane oriented normal to the blade's spanwise axis at point j .

2.3 | Adjoint gradients

For high-dimensional optimization problems, design gradients provide crucial information that can reduce the number of expensive function evaluations required. Adjoint methods provide an efficient way to obtain high-dimensional gradients that avoids the curse of dimensionality. Although deriving and solving the adjoint equations are typically quite cumbersome, NREL's WindSE tool provides tools to automate this process. WindSE was originally developed for adjoint-based optimization of wind plant layouts³¹; however, the adjoint capability has been expanded to develop data-driven flow model corrections,³² learn wind plant surrogate models,³³ and model the effects of terrain.³⁴

The WindSE python package (<https://github.com/nrel/windse>) uses FEniCS as the finite-element backend.³⁵⁻³⁸ FEniCS handles the mesh generation, assembly of the finite-element system of equation, and approximating the solution to that system of equation using the PETSc solver library.³⁹ The backbone of the optimization is performed by a python package known as `dolfin-adjoint`.⁴⁰ This package is also built with FEniCS and allows for the computation of gradients with respect to partial differential equations. It achieves this by using what is referred to as the "tape." The tape records all of the operations required to calculate some objective function as well as all the connections between these operations. Three items are required to perform an optimization, namely, the design variables, a forward problem, and an objective function. The design variables are the parameters that are to be optimized. The forward problem is the series of operations that are required to calculate the value of the objective function. These operations can be as simple as addition or multiplication and can be as complex as a full solve of a system of partial differential equations. Finally, the objective is the value to be optimized.

As the forward problem is being performed for the first time, each operation is recorded to the tape as a series of "block" objects. Figure 1 contains a representation of the tape which is made up of these blocks. These blocks contain all the information required to recompute their output as well as which design variables they depend on. The key feature of a block is that it can compute the local derivative of its output with respect to the relevant design variables. For the unsteady actuator line solver, a model of the tape is presented in Figure 1. This diagram shows three timesteps of the unsteady solver starting with timestep $k - 1$. Each timestep contains the three solves required for the pressure-correction scheme, which are described in Section 2.1. These solves are represented by the blocks labeled S_1 , S_2 , and S_3 . The block labeled \mathcal{F} is the function that produces the external forces computed using the actuator line model from Section 2.2. The last block of interest is the one labeled J . This is the block that uses the solution from the unsteady solve to compute the objective function whether it is the wake deflection, the vertical kinetic entrainment, or power generated. Some other minor blocks include the time averaging of the objective or assignment operations required to keep track of the flow of information.

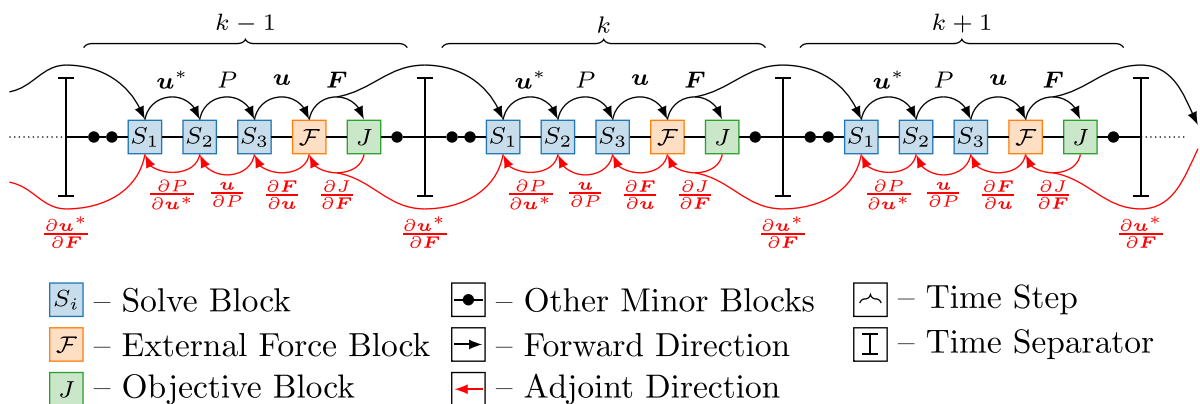


FIGURE 1 Scheme depicting how `dolfin-adjoint` records the forward problem and objective function onto a tape. The reverse process propagates the gradient backwards through this tape

After the forward problem is finished, the adjoint operation works backwards by “rewinding” the tape. Because the final objective is essentially the output of a chain of functions, the gradient can be found using a massive chain-rule derivative. The adjoint of this chain-rule allows each block to only care about computing their local derivatives as a contribution to this chain-rule. For more details, see Funke and Farrell.⁴¹ Working backwards through the tape in this manner allows for the computation of the objective functions gradient with respect to the design variables, $\partial J/\partial c_i$. This gradient can be fed into an optimization routine. All the case studies in this paper use the Sequential Least-Squares Programming (SLSQP) algorithm accessible with the Python library `scipy`,⁴² which was observed to converge well and exhibit relatively fast performance versus other optimization algorithms.

2.4 | Baseline rotor configuration

All the studies presented in the current work use the International Energy Agency (IEA) Wind Task 37 land-based reference wind turbine, which is characterized by a nameplate power of 3.4 MW, a rotor diameter, D , of 130 m, and a hub height, HH , of 110 m.⁴³ Before beginning an optimization study designed to test the effect of changing blade design on plant power output, it was necessary to establish a performance baseline. Rather than re-optimizing the blade geometry within WindSE, the rated operating conditions were optimized such that the initial chord distribution generated the maximum power output for a single turbine and that any variation in chord resulted in diminished power performance. This baseline case uses the same rotor geometry, chord measurements, and airfoil polars published for the IEA Wind reference wind turbine,⁴³ but the tip-speed ratio has been increased from the rated value of 8.01 to 10.5 through the combined effects of lowering the wind speed at hub height to 7.5 m/s and increasing the rotor speed to 11.6 rpm. A summary of the reference conditions originally prescribed and the baseline values adopted by every case study in Section 3 are reported in Table 1.

This increased tip-speed ratio results in the *individual* output power being maximized only when using the baseline chord, even on a relatively coarse computational grid. Deviations in chord away from these dimensions, although potentially favorable for downstream turbines, result in reduced power for an individual turbine as predicted by blade element models.⁴⁴ The discrepancy in the optimal tip-speed ratio can be explained by the fundamentally different formulations present in the actuator line model used by WindSE and the conventional blade element momentum solver used to optimize the aerodynamic design of the wind turbine rotor.⁴³

3 | CASE STUDIES AND RESULTS

The three design optimization case studies presented each have a different objective function and slightly different problem setup to showcase different aspects of individual turbine blade design for plant controls. Nevertheless, each generally has a similar optimization problem formulation,

$$\begin{aligned}
 & \text{maximize} && \text{Wake deflection} && \text{or} \\
 & && \text{Vertical entrainment} && \text{or} \\
 & && \text{Plant power} && \\
 & \text{w.r.t.} && c_i && i = 1, \dots, N \\
 & \text{subject to} && c_{i,\min} \leq c_i \leq c_{i,\max} &&
 \end{aligned} \tag{14}$$

where the objective functions will be mathematically defined in each of the individual case study discussions and c_i are the design variables, which are the chord lengths at the $N = 10$ actuator line nodes along the blade span. The three blades adopt the same chord distributions. The upper and lower limits for the blade chord are derived from $2\times$ and $0.5\times$ offsets applied to the IEA 3.4-MW reference wind turbine chord profile, respectively, with a slight modification to allow additional design freedom at the narrow blade tip.

To verify the convergence of the actuator line implementation and unsteady solver used in the following case studies and to further support our discretization choices, a representative simulation of a single, yawed turbine was carried out using varying levels of temporal and spatial

TABLE 1 Simulation setup for the IEA 3.4-MW reference wind turbine

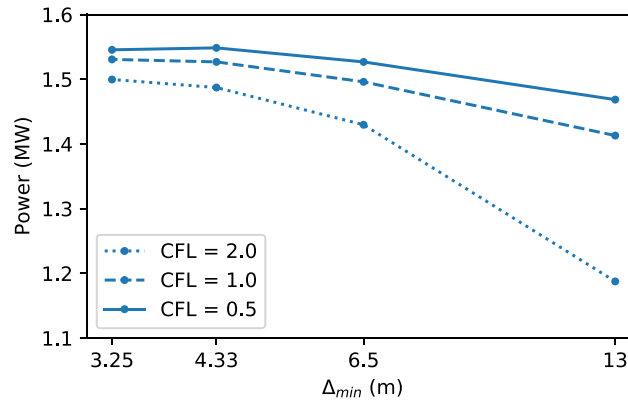
Quantity	Reference	Baseline
Rotor speed (rpm)	10.6	11.6
Tip-speed ratio (-)	8.0	10.5
Hub height wind speed (m/s)	9.0	7.5

Note: Reference corresponds to the rated conditions of the original design,⁴³ whereas Baseline corresponds to the values used by all case studies.

TABLE 2 The yawed, single-rotor power output and the required simulation wall-clock time as a function of varying temporal and spatial refinement

CFL	Turbine power (MW) and wall-clock time			
	$\Delta_{min} = 13.0$ m	$\Delta_{min} = 6.50$ m	$\Delta_{min} = 4.33$ m	$\Delta_{min} = 3.25$ m
2.0	1.19 (13.60 min)	1.43 (2.76 h)	1.49 (17.11 h)	1.50 (55.08 h)
1.0	1.41 (25.58 min)	1.50 (5.46 h)	1.53 (32.87 h)	1.53 (99.32 h)
0.5	1.47 (50.07 min)	1.53 (11.12 h)	1.55 (68.54 h)	1.55 (200.15 h)

Note: All times were computed to reflect single core performance.

**FIGURE 2** Visualizing the yawed, single-rotor power output as a function of varying temporal and spatial refinement

refinement. The turbine is operated at the baseline conditions shown in Table 1 with a yaw angle of $\theta = 20^\circ$. The average rotor power is calculated according to Equation (13) over an averaging window of [200s, 250s]. The results of this testing are compiled in Table 2 and visualized in Figure 2.

These values are obtained with a single pass of the forward solve only; that is, no optimization or alteration of baseline controls is carried out. As the value of Δ_{min} is reduced, the turbine power converges to a steady-state value ~ 1.55 MW. A similar convergence pattern emerges as the result of temporal refinement, with smaller timesteps resulting in a power output approaching ~ 1.55 MW. It is important to note the other effect of refinement, both temporally and spatially, which is the increased simulation time required. This increased time to solution means that the optimization algorithm is limited to fewer and fewer optimization iterations within the 10-day wall-clock limit imposed by the high performance computing (HPC) resources carrying out our studies. Since each optimization iteration often requires as many as 10 iterations of the forward solution, every effort is taken to minimize the time to solution through coarsening the temporal and spatial discretization while still preserving the fundamental physics required for a meaningful optimization. With this in mind, many of the following case studies adopt a CFL number of 2.0 with $\Delta_{min} = 6.50$ m, which maintains the physical fidelity to recover $>92\%$ of the converged objective function shown here while also being relatively quick to solve; even a study where each optimization iteration requires ~ 10 forward solves (~ 27.6 h per optimization iteration) can complete multiple optimization iterations and exhibit convergence within the 10-day limit. In the following case studies, these discretization effects will be periodically revisited in the form of problem-specific tests designed to verify the implementation of specific objective functions operating under specific conditions. However, this power-only study provides a good rule of thumb for developing further modifications and a measure of confidence that our unsteady solver and actuator line implementation propagate the crucial fluid pressure and velocity appropriately.

3.1 | Case Study 1: Maximizing wake deflection

The goal of Case Study 1 is to investigate the possibility of increasing the wake steering ability of a yawed turbine by optimizing blade geometry to maximize the lateral wake deflection. This extends the work done in a previous WindSE study to characterize the sensitivity of wake position with respect to changes in chord profile along the length of the blade.²⁶ Increasing the ability of an upstream turbine to deflect wake away from a downstream turbine may allow for the creation of more favorable downstream flow conditions, a corresponding increase in power output from waked turbines, or less power sacrificed from the upstream turbine.

3.1.1 | Simulation setup

Case Study 1 uses the baseline operating conditions listed in Table 1 and yaws the turbine at an angle of 20° . The computational domain employed in this study is shown in Figure 3.

The size of the domain is chosen to be sufficiently large to model the effects of actuator forces across the total permissible range of design values without incurring problematic edge effects while minimizing the total degrees of freedom and thereby simulation time for the reasons presented previously. This sizing process is carried out by iteratively reducing the separation distance between the turbine and the domain walls and observing the behavior of the objective function, where a large difference in objective function associated with a restricted domain indicates the introduction of non-physical edge effects. Not depicted in Figure 3 is the domain width in the Y-direction which is set to $4D$ where D is the diameter of the blade. The inflow velocity, U_{in} , is a steady-state logarithmic profile that satisfies the hub height velocity of 7.5 m/s. The objective function is measured on a YZ-plane located 5 rotor diameters downstream to showcase the results of enhanced steering at a representative location for an un-modeled, downstream turbine. The value of the objective function is given by a calculation of the wake centroid in the Y-direction:

$$y_c = \frac{1}{t_2 - t_1} \int_{t_1}^{t_2} \frac{\int_{x=5D} Y_i u_{wd} dA}{\int_{x=5D} u_{wd} dA} dt \quad (15)$$

where $5D$ is the downstream distance expressed in rotor diameters and the wake deficit magnitude, u_{wd} , is given by $u_{wd} = |u_{inflow} - u|$. Although there are plenty of metrics that can be used to approximate the location of a wake, this integral based method was the most amenable to how `dolphin-adjoint` handles computation of the derivatives with respect to the objective function, namely, that it maintains the connections between blocks as discussed in Figure 1.

Due to the unsteady nature of the flow, this objective function is evaluated as a time average using a temporal window that begins when the wake from the turbine is well established at the objective measurement plane (t_1) and ends when the average centroid location is suitably converged (t_2), namely, from 200 to 250 s. These values were informed by a combination of tests, one in which the forward problem was run for 350 s and a calculation of the running average was computed for various time windows, and one in which the standard deviation of the raw centroid data was computed to determine the point at which the statistical time-varying behavior became bounded. The results of this testing are found in Figure 4. The time window of 200 to 250 s seems to be a good approximation of the final converged average.

3.1.2 | Results

The chord profile that maximizes the Y-position of the wake centroid is shown in Figure 5. The optimized chord profile is significantly changed from the baseline definition. Interestingly, the change in chord, Δc , is not monotonic across the span of the blade; rather, it is increased in the first half, $y/R < 0.5$, decreased for the majority of the second half, $0.5 < y/R < 0.85$, and increased sharply at the blade tip, $0.85 < y/R$. It is also important to note that the control node at the tip of the blade remained stationary. This is because the model includes a tip-loss term that adjusts the forces at the tip of the blade, which remove the tip's contribution to the power objective function. This results zero sensitivity in the gradient with




FIGURE 3 The computational domain, simulation definitions, and the objective measurement location used in Case Study 1




FIGURE 4 A quick investigation on how the y -coordinate of the wake centroid at 5D downstream varies over time. Left: The line labeled “Full Avg” is the running averages starting at 0 s and ending at 350 s. The “50 s Avg” and “150 s Avg” are running average that both start at 200 s but end at 250 and 350 s, respectively. The full average line takes a significant time to recover from the first 100 s where the wake has not yet propagated downstream yet. Right: The standard deviation of y_c computed looking backward over a sliding, 50-s window. This statistic settles into a bounded behavior at the start of the 50- and 150-s measurements

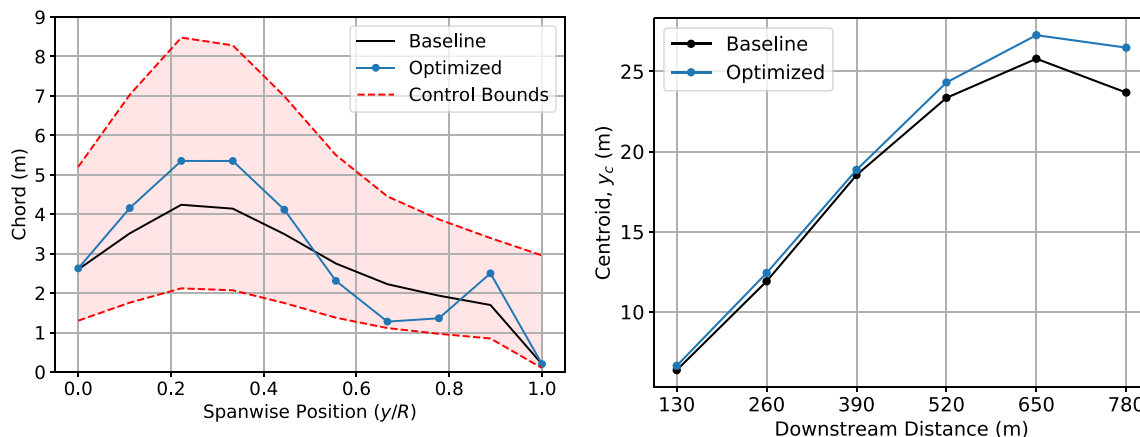


FIGURE 5 (left) The chord profile that maximizes lateral wake deflection is shown relative to the baseline chord profile for the IEA 3.4-MW reference wind turbine; the dotted lines represent upper and lower limits for the chord and the shaded region encompasses all permissible chord profiles. (right) The location of the centroid y -coordinate downstream from the turbine

respect to chord for the tip node, so even though the optimizer can adjust that node, the gradient will always be zero. Additionally, in Figure 5 is a plot of the wake centroid's y -coordinate as the wake progresses downstream. This shows that optimizing the centroid position at 5D also diverts the centroid at 4D and 6D as well.

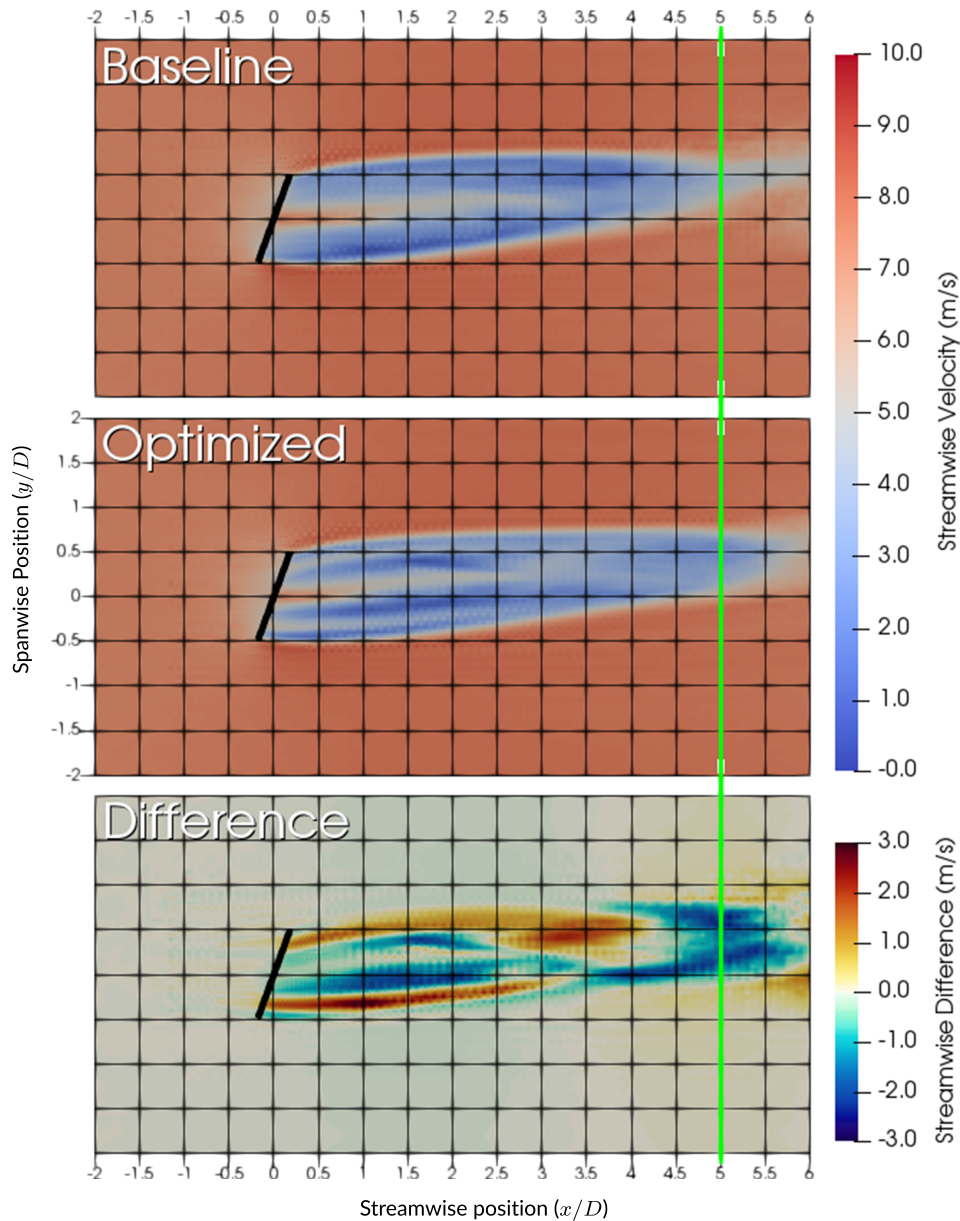
To increase the speed of the optimization, it was performed on a coarse grid. This results in approximate solutions that follow the trends of highly resolved simulation. Additionally, the centroid deflection for a simulation with a more resolved mesh was also calculated for comparison.

The centroid values for the baseline and optimized chord profiles for this coarse mesh are found in Table 3. The value of the objective function for the optimized chord profile on the coarse mesh is 27.48 m. Compared to the wake deflection obtained using the baseline chord, 25.51 m, the optimized chord profile attains a roughly 2 m of additional wake deflection, namely, a 7.7% increase. The wake generated by this altered blade geometry is shown alongside the unaltered wake in Figure 6. When the optimized chord is used in a more refined simulation, there is still an increase in deflection, but it is reduced to only 1.13 m.

To investigate the cause of this increase in deflection, Figures 6 and 7 show cross sections of the wake at the hub height and downstream through-plane slice, respectively. Based on the bottom plot in Figure 6, the modification to the baseline chord found in Figure 5 results in significant changes in the wake. To get a better look at this difference, Figure 7 presents three vertical slices at 1D, 3D, and 5D. Using the optimized chord, the wake develops a ring of high velocity, which can be seen in the 1D slice. This ring has a radius of approximately 50 m, which corresponds to when the optimized chord thickens at the tip, $y/R \approx 0.75$. At 3D, this ring is distorted by the counter-rotating vortices set up by

TABLE 3 The average y -coordinate of the wake centroid for the baseline and optimized chord profiles with two different mesh resolutions

DOFs	Δ_{min} (m)	CFL	Case	Deflection (m)	% Change
178 K	6.5	1.0	Baseline	25.51	
			Optimized	27.48	+7.72%
1152 K	3.25	1.0	Baseline	27.05	
			Optimized	28.18	+4.18%

**FIGURE 6** Hub height slices of the streamwise velocity averaged over the simulation's final 50 s in the XY-plane where the green line indicates where the objective (wake centroid) recorded; top, the wake generated by the baseline chord profile; middle, the wake generated by the optimized chord profile; bottom, $(\mathbf{u}_{optimized} - \mathbf{u}_{baseline})$. The yellow line indicated the location of centroid averaged over the 50-s window

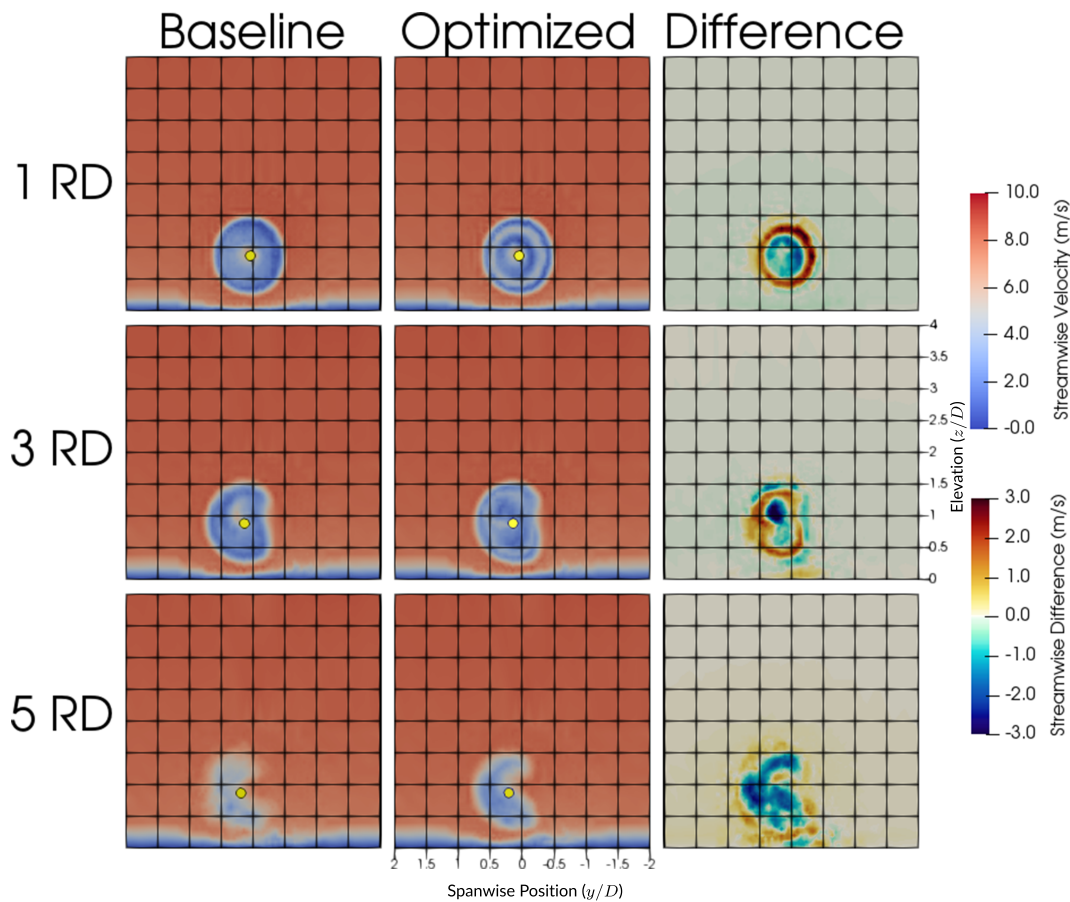


FIGURE 7 Slices of the streamwise velocity averaged over the simulation's final 50 s in the YZ-plane at 1, 3, 5, rotor diameter downstream of the turbine; left, the wake produced by the baseline chord profile; center, the wake produced by the optimized chord profile; right, ($\mathbf{u}_{\text{baseline}} - \mathbf{u}_{\text{optimized}}$). The yellow dots indicated the location of centroid averaged over the 50-s window

the yawed turbine entraining the higher velocities toward the center. This ring then dissipates downstream eventually leading to an outer “shell” of high speeds that surround the lower velocity “core” of the wake. This creates a large amount of off-center, slower velocity, shifting the centroid to the left.

However, even with these differences, the optimized chord results in a relatively modest 2 m of additional deflection. These results suggest only a moderate control authority over wake deflection, but that the chord can significantly affect the structure of the wake. The enhanced wake deflection is accomplished in part by a new induction profile that results in greater thrust and therefore deeper velocity deficits and a longer wake. The trade-offs between wake deflection and velocity deficit are addressed in the next two case studies in which the objective function are vertical entrainment and overall power production, respectively.

3.2 | Case Study 2: Maximizing vertical entrainment of kinetic energy

The interpretation of wake steering strategies originally emphasized lateral deflection of the wake away from downstream rotors; however, a deeper understanding of the flow control mechanisms involved in wake steering has emerged. Recent work^{17–20,45} has shown that counter-rotating vortices created by rotors operated in yawed conditions act to distort the wake into a “kidney bean” shape and also entrain higher momentum flow from aloft. This new understanding is consistent with our earlier results that showed limited ability to change the lateral wake deflection and suggest that kinetic energy entrainment is the more important flow control mechanism. In the limit of a very large wind farm, the kinetic energy to sustain turbine power generation is entrained from above rather than advected through the frontal inflow plane,^{21,22,46,47} so a flow control strategy that maximizes the vertical entrainment of kinetic energy would be important in this limiting case. Introducing a synthetic forcing term was found to have a significant impact on wind farm power output and kinetic energy entrainment,⁴⁸ suggesting that blades might be designed to enhance kinetic energy entrainment.

3.2.1 | Simulation setup

The goal of this study is to demonstrate a blade optimization that enhances the entrainment of kinetic energy in yawed conditions. The turbulence-induced vertical flux of mean kinetic energy is

$$\Phi(x, y, z) = \overline{u'w'} \bar{u} \quad (16)$$

where u , v , and w are the streamwise, spanwise, and vertical velocity components, the overbar denotes time averaging and the primes indicate fluctuating velocity components from the decomposition $u = \bar{u} + u'$. The net gain in kinetic energy available to be extracted due to vertical fluxes is

$$\mathcal{P}_\Phi(x, y) = |\Phi(x, y, z_h + D/2) - \Phi(x, y, z_h - D/2)| \quad (17)$$

Similarly, the lateral flux of mean kinetic energy is

$$\Psi(x, y, z) = \overline{u'v'} \bar{u} \quad (18)$$

and the net gain in kinetic energy available to be extracted from lateral fluxes is

$$\mathcal{P}_\Psi(x, z) = |\Psi(x, y_h + D/2, z) - \Psi(x, y_h - D/2, z)| \quad (19)$$

where z_h and y_h refer to the coordinates of the turbine hub in vertical and spanwise directions, respectively.

The computational domain for Case Study 2 is shown in Figure 8. We consider the entrainment in a region 6 rotor diameters long in the downstream direction and 3 rotor diameters wide in the spanwise direction and impose a fixed yaw angle $\theta = 20^\circ$. The optimization is performed with an aggressive timestep corresponding to CFL = 2 and simulation time of 300 s, but final reported power production and entrainment is calculated with a higher-fidelity simulation run at CFL = 1 for 900 s. The optimization is performed to maximize an objective function corresponding to the total integrated vertical entrainment $\int \int \mathcal{P}_\Phi(x, y) dx dy$, which results in substantial benefits to lateral entrainment as well.

3.2.2 | Results

The optimized blade profile for kinetic energy entrainment is shown in Figure 9. The chord is reduced in the inner third of the blade span and significantly increased in the outer half of the blade span. As in the previous case, the chord of the outermost node is unchanged because the tip-loss correction removes any sensitivity of the objective function to changes at that location. Moving the thrust outboard in this manner produces stronger mixing in the shear layer of the wake and also enhances the kinetic energy flux. Figure 10 shows the lateral entrainment \mathcal{P}_Ψ and vertical




FIGURE 8 Conceptual schematic of the vertical entrainment objective function, where increased mixing due to controls applied to the modeled turbine results in faster wake dissipation and a more energized flow at the location of an un-modeled downstream turbine yielding a net increase in power output

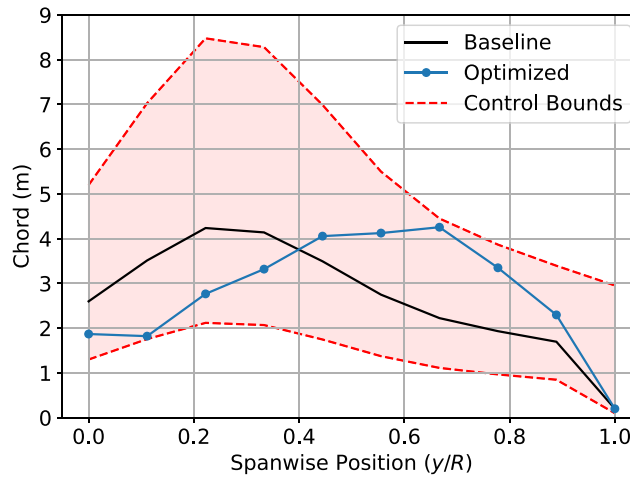


FIGURE 9 The optimized chord profile for maximizing vertical entrainment of kinetic energy




FIGURE 10 Lateral and vertical kinetic energy entrainment from the baseline rotor design

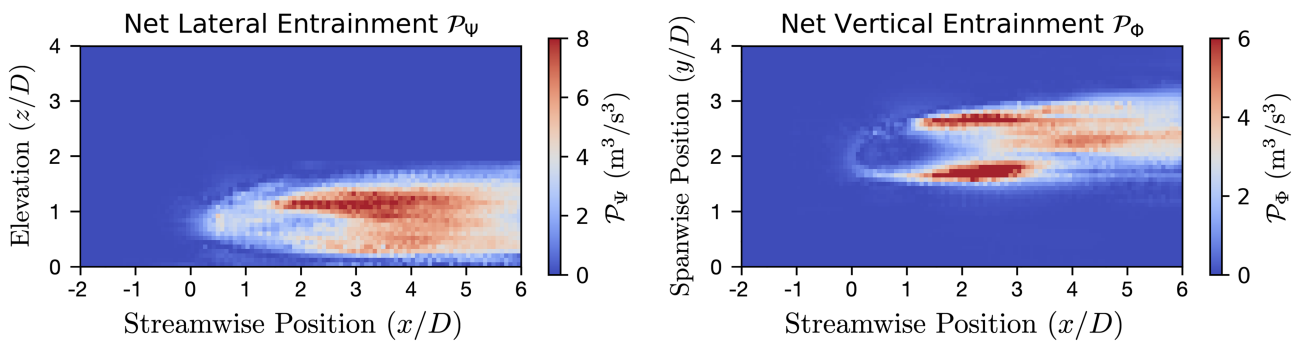


FIGURE 11 Lateral and vertical kinetic energy entrainment from the optimized rotor design

TABLE 4 Entrainment-based optimization results for the $\Delta_{min} = 6.5$ m

Case	Vertical entrainment \mathcal{P}_ϕ (m^3/s^3)	Lateral entrainment \mathcal{P}_ψ (m^3/s^3)	Power (MW)
Baseline profile	960.9	692.6	1.87
Entrainment optimized	3955.0	5741.7	1.61

entrainment for the baseline rotor profile and Figure 11 shows similar plots for the optimized rotor. The additional outboard thrust produces considerably stronger counter-rotating vortex pairs that result in substantially more entrainment. Although our optimization targeted vertical entrainment, we found additional benefits to the lateral kinetic energy fluxes. Table 4 summarizes the performance of the optimized and baseline rotor

configurations. The optimized rotor entrains roughly $4\times$ more kinetic energy in the vertical direction and $8\times$ more kinetic energy in the horizontal direction than the baseline rotor configuration. However, this comes with a penalty of a 13.9% decrease in power output. It should be noted that the specific improvements relayed here are potentially dependent on the mesh resolution and chosen volume of integration for the objective value.

The optimized rotor also enhances wake recovery rates. The increased thrust results in deeper velocity deficits in the near wake, but the entrained kinetic energy results in greater wake recovery downstream. This can be seen in Figure 12 which shows the difference in flow fields between the optimized and baseline rotors. The optimized rotor results in faster wake velocities after about 2 rotor diameters downstream, with gains especially pronounced at the top of the rotor and, to a lesser degree, at hub height. We note that the resolution demands for the fluxes are greater than for the mean velocity profile, and a full uncertainty quantification study of these results would be warranted in future work.

3.3 | Case Study 3: Maximizing power output

In Case Study 3, an optimization using a similar set of design variables is performed to maximize the cumulative power output of a two-turbine farm. In Case Studies 1 and 2, the wake deflection and vertical entrainment of kinetic energy objective functions are presented as proxies for this general goal, as moving the wake or energizing the flow from above to interact more favorably with a fictitious rotor located 5 diameters downstream seems an intuitive way to improve power output from the waked turbine. However, using Case Study 1 as an example, Figures 6 and 7 show that the blade changes used to achieve this enhanced deflection produce additional flow effects not accounted for with a simple translational shift of the low-energy wake. The effect of these changes on a downstream turbine's power output is not readily apparent, nor is the effect of potentially detrimental blade changes made to the leading turbine with respect to its own power production. A cumulative power objective function takes both these complicating factors into account by calculating and summing the individual powers using Equation (13). This allows for the possibility of an altruistic optimization in the sense that a leading turbine may do something individually unfavorable from a power output perspective to enable a net gain via the generation of a very favorable downstream flow.

3.3.1 | Simulation setup

The extended computational domain required to capture two turbines is shown in Figure 13.




FIGURE 12 Differences in the wake show that the rotor optimized for kinetic energy entrainment produces greater deficits in the near wake, but the added kinetic energy causes greater wake recovery, especially at hub height and the top of the rotor plane




FIGURE 13 Top-down view of the simulation design for Case Study 3, designed to increase the combined power output of turbines A and B through design change applied on A

With the few noteworthy differences outlined below, Case Study 3 adopts the same problem definition as presented for Case Study 1. The spacing between turbines is again $5D$, consistent with the measurement plane for the fictitious turbine layout in Case 1, but because the downstream turbine must now be modeled for an accurate power calculation, the outflow region is appropriately extended. Turbine A is operating in a yawed configuration with $\theta = 20^\circ$ while turbine B is operating in an un-yawed configuration, that is, $\theta = 0^\circ$. The design variables for this study are again the chord lengths defined along the blade axis, one per each actuator node, with the same constraints as described above. For simplicity, these design variables are only applied to turbine A; turbine B uses the baseline chord profile. Although the inverse case—optimizing the chord profile for turbine B to maximize output power under waked conditions—is also a study of some interest, the experience generated from controlling wakes in Case Study 1 makes turbine A the natural choice for this extended study. Additionally, the optimization was restricted to turbine A in order to isolate any changes in turbine B to only wake effects. The objective function in this study is the cumulative power, that is, $W_A + W_B$, where no weighting or bias is applied to the output of either rotor. The value of this objective function is averaged in time from the moment the wake from turbine A is well established and interacting with turbine B, 200 s, to the point where the time-averaged power output produced by the now-waked state of turbine B has converged to a near-steady state, 250 s.

3.3.2 | Results

The results of optimizing the upstream turbine's chord profile for cumulative power in this two-turbine configuration are shown in Figure 14.

The optimized chord profile for turbine A employs a slightly increased chord length approaching the blade tip, with a maximum increase of $\sim 24\%$ above the baseline chord at the $N - 1$ actuator point. It is worth noting that for the first half of the blade span, $y/R < 0.5$, only minor chord increases are made with none greater than 3%. As expected, the chord changes are much less dramatic when compared to the optimized profiles generated in Case Studies 1 and 2. This is due to the objective function being directly related to the calculated rotor power, which inhibits the sort of large deviations away from the optimized IEA Wind 3.4-MW reference wind turbine geometry that would lead to diminished power output. However, for a two-turbine system in which it is the *cumulative* power that is of interest, turbine A does have a measure of freedom to affect changes if it means that its individual losses are more than offset by gains from turbine B, that is, $(\Delta W_A + \Delta W_B) > 0$. Also shown in Figure 14 is a manually prescribed test case designed to compare the optimized result with the performance of a nearby solution. “Optimized++,” where the “++” is used in the spirit of an increment operator, is the chord profile that results from adding an additional +50% change on top of the already optimized profile. Conceptually, it represents moving slightly further along the gradient exploited by the optimized result as a simple way to test the hypothesis that further movement along this direction no longer yields improvements. The individual and cumulative power outputs for this study are summarized in Table 5.

The top portion of Table 5 summarizes the values obtained using the coarse timestep size employed by the actual optimization solution, $CFL = 2.0$. The bottom portion of Table 5 shows the values obtained as part of a post-processing exercise to compare the results from the coarse temporal discretization used by the optimizer with a more finely discretized solution at $CFL = 0.5$. Beginning with the results from the optimization at $CFL = 2.0$, the optimized cumulative power output for this two-turbine system is 2.289 MW. When compared to the baseline power output, 2.238 MW, this means the optimized profile has affected a 2.27% increase in the power output from this two-turbine farm. It is important to highlight that in the case with $CFL = 2.0$, turbine A makes these changes at the expense of its own power, -0.53% , as deviation from the baseline




FIGURE 14 The optimized chord profile for upstream turbine A that maximizes power in a two-turbine configuration is marked with blue circles; a synthesized test case designed to test the performance of a nearby solution, “Optimized++,” is marked with orange triangles

TABLE 5 The average power output from upstream turbine A, downstream turbine B, and the total when using the baseline rotor in an unyawed configuration, and the baseline, optimized, and Optimized++ chord profiles at turbine A in a yawed configuration using a mesh with $\Delta_{min} = 6.5$ m; percent changes, shown in brackets, are relative to the baseline value for each case

Δt	Case	Turbine A (MW)	Turbine B (MW)	Total (MW)
CFL = 2.0	No yaw	1.569	0.132	1.701
	Baseline	1.429	0.808	2.238
	Optimized	1.422 (−0.53%)	0.867 (+7.23%)	2.289 (+2.27%)
	Optimized++	1.401 (−1.98%)	0.882 (+9.06%)	2.283 (+2.01%)
CFL = 0.5	No yaw	1.664	0.072	1.736
	Baseline	1.525	0.860	2.386
	Optimized	1.528 (+0.17%)	0.944 (+9.67%)	2.471 (+3.59%)
	Optimized++	1.510 (−0.98%)	0.964 (+12.07%)	2.474 (+3.73%)

chord profile is generally unfavorable from the perspective of a greedy optimization; it is the presence of a relatively larger improvement in the flow conditions at turbine B and the notion of a cumulative objective function that serve to justify the performance penalty at turbine A. When testing the manually prescribed Optimized++ case with CFL = 2.0, turbine A's output is further diminished but without a sufficient boost to turbine B, resulting in slightly diminished performance overall. This provides confidence that the optimizer converged to an appropriate result and did not neglect further improvements by not further pursuing this line of design variable change.

The results from the smaller timestep test, CFL = 0.5, show that the optimized results obtained with the coarse timestepping are still applicable to a more finely resolved problem. In this case, testing the optimized blade shape obtained from the coarse study yields a small increase in turbine A's power and a significant increase in the power produced from turbine B. When further extending this solution to the Optimized++ case, the larger chord follows the trend of reducing turbine A's output and further enhancing turbine B's output. The additional power increase compared to the optimized case is quite small, 0.003 MW, suggesting that although an optimization run with a smaller timestep might have pursued this gradient slightly further and thus converged to an optimal chord profile slightly larger than the one obtained with a coarse timestep, the coarse optimization converged very near the same maximum configuration being identified by the more finely discretized study. This supports the idea that optimization studies, which principally rely on trends in the data as opposed to actual magnitudes, can be run using a coarser temporal discretization and still achieve results that match well with a more finely resolved view of the problem. Caution should still be used and spot-checks with finer discretizations are advised as evidenced by the turbine A results where the Optimized solution at the coarse timestep showed a relative power decrease of −0.53% but a +0.17% increase when evaluated at the finer timestep. Nevertheless, when considering the computational time and memory savings associated with using the larger timestep size, $\sim 4\times$ improvement, a coarse optimization followed by a more refined verification may offer significant benefits even with the added verification step.

An instantaneous snapshot taken near the middle of the objective averaging window of the wakes generated by both the baseline and optimized chord designs is shown in Figure 15.

Comparing the left and middle frames of Figure 15, the optimized chord profile with its thicker blade tip seems to generate additional mixing during the interaction with turbine B as evidenced by the increased presence of vortical structures identified with Q-criterion isosurfaces. When looking at the differences between both streamwise velocity fields, it is clear that the increased chord length near the blade tip results in increased thrust and a corresponding slowdown of flow behind the tip-swept area, $u_x(\text{optimized}) - u_x(\text{baseline}) < 0$. As seen in the vorticity, this increased induction seems to result in a stronger mixing effect ahead of turbine B and a region of transiently accelerated flow available to the waked rotor, $u_x(\text{optimized}) - u_x(\text{baseline}) > 0$. Although this phenomenon is unsteady in nature, the optimized chord profile at turbine A seems to take advantage of the effect over a relatively long time window to attain the improvements summarized in Table 5.

To determine the effect of turbine A's yaw angle, this optimization study was repeated holding all previously discussed parameters constant with the exception of θ , which was varied to span the range [10, 25]. The relatively coarse timestep size, CFL = 2.0, was again employed for computational speed. The chord profiles resulting from these optimizations are shown in Figure 16.

There are significant similarities between the optimal chord profiles for turbine A when $\theta > 10^\circ$, indicating that the design trend of increasing chord length by $\sim 24\%$ at the blade tip is relatively robust to upstream yaw angle. It is evident that the optimized chord profile obtained when turbine A is yawed at $\theta = 17.5^\circ$ shows the smallest deviation from the baseline profile, while increasing yaw, for example, to $\theta \geq 20.0^\circ$, and decreasing yaw, for example, to $\theta \leq 15.0^\circ$, result in larger deviations in the chord profiles away from the baseline profile. However, it is worth reiterating that with the exception of $\theta = 10.0^\circ$, these changes with respect to yaw are small and thus, a single, representative chord profile could be selected or averaged from this set to produce a blade that behaves well across a range of upstream yaw conditions. The cumulative power outputs resulting from these optimized profiles are summarized alongside their baseline outputs in Table 6.

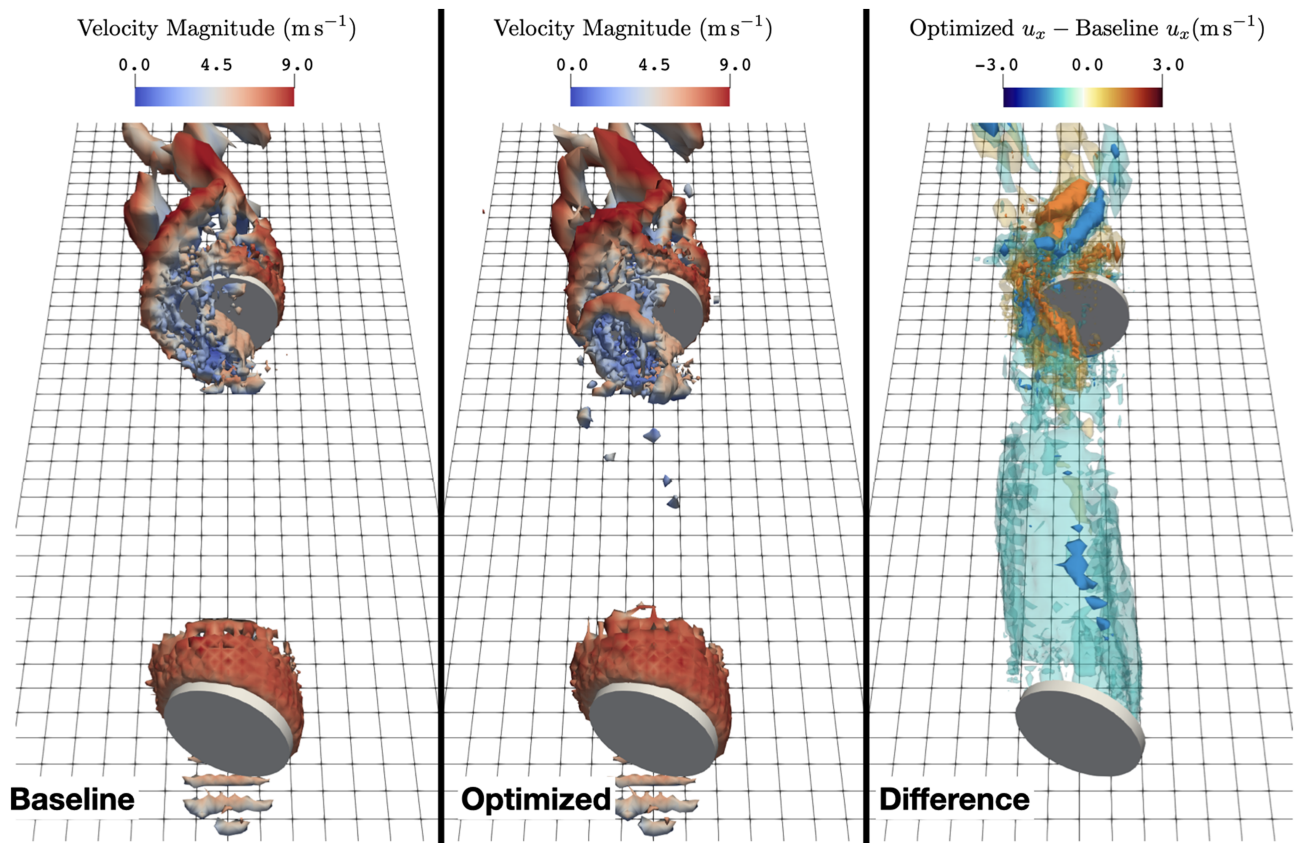


FIGURE 15 The wake generated by the baseline chord profile (left), optimized chord profile (middle), and the streamwise velocity differences established by both rotors (right); the baseline and optimized views depict an isosurface of Q-criterion colored by velocity and the difference view shows isosurfaces of streamwise velocity change, $u_x(\text{optimized}) - u_x(\text{baseline})$, where a relative increase in streamwise velocity ahead of the second turbine due to increased entrainment can be seen




FIGURE 16 The result of optimizing chord profiles for total power output as a function of the yaw angle of upstream turbine A

Interestingly, when $\theta = 17.5^\circ$, which corresponds to the smallest increase in chord profile shown in Figure 16, the optimized rotor offers the smallest relative improvement in terms of the total power output. As turbine A is yawed either more or less aggressively, the power increase attainable when using the optimized rotor increases. This suggests a region of yawed operation around 17.5° in which the baseline IEA 3.4-MW reference wind turbine chord profile is near optimal and that altering the yaw angle of turbine A away from this value brings about slightly more

TABLE 6 The total power, $W_A + W_B$, obtained from the two-turbine system using the baseline chord profile and the optimized chord profile with $CFL = 2.0$

Yaw angle, θ ($^\circ$)	10.0	12.5	15.0	17.5	20.0	25.0
Baseline (MW)	1.839	1.946	2.030	2.154	2.238	2.379
Optimized (MW)	1.917	1.985	2.085	2.188	2.289	2.436
Difference (MW)	0.078 (+4.26%)	0.039 (+1.98%)	0.055 (+2.71%)	0.034 (+1.58%)	0.051 (+2.27%)	0.057 (+2.38%)

significant chord changes with slightly greater potential benefit. The smallest angle tested, $\theta = 10^\circ$, results in a chord profile that is markedly larger than the others, albeit with the same trend of increasing at the blade tip. This sudden change may mark the transition into a different regime of optimal blade shapes when the leading turbine is yawed less than a certain threshold, but further study is needed to test this hypothesis.

4 | CONCLUSIONS

Overall, this study successfully demonstrated multiple possibilities for optimizing one or more turbine rotors using different objective functions each formulated to target enhancing wind plant controls in slightly different ways. In Case Study 1, the goal is to maximize wake deflection during yawed operation by controlling spanwise chord profile. The objective function measuring progress toward this goal attempts to characterize the position of the wake centroid through a weighted average of the magnitude of the wake deficit. The motivation for this objective function is the underlying assumption that steering the wake to generate less interference with a downstream turbine will be favorable for the power production of the downstream turbine. The optimized chord profile provides an additional deflection of roughly 1–2 m, but the changes in chord required to achieve this enhancement and the corresponding changes to the flow field are surprisingly complex. Without a calculation of power in the objective function, the optimizer is relatively uninhibited to explore the total design space; this results in the optimizer selecting a non-intuitive blade design that exploits the formulation of the centroid objective function to alternately accelerate/decelerate the flow to shift the centroid as calculated by Equation (15). We also note that the relatively small improvements in deflection suggest lower control authority over wake deflection and motivate consideration of other flow control mechanisms.

Case Study 2 was motivated by observations that yawed conditions during wake steering create a counter-rotating vortex pair which serves to distort the wake and entrain kinetic energy. The optimizer was able to improve kinetic energy entrainment in the vertical direction by $4\times$ and in the lateral direction by over $8\times$, suggesting much greater control authority for entrainment than wake deflection. This was achieved by increasing the chord length in the outer 50% of the blade and reducing it in the inner third. Although this profile is beneficial from a plant controls perspective, the structural ramifications of such a design are not considered here. The higher thrust creates deeper velocity deficits in the near wake, but the enhanced entrainment results in greater wake recovery beyond 2 rotor diameters downstream. The optimizer rotor in Case Study 2 produced roughly 14% less power than the baseline rotor, naturally leading to questions about net power gain which were addressed in Case Study 3.

Case Study 3 considered an objective function corresponding to the total power produced from two turbines by optimizing the rotor in the upwind yawed turbine. This new objective function removes some of the flexibility in the design space as the IEA 3.4-MW reference wind turbine chord profile is already near optimal with respect to power production and any significant deviation away from this baseline is likely to incur reductions in power. For the two-turbine farm in Case Study 3 with $CFL = 2.0$, the optimized upstream turbine lengthened its chord near the blade tip to increase the waked turbine's power output by 7.23% at the cost of a 0.53% decrease in its own power, resulting in a net gain of 2.27% for the plant when using the low fidelity coarse timestep and a 3.59% net power increase when using a finer timestep with $CFL = 0.5$. It is expected that the relatively modest improvements attainable from a two-turbine system like the one studied here will be increasingly beneficial when moving to optimizations of larger farms with much deeper blockage regions where the fluid interactions produced by the leading turbine will likely have cascading effects. This more complicated optimization problem with second-level, third-level, and so forth, turbine interactions is the focus of future work, and strong motivation for WindSE to add support for parallel optimization studies that can mitigate the cost of increasing degrees of freedom—both from the increased size of the computational grid and the increased length of the adjoint tape necessary to record flow through on a larger farm—by taking advantage of distributed computing resources.

These case studies demonstrate the power of an adjoint-enabled actuator line CFD model that can flexibly optimize a variety of flow control objectives. We find convincing evidence that blade design choices can enhance plant-level flow control strategies to deflect wakes, entrain momentum, and increase power production. Although blade structural and manufacturing considerations are not addressed here, this new design tool provides a pathway for exploring many new innovations in blade design and wind plant controls.

Finally, it is important to properly contextualize the results of this paper. While Table 2 and Figure 2 show the overall convergence trends for the power calculation with respect to decreasing CFL and mesh spacing, it is important to note that due to the computational limitations outlined

in Section 3, the studies performed in this paper could not be carried out at the highest, desired fidelity. This means that, although the quantitative results may not perfectly reflect the ground truth, the trends that these quantitative results expose are worth considering, for example, how increasing the chord near the tip can increase overall power production. Supporting higher-fidelity simulation and optimization studies in the future is the motivation for ongoing research to improve computational performance and scalability.

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 was 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.

This research was performed using computational resources sponsored by the Department of Energy's Office of Energy Efficiency and Renewable Energy and located at the National Renewable Energy Laboratory.

The authors would like to acknowledge the help of Luis Martinez Tossas in setting up the actuator line model within WindSE.

CONFLICT OF INTERESTS

The authors declare no potential conflict of interests.

PEER REVIEW

The peer review history for this article is available at <https://publons.com/publon/10.1002/we.2699>.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from the corresponding author upon reasonable request. The open-source software used to perform the optimizations in this article is publicly available at <https://github.com/NREL/WindSE>.

ORCID

Jeffery Allen  <https://orcid.org/0000-0003-2815-6835>

Ethan Young  <https://orcid.org/0000-0003-1106-7406>

Pietro Bortolotti  <https://orcid.org/0000-0002-8515-9488>

Garrett Barter  <https://orcid.org/0000-0001-7757-8853>

REFERENCES

1. Veers P, Dykes K, Lantz E, et al. Grand challenges in the science of wind energy. *Science*. 2019;366(6464):eaau2027.
2. Steinbuch M, de Boer WW, Bosgra OH, Peters SAWM, Ploeg J. Optimal control of wind power plants. *J Wind Eng Ind Aerodyn*. 1988;27(1):237-246. [https://doi.org/10.1016/0167-6105\(88\)90039-6](https://doi.org/10.1016/0167-6105(88)90039-6)
3. Johnson KE, Thomas N. Wind farm control: addressing the aerodynamic interaction among wind turbines. In: 2009 American Control Conference; 2009; St. Louis, MO, USA:2104-2109.
4. Bitar E, Seiler P. Coordinated control of a wind turbine array for power maximization. In: 2013 American Control Conference. Washington, DC, USA; 2013:2898-2904.
5. Medici D, Dahlberg JAA. Potential improvement of wind turbine array efficiency by active wake control (AWC). In: Proceedings of European Wind Energy Conference. Madrid, Spain: 2003:65-84.
6. Medici D, Alfredsson PH. Measurements on a wind turbine wake: 3D effects and bluff body vortex shedding. *Wind Energy*. 2006;9(3):219-236. <https://doi.org/10.1002/we.156>
7. Wagenaar JW, Machielse L, Schepers J. Controlling wind in ECN's scaled wind farm. *Proc. Europe Premier Wind Energy Event*. 2012;1:685-694.
8. Park J, Kwon S, Law KH. Wind farm power maximization based on a cooperative static game approach. In: Sodano H, ed. *Active and Passive Smart Structures and Integrated Systems 2013*, Vol. 8688: International Society for Optics and Photonics. SPIE; 2013:204-218. <https://doi.org/10.1117/12.2009618>
9. Fleming PA, Ning A, Gebraad PMO, Dykes K. Wind plant system engineering through optimization of layout and yaw control. *Wind Energy*. 2016; 19(2):329-344. <https://doi.org/10.1002/we.1836>
10. Gebraad PMO, Teeuwisse FW, van Wingerden JW, Fleming PA, Ruben SD, Marden JR, Pao LY. Wind plant power optimization through yaw control using a parametric model for wake effects—a CFD simulation study. *Wind Energy*. 2016;19(1):95-114. <https://doi.org/10.1002/we.1822>
11. Campagnolo F, Petrović V, Schreiber J, Nanos EM, Croce A, Bottasso CL. Wind tunnel testing of a closed-loop wake deflection controller for wind farm power maximization. *J Phys Conf Ser*. 2016;753:032006. Publisher: IOP Publishing.
12. Fleming P, Annoni J, Shah JJ, et al. Field test of wake steering at an offshore wind farm. *Wind Energy Science*. 2017;2(1):229-239. <https://doi.org/10.5194/wes-2-229-2007>

13. Fleming P, King J, Dykes K, et al. Initial results from a field campaign of wake steering applied at a commercial wind farm—part 1. *Wind Energy Sci*. 2019;4(2):273-285. <https://doi.org/10.5194/wes-4-273-2019>
14. Fleming P, King J, Simley E, et al. Continued results from a field campaign of wake steering applied at a commercial wind farm—part 2. *Wind Energy Sci*. 2020;5(3):945-958. <https://doi.org/10.5194/wes-5-945-2020>
15. Doekemeijer BM, Kern S, Maturu S, et al. Field experiment for open-loop yaw-based wake steering at a commercial onshore wind farm in Italy. *Wind Energy Sci Discuss*. 2020;2020:1-22.
16. Bastankhah M, Porté-Agel F. Experimental and theoretical study of wind turbine wakes in yawed conditions. *J Fluid Mech*. 2016;806:506-541. Publisher: Cambridge University Press. <https://doi.org/10.1017/jfm.2016.595>
17. Howland MF, Bossuyt J, Martínez-Tossas LA, Meyers J, Meneveau C. Wake structure in actuator disk models of wind turbines in yaw under uniform inflow conditions. *J Renew Sustain Energy*. 2016;8(4):043301.
18. Howland MF, Lele SK, Dabiri JO. Wind farm power optimization through wake steering. *Proc Natl Acad Sci*. 2019;116(29):14,495-14,500.
19. Martínez-Tossas LA, Annoni J, Fleming PA, Churchfield MJ. The aerodynamics of the curled wake: a simplified model in view of flow control. *Wind Energy Sci (Online)*. 2019;4(NREL/JA-5000-73451).
20. Fleming P, Annoni J, Shah JJ, Wang L, Ananthan S, Zhang Z, Hutchings K, Wang P, Chen W, Chen L. Field test of wake steering at an offshore wind farm. *Wind Energy Sci Discuss*. 2017;2:229-239.
21. Cal RB, Lebrón J, Castillo L, Kang HS, Meneveau C. Experimental study of the horizontally averaged flow structure in a model wind-turbine array boundary layer. *J Renew Sustain Energy*. 2010;2(1):013106.
22. Calaf M, Meneveau C, Meyers J. Large eddy simulation study of fully developed wind-turbine array boundary layers. *Phys Fluids*. 2010;22(1):015110.
23. Bortolotti P, Bottasso CL, Croce A. Combined preliminary—detailed design of wind turbines. *Wind Energy Sci*. 2016;1(1):71-88.
24. Stanley APJ, Ning A. Coupled wind turbine design and layout optimization with nonhomogeneous wind turbines. *Wind Energy Sci*. 2019;4(1):99-114. <https://doi.org/10.5194/wes-4-99-2019>
25. Energy GERenewable. GE renewable energy launches cypress onshore wind platform, designed to grow, adapt, thrive and lower the cost of electricity for customers. <https://www.ge.com/news/press-releases/ge-renewable-energy-launches-cypress-onshore-wind-platform-designed-grow-adapt;2018>.
26. Young E, Allen J, Bortolotti P, King R, Barter G. An unsteady actuator line solver to enable adjoint sensitivity studies. *J Phys Conf Ser*. 2020;1618. Publisher: IOP Publishing.
27. Chorin A. A numerical method for solving incompressible viscous flow problems. *J Comput Phys*. 1967;2(1):12-26.
28. Kim J, Moin P. Application of a fractional-step method to incompressible Navier-Stokes equations. *J Comput Phys*. 1985;59(2):308-323.
29. Sørensen JN, Shen WZ. Numerical modeling of wind turbine wakes. *J Fluids Eng*. 2002;124:393-399.
30. Churchfield M, Lee S, Moriarty P, Martínez L, Leonardi S, Vijayakumar G, Brasseur J. A large-eddy simulation of wind-plant aerodynamics. In: Proceedings of the 50th AIAA Aerospace Sciences Meeting; 2012; Nashville, Tennessee.
31. King RN, Dykes K, Graf P, Hamlington PE. Optimization of wind plant layouts using an adjoint approach. *Wind Energy Sci*. 2017;2(1):115-131. <https://doi.org/10.5194/wes-2-115-2017>
32. King RN, Adcock C, Annoni J, Dykes KL. Data-driven machine learning for wind plant flow modeling. *J Phys Conf Ser*. 2018;1037:072004.
33. King R, Glaws A, Geraci G, Eldred MS. A probabilistic approach to estimating wind farm annual energy production with Bayesian quadrature The organization; 2020:1951.
34. Allen J, King R, Barter G. Wind farm simulation and layout optimization in complex terrain. *J Phys Conf Ser*. 2020;1452:012066.
35. Alnæs MS, Blechta J, Hake J, et al. The FEniCS project version 1.5. *Archive of Numerical Softw*. 2015;3:100. <https://doi.org/10.11588/ans.2015.100.20553>
36. Logg A, Mardal K-A, Wells GN, et al. *Automated Solution of Differential Equations by the Finite Element Method*: Springer; 2012.
37. Logg A, Wells GN, Hake J. Dolfin: a C++/Python finite element library. In: Logg A, Mardal K-A, Wells GN, eds. *Automated Solution of Differential Equations by the Finite Element Method, Volume 84 of Lecture Notes in Computational Science and Engineering*. Springer; 2012.
38. Logg A, Wells GN. Dolfin: automated finite element computing. *ACM Transactions on Mathematical Software*. 2010;37(2).
39. Balay S, Abhyankar S, Adams MF, et al. PETSc users manual. ANL-95/11—Revision 3.13, Argonne National Laboratory; 2020. <https://www.mcs.anl.gov/petsc>
40. Mitusch SK, Funke SW, Dokken JS. dolfin-adjoint 2018.1: automated adjoints for FEniCS and Firedrake. *J Open Source Softw*. 2019;4(38):1292. <https://doi.org/10.21105/joss.01292>
41. Funke SW, Farrell PE. A framework for automated PDE-constrained optimisation. *CoRR*. 2013;abs/1302.3894. <http://arxiv.org/abs/1302.3894>
42. Virtanen P, Gommers R, Oliphant TE, et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods*. 2020;17:261-272.
43. Bortolotti P, Tarres HC, Dykes K, Merz K, Sethuraman L, Verelst D, Zahle F. IEA Wind Task 37 on systems engineering in wind energy—WP2.1 Reference wind turbines, National Renewable Energy Lab Report NREL/TP-73492; 2019.
44. Ning SA. A simple solution method for the blade element momentum equations with guaranteed convergence. *Wind Energy*. 2014;17(9):1327-1345.
45. Fleming P, Annoni J, Churchfield M, Martínez-Tossas LA, Gruchalla K, Lawson M, Moriarty P. A simulation study demonstrating the importance of large-scale trailing vortices in wake steering. *Wind Energy Sci Discuss*. 2018;3:243-255.
46. Stevens RichardJAM, Meneveau C. Flow structure and turbulence in wind farms. *Ann. Rev Fluid Mech*. 2017;49:311-339.
47. VerHulst C, Meneveau C. Large eddy simulation study of the kinetic energy entrainment by energetic turbulent flow structures in large wind farms. *Phys Fluids*. 2014;26(2):025113.
48. VerHulst C, Meneveau C. Altering kinetic energy entrainment in large eddy simulations of large wind farms using unconventional wind turbine actuator forcing. *Energies*. 2015;8(1):370-386.

How to cite this article: Allen J, Young E, Bortolotti P, King R, Barter G. Blade planform design optimization to enhance turbine wake control. *Wind Energy*. 2022;25(5):811-830. doi:10.1002/we.2699