An Unsteady Actuator Line Solver to Enable Adjoint Sensitivity Studies for Wake Steering

Ethan Young[1], Jeffery Allen[1], Pietro Bortolotti[1], Ryan King[1], and Garrett Barter[1]

[1]National Renewable Energy Laboratory, Golden, CO, USA

E-mail: ethan.young@nrel.gov

Abstract. This study demonstrates the sensitivity of wind turbine wake steering performance to blade design. An actuator line model was implemented within an unsteady adjoint solver that enables efficient execution of gradient-based optimization and sensitivity studies. After first confirming the feasibility of wake steering by controlling actuator line chord profiles and formulating a suitable objective function for wake position, a sensitivity study was conducted to determine the relative importance of chord length as a function of spanwise position on the resulting turbine wake deflection. The results presented here support the idea that blade design choices play a role in wake control. In a larger context, this study demonstrates a computational framework in which turbine and blade designs can be studied at the individual and farm-wide level to enhance wind plant controllability and manage power output.

1. Introduction

When planning a new wind farm, one of the major considerations is how to optimize turbine design, controls, and farm layout to maximize expected power output. A variety of computational models exist to solve such problems, but traditionally, engineering flow models have simplified potentially important flow features through linearization or analytical wake superposition techniques in the interest of computational speed. To address these shortcomings, the wind energy community increasingly employs higher-fidelity computational fluid dynamics (CFD) methods in optimization studies. The National Renewable Energy Laboratory (NREL) has shown recent success in this area with the development of WindSE [1], an open source tool that uses an adjoint Reynolds-averaged Navier-Stokes (RANS) formulation for optimization with the ability to capture terrain-induced effects [2] [3]. In the present work, we demonstrate an unsteady flow solver and higher-fidelity turbine representations using an actuator line model. When implemented within WindSE's analytical gradient and adjoint capability, a new regime of problems opens up, including the optimization of turbine blade characteristics for wake steering, time varying plant controls, and aerodynamic loads mitigation.

The long-term objective for this project is to develop a complete software environment for the optimization of wind farms through CFD simulations. This software environment, WindSE, allows the user to simulate both steady and unsteady flows around wind turbines in complex terrain. These capabilities allow WindSE users to optimize control variables like yaw angle or axial induction and farm-wide parameters like turbine location to harvest peak power under the specified conditions. The former cases are both examples of the idea of controllability, in

which parameters describing turbine features can be studied and tuned to optimize performance or achieve a desired aerodynamic effect. The focus of this work is the recent addition of the unsteady solver, the development of higher-fidelity, adjoint-preserving turbine representations, and a gradient-based local sensitivity study on the effect of turbine blade characteristics on wake steering efficacy. This sensitivity study provides insights on the link between blade design parameters and wake steering performance.

2. Methodology

WindSE is written using FEniCS, an open-source collection of software for the automated, finite-element solution of partial differential equations [4]. Pairing FEniCS with `dolfin-adjoint` allows WindSE to automatically compute derivatives for gradient-based optimization with respect to control parameters [5].

2.1. Unsteady Fluid Solver

The fluid physics are governed by the filtered Navier-Stokes and continuity equations, Equations 1 and 2, respectively.

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

$$\nabla\cdot\mathbf{u} = 0 \tag{2}$$

The unsteady flow solver uses a pressure correction approach to numerically solve the Navier-Stokes equation while enforcing the incompressibility constraint. This widely-used numerical scheme involves first discretizing the viscous, convective, pressure gradient, and forcing terms in the Navier-Stokes governing equation to calculate a predicted velocity, $\mathbf{u}^*$ at the next timestep [6] [7]

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

where $\mathbf{u}$, P, and $\mathbf{F}$ are the filtered velocity, pressure, and forcing terms, respectively, solved from the previous time step k, and ν is the viscosity, which includes contributions from kinematic viscosity and a sub-grid scale eddy viscosity, $\nu = \nu_0 + \nu_T$. The convecting velocity, $\mathbf{A}$, and the convected velocity, $\mathbf{B}$, are expressions of the filtered velocity field chosen for numerical stability. $\mathbf{A}$ is an explicit Adams-Bashforth projection which returns a velocity corresponding to time level $k + 1/2$, $\mathbf{A} = 1.5\mathbf{u}^k - 0.5\mathbf{u}^{k-1}$, and the convected velocity, $\mathbf{B}$, is an implicit average that also exists at time level $k + 1/2$, $\mathbf{B} = 0.5\left(\mathbf{u}^* + \mathbf{u}^k\right)$. The timestep, Δt, used to discretize the problem is adaptive, proportional to the ratio of the minimum cell size to the maximum fluid velocity, $\propto \Delta_{\min}/|\mathbf{u}_{\max}|$. This allows the fluid solver to make stabilizing adjustments in response to highly dynamic fluid features while decreasing the time to solution overall.

The predicted velocity resulting from Equation 3 is, in general, not divergence free; correcting this to enforce the incompressibility constraint requires constructing and solving a Poisson equation for the change in the pressure field, ϕ.

$$\frac{\mathbf{u}^{k+1} - \mathbf{u}^*}{\Delta t} = -\nabla\phi \tag{4}$$

$$\frac{\nabla\cdot\mathbf{u}^{k+1} - \nabla\cdot\mathbf{u}^*}{\Delta t} = -\nabla^2\phi \tag{5}$$

$$\nabla\cdot\mathbf{u}^* = \Delta t\nabla^2\phi \tag{6}$$

After solving Equation 6 for ϕ, the predicted fluid velocity can then be updated with the effects of the pressure change yielding the true filtered velocity at the next timestep, while the pressure field can be updated to reflect the changes required to enforce continuity.

$$\mathbf{u}^{k+1} = \mathbf{u}^* - \Delta t \nabla \phi \tag{7}$$

$$P^{k+1} = P^k + \phi \tag{8}$$

Direct numerical simulation of the turbulent wind turbine regime would require prohibitively expensive refinement of the computational grids. To avoid this pitfall, the fluid solver employs the Smagorinsky eddy viscosity model to dissipate energy at the sub-grid scale.

$$\nu_T = (C\Delta)^2 |\mathbf{S}| \tag{9}$$

$$\mathbf{S} = \frac{1}{2}\left(\nabla \mathbf{u} + (\nabla \mathbf{u})^T\right) \tag{10}$$

where C is the Smagorinsky constant, 0.17 in the current work, Δ is a characteristic length associated with each cell, calculated as (cell volume)$^{1/3}$ in three dimensions, and $\mathbf{S}$ is the resolved strain rate tensor, a function of the filtered velocity gradient. This approximation provides the simulations with both speed and stability while preserving the physics of the simulated, large-scale wake features arising from the turbine actuator forces.

2.2. Turbine Actuator Models

A wind turbine in WindSE can be modeled in two different ways: actuator disks or actuator lines. In the actuator disk model, the force on the fluid from the turbine is given by

$$\mathbf{F}(x, y, z, t) = \frac{1}{2}\rho A C_t' \mathbf{u}\,(t)^2\, T(x) D(y, z) \tag{11}$$

where ρ is the fluid density, A is the circular area defined by the turbine radius, R, and $C_t' = C_t/(1-a)^2$ is a modified thrust coefficient defined using the axial induction factor a [2]. This force is distributed around the continuous disk swept by the rotor using a pair of Gaussian kernels, T and D, for the normal and in-plane directions, respectively.

In the recently-added actuator line model, the force on the fluid from the blade, $\mathbf{F}$, follows a Gaussian distribution defined along a line with

$$\mathbf{F}(x, y, z, t) = -\sum_{j=1}^{N} \mathbf{f}_j(x_j, y_j, z_j, t) \frac{1}{\varepsilon^3 \pi^{3/2}} \exp\left(-\frac{\mathbf{r}_j \cdot \mathbf{r}_j}{\varepsilon^2}\right) \tag{12}$$

where N is the number of blade discretization segments, defined by $\lceil R/\Delta_{\min} \rceil = 10$ in the current work, ε controls the width of the Gaussian, currently $2\Delta_{\min}$, $\mathbf{r}_j$ is the distance vector between any fluid point (x, y, z) and actuator point (x_j, y_j, z_j), and the force at each actuator point, $\mathbf{f}_j$, is the result of projecting the lift and drag forces, $\mathbf{f}_l$ and $\mathbf{f}_d$, onto directions perpendicular and parallel, respectively, to the relative fluid velocity. The lift and drag forces are given by

$$\mathbf{f}_l(x, y, z, t) = \frac{1}{2} C_l \rho A u_{\text{mag}}(x, y, z, t)^2 \tag{13}$$

$$\mathbf{f}_d(x, y, z, t) = \frac{1}{2} C_d \rho A u_{\text{mag}}(x, y, z, t)^2 \tag{14}$$

where C_l and C_d are the lift and drag coefficients defined at each actuator point, A is the product of chord length at the actuator point, c, and actuator segment width, $R/(N-1)$, and u_{mag} is the magnitude of the relative velocity between the local flow and the rotating blade segment [8] [9]. For the purpose of simplifying the adjoint formulation, the local flow is currently approximated using the inflow conditions. In the current work, the coefficients of lift and drag are not dependent on angle of attack; instead, they are assigned a value which varies only with spanwise position.

2.3. Adjoint Framework

The derivation of the adjoint system is handled by `dolfin-adjoint`, which is a python package built on top of FEniCS. While performing the initial forward solve, `dolfin-adjoint` records the control-dependent steps to a "tape". This tape contains all of the steps required to recompute the forward problem as well as the dependencies of each step on the control variables. Then an objective function is defined that is dependent on the results of the forward solve. The control(s) is the parameter(s) that can be adjusted to optimize the objective functional. With all of this information, `dolfin-adjoint` can "rewind" the tape to compute gradients of the objective function with respect to the controls. The rewinding process uses adjoint methods and the chain rule to propagate the derivatives backwards through the forward solve [10, 11]. These gradients are then provided to `scipy`, which is the optimization package responsible for updating the control accordingly.

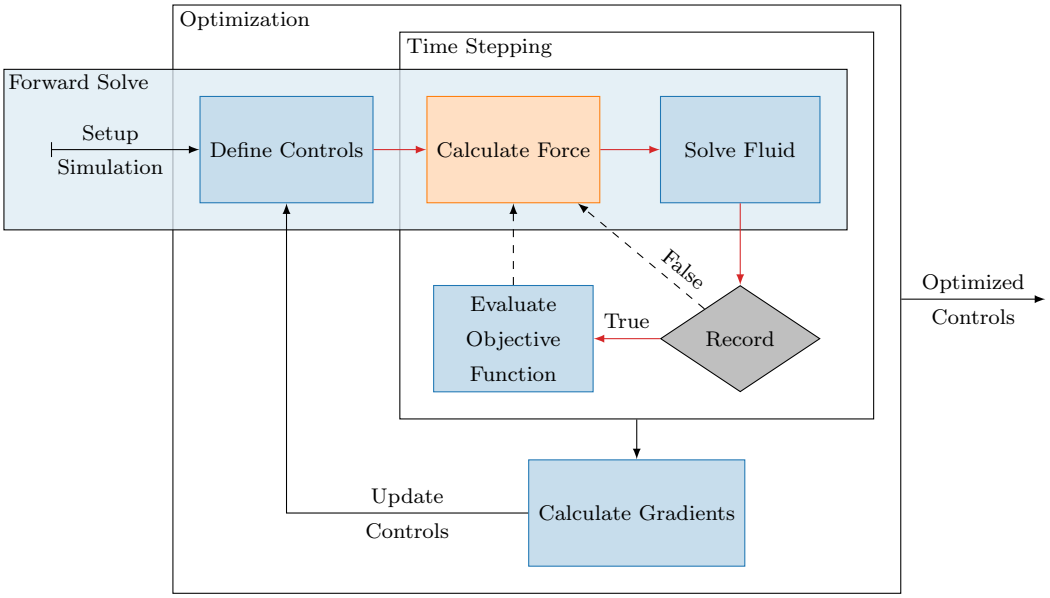

Figure 1. Block diagram depicting the sequence of operations that define the time stepping and optimization loops. Solid arrows indicate information flowing between steps. Dashed arrows simply show which step is next. Red arrows indicate a flow of information that needs to be tracked by `dolfin-adjoint`. Orange blocks indicate custom derivative calculation methods.

Figure 1 shows a simplified flow chart of the algorithm WindSE uses to optimize with respect to a time varying simulation. The optimization starts by performing an initial forward solve, which consists of defining the controls, calculating the body force imparted by the actuator lines, and then solving for the fluid fields with respect to time. During each time step, it is possible to record a contribution to the objective function. If this is the case, a separate routine is called to evaluate the objective function.

The red arrows in Figure 1 indicate steps where information is transferred that depends on the controls, and therefore need to have associated methods that can compute the local gradients. For the most part, `dolfin-adjoint` handles the construction of these gradients using automatic differentiation. However, due to the complex nature of the actuator line model, it became useful to construct an analytic method for computing the derivatives with respect to the chord length control.

3. Results

3.1. Case Description

Testing the controllability of the turbine wake proceeds in two parts. First, the effect on wake deflection of yawed operation combined with changing chord profiles is investigated and this data is used to define a suitable objective function. Second, the sensitivity of the chord control is quantified using local adjoint gradient information. All the studies presented in the current work use the International Energy Agency (IEA) 3.4 MW land-based reference wind turbine (RWT) properties and rated operational conditions shown in Table 1 [12].

Table 1. Simulation Setup for the IEA 3.4 MW RWT.

Property	Value
Hub Height	110 m
Rotor Diameter, D	130 m
Rotor Rated Speed	10.6 rpm
Hub Height Wind Speed	$9.0\,\mathrm{m\,s^{-1}}$
Tip Speed Ratio	8.01

The lift and drag coefficients vary along the length of the blade as shown in the left portion of Figure 3. These values were obtained using `CCBlade`, a blade element momentum method for analyzing wind turbine aerodynamic performance [13]. In the current study, the coefficient of lift and drag are held constant with respect to relative fluid velocity and vary only with blade span. This study adopts a wind speed of $9.0\,\mathrm{m\,s^{-1}}$ at hub height. The turbine is operated at its rated tip speed ratio of 8.01, which corresponds to a rotor speed of 10.6 rpm and a tip speed of $72.1\,\mathrm{m\,s^{-1}}$. The computational domain is chosen to be sufficiently large to model actuator forces and the deflected wake without spurious edge effects, $8D \times 4D \times 4D$ in the x, y, and z-directions, as shown in Figure 2.

Using Figure 2 as a reference, a log inflow profile $u_x(z) = \frac{u^\star}{\kappa} \ln\left(\frac{z}{z_0}\right)$ is specified for velocity along the left-hand wall, where $z_0 = 0.0001$ and $\kappa = 0.4$. Simulations are carried out on a regular Cartesian mesh with global mesh refinement in the near-ground region and local mesh refinement in the cylindrical region aligned with the rotor axis and extending to the outflow plane. No freestream turbulence is added to the inflow, but turbulence is created by the turbine and propagated in the wake. A no-slip condition is specified at the bottom wall, and a free slip condition is specified at the top wall and both spanwise walls, prohibiting velocity in the direction normal to the boundary. At the outlet on the right-hand side of the domain, a Dirichlet pressure condition is set, $P = 0$, which removes any additional computational complexity involving solving a singular system of equations for ϕ.

3.2. Effect of Chord Length on Turbine Wake

Using the problem specified above, the controllability study begins by first confirming the possibility of changing wake position by varying the chord profile enforced at the actuator nodes

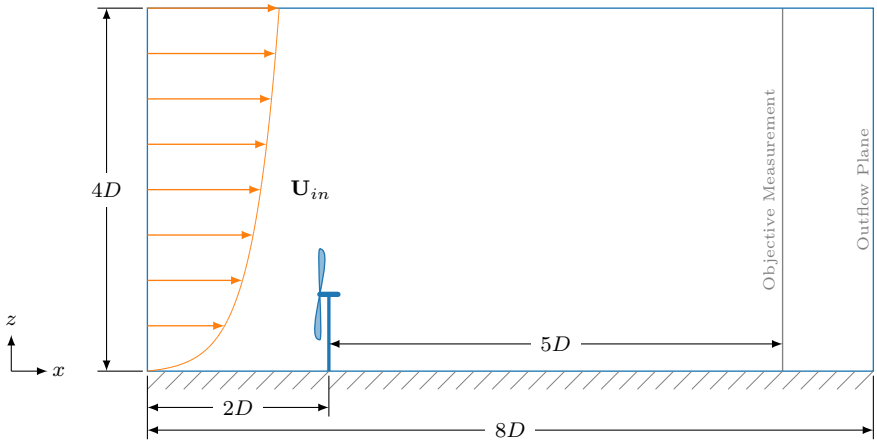

Figure 2. The computational domain, simulation definitions, and a representative objective measurement plane used in the current work.

defining the blade span. Changes in chord length affect the lift and drag forces, and in turn the wake deflection, through the quantity A in Equations 13 and 14. Wake steering was chosen as a target due to its natural observability in the velocity field as well as an increased focus on wake steering in wind farm optimization studies [14]. At this stage, no gradient calculation or sensitivity analysis takes place; rather, the chord profile along the blade span is manually changed to produce multiple, separate simulations. Four tests were chosen to demonstrate the control authority over wake behavior with the distinguishing design parameters listed in Table 2.

Table 2. Wake Steering Parameter Study.

Case	Yaw	Chord Modification
A	0°	None
B	20°	−40% Chord at tip
C	20°	None
D	20°	+40% Chord at tip

In Cases A and C, the chord length is set to the baseline profile specified for the IEA 3.4 MW RWT. Comparing simulations A and C shows the effect of yawing a turbine on the wake, while simulations B and D are designed to illustrate the effect of modifying chord length on the wake produced during yawed operation. The changes to the chord profile are illustrated in the right portion of Figure 3.

Each simulation was run until the wake reached the outflow wall of the computational domain, at which point recording of the time averaged velocities began. These time averaged velocities were recorded for 35 complete blade revolutions, 200 s, to ensure the sample window was large enough for the turbulent flow statistics to begin to stabilize. Full stabilization of the flow statistics requires a significantly larger measurement window [15], but this relatively coarse truncation yields smoothly differentiable fields for the present study. The averaged velocities along a hub-height xy-plane are shown in Figure 4.

Studying the time averaged flows in Figure 4, it is evident that the chord profile defined on the actuator lines does have an impact on the turbine wake. The most obvious effect is a larger

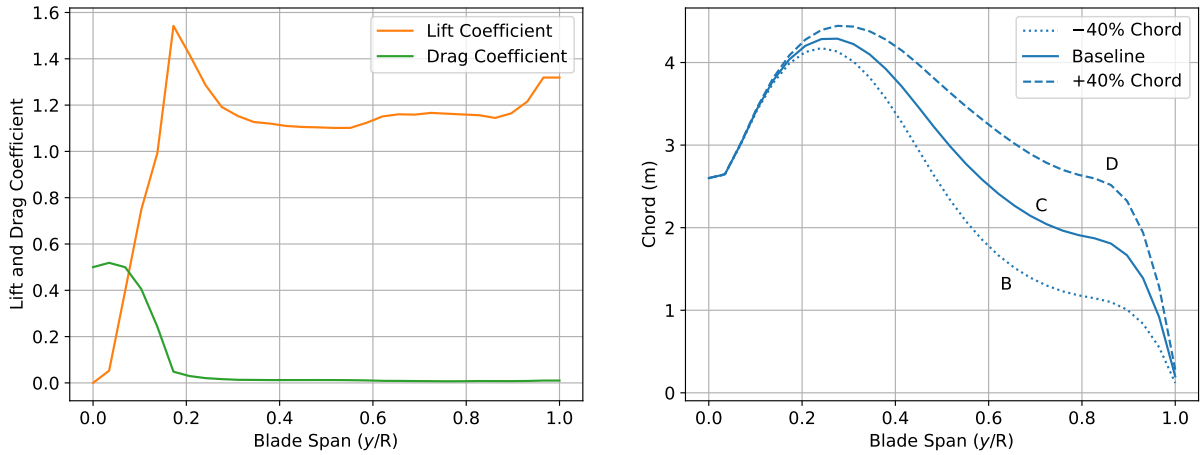

Figure 3. Left: The lift and drag coefficients as a function of blade span; Right: The modified chord profiles used in this pilot study, where the +40% and −40% modifications were generated through simple scaling of the baseline chord values.

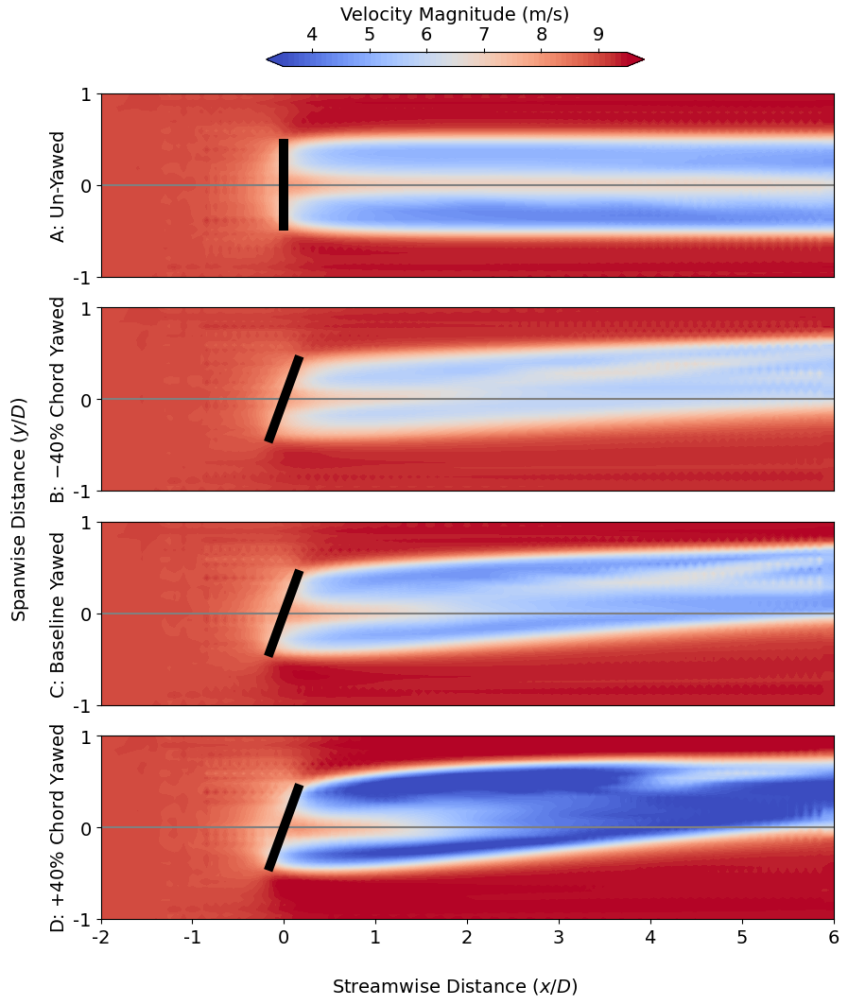

Figure 4. The time-averaged wakes produced behind the four turbine cases; all axis tick marks located at integer multiples of D.

wake deficit associated with increasing chord, but there also appear to be more subtle changes in both the shape and path of the wake. To quantify this effect and enable a more thorough sensitivity study, it is necessary to define a metric for the changing wake position. In the current problem definition, it is the displacement in the y-direction that is of principal interest in the larger context of wind farm optimization; for example, shifting a wake such that it interacts favorably with downwind turbines. Enhancing the wake steering of an upstream turbine in this way may increase the net power output for the farm, even if the upstream turbine's output is somewhat diminished. A weighted centroid calculation is used to characterize the wake position, calculated as the average of all y-positions included on the yz-plane of interest weighted by the magnitude of the wake deficit

$$y_c = \frac{\int_{x=n\mathrm{D}} y_i u_{\mathrm{wd}}\, dA}{\int_{x=n\mathrm{D}} u_{\mathrm{wd}}\, dA}, \tag{15}$$

where nD is the the distance that is n rotor diameters downstream and the wake deficit magnitude, u_{wd}, is given by, $u_{\mathrm{wd}} = |\mathbf{u}_{\mathrm{inflow}} - \mathbf{u}|$. In this test case, the centroid calculation is performed only once as a post-processing step after the flow field has already been averaged in time. Performing this calculation on measurement planes located $1D$, $2D$, ..., $6D$ downstream from the turbine generates the centroid paths shown in the bottom portion of Figure 5.

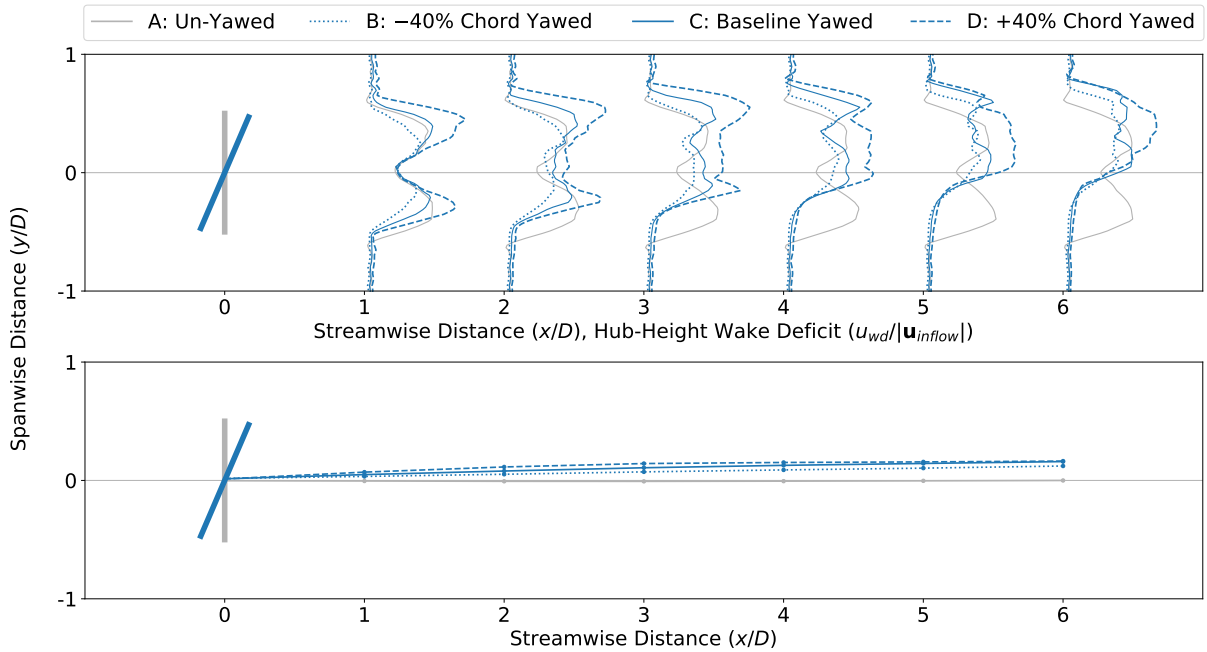

Figure 5. Top: Hub-height slices of the wake deficits produced by the four turbine cases; Bottom: The movement of the wake centroid, y_c, in the spanwise direction.

Although the centroid calculations in the bottom portion of Figure 5 show relatively small changes, there is a distinct shift in the $+y$-direction associated with larger chord lengths and a shift in the $-y$-direction associated with smaller chord lengths. The upper portion of Figure 5 confirms the differences in wake shape suggested by Figure 4. Based on these test cases, it appears the calculation of the wake centroid, y_c, will provide a suitable objective function for controlling wake steering through modification of the chord profile. For large wind plants, the benefits of wake steering are in large part due to the enhanced vertical entrainment caused by counter-rotating vortices created by the yawed rotor. Figure 6 shows that the characteristic

kidney bean shape of a distorted wake is accentuated in the +40% chord case, reflecting additional plant-level benefits due to vortex-induced entrainment not directly captured by the wake centroid metric.

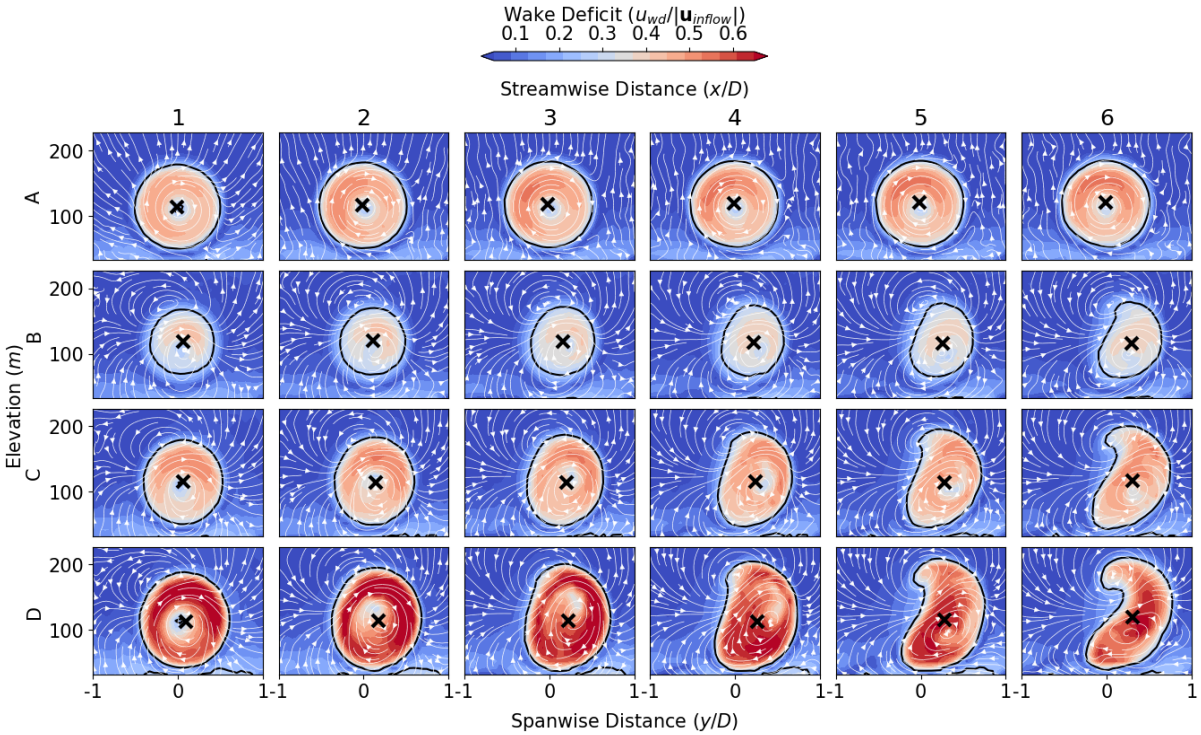

Figure 6. The normalized wake deficit and streamlines in yz-planes in the turbine wake where the 2D centroid, (y_c, z_c), is marked with an ✖; z_c is calculated by replacing y_i with z_i in Equation 15.

3.3. Verification of Adjoint Gradients

Before proceeding to the results of the sensitivity study, the expected behavior of the objective function with respect to control variable changes is verified. Because of the inherently unsteady nature of the problem, using an instantaneous centroid calculation as the objective function leads to problematic oscillations. To ensure that the objective function, J, for the sensitivity analysis is well posed, the centroid position returned by Equation 15 is added to a cumulative temporal average at each timestep during the simulation to form an objective function more indicative of the overall state of the system.

In the current work, this objective function depends on the value of the chord, $J(c)$, with other variables being held constant. Changes to this control must produce a measurable change in the objective function. With this in mind, verification of the objective function and its analytically-derived gradient, ∇J, is carried out by comparing the value of the analytically-derived gradient with its finite-difference equivalent as shown in Equation 16.

$$|J(c + h\ \delta c) - J(c) - h\ \delta c \nabla J| \to 0 \tag{16}$$

This difference, the second-order Taylor remainder, is expected to converge at $\mathcal{O}(h^2)$; that is, halving the perturbation amount, h, for a fixed change in the control variables, δc, should yield a factor of four decrease in the difference expressed by Equation 16. In the current study,

repeatedly halving h recovers the expected second-order convergence, verifying that both the objective function and its analytically-derived gradients are behaving as expected.

3.4. Wake Sensitivity to Chord Profile

After verifying the convergence of the Taylor remainder, attention turns to measuring the sensitivity of the wake with respect to the chord profile. As described in Section 3.3, the objective function for this study is the time-averaged wake center, y_c, computed using Equation 15. To explore how this wake center depends on the chord profile, several tests were performed the results of which are shown in the left portion of Figure 7. Additionally, the recorded wake centroid values underlying these sensitivities can be found in the right portion of Figure 7. In both cases, the results shown correspond to an objective measurement plane located $5D$ downstream from the turbine. All of the following tests use the unmodified IEA 3.4 MW RWT turbine properties and model the rotor yawed at 20°, identical to Case C from Table 2. The summary for each sensitivity test is found in Table 3

Table 3. Sensitivity Test parameters. The objective function is first recorded at the "Start Record Time" and is then time averaged until "Final Record Time" resulting in the final averaged wake center in the "Objective Value" column.

Trial	Start Record Time	Final Record Time	Chord Modification	Objective Value
I	0 s	300 s	None	13.00 m
II	96.3 s	150 s	None	19.88 m
III	96.3 s	150 s	+1 m uniformly	26.46 m

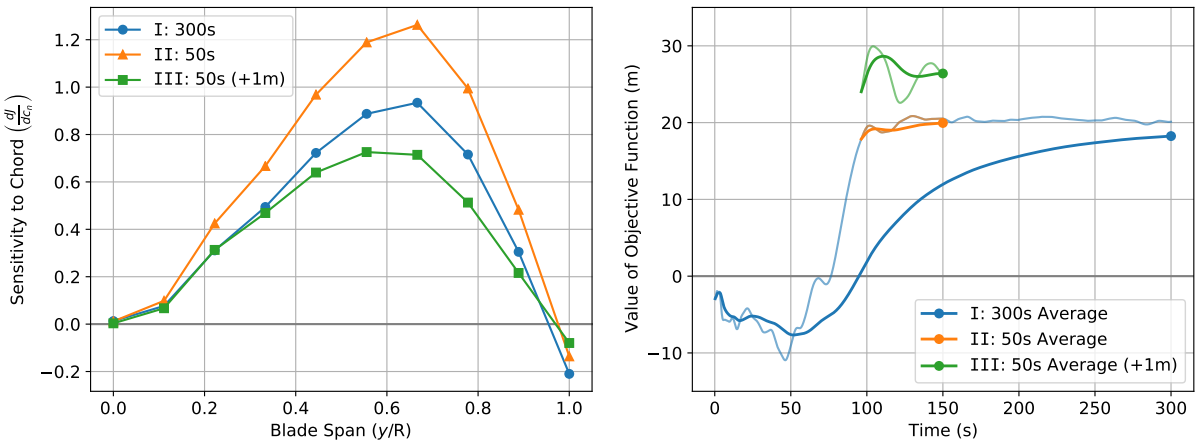

Figure 7. Left: Adjoint-based local sensitivities for the chord length corresponding to the cases listed in Table 3. The sensitivities show how much the wake centroid is expected to move, in meters, when the chord at a node is increased by one meter. Right: The value of the objective function over time; lighter lines correspond to the raw objective value, darker lines are time-averaged values.

The sensitivities shown in Figure 7 illustrate how modifying the chord can affect the wake location. Our sensitivity results show the greatest impact on wake deflection comes from increasing chord at roughly 2/3 span, and that increases in chord increase wake deflection

everywhere except for the blade tip where decreasing chord is beneficial. While the magnitude of the sensitivities corresponding to an individual node are relatively small ($\mathcal{O}(1\,\mathrm{m})$), the cumulative effect of modifying chord at multiple spanwise positions can achieve relatively large enhancements to wake deflection.

Additionally, we found that the sensitivity results were affected by the objective recording time window. In Trial I, the recording window was the entire 300 s simulation, including the approximately 100 s at the start of the simulation before the wake reached the objective measurement plane, resulting in two problems. First, as shown in the left portion of Figure 7 comparing Trial I to Trial II where the objective was recorded only after this 100 s period, the average wake location becomes less sensitive. This reduced sensitivity reflects the fact that for the first 100 s, the chord has no effect on the objective function. The second issue with Trial I is that the right portion of Figure 7 shows how the average wake center is slow to recover and thus requires a longer simulation time for stabilization of the objective function. Trial II, with a recording window of just over 50 s and a total simulation time of 150 s, recovers the general behavior and wake center seen in Trial I. Both of these problems show why care must be taken when choosing a recording window.

In Trial III we also consider a case where the chord values are uniformly increased by 1 m. Trial III serves as an illustration of how the chord sensitivities at each node along the rotor have a cumulative effect. Even though the maximum sensitivity for Trial II is slightly more than 1.2 m of wake deflection per 1 m of increased chord, increasing the chord value of all nodes along the blade span results in a total increase of 6.58 m of wake deflection. This is approximately a 33% increase over Trial II. This additional deflection might allow the upstream turbine to be yawed less aggressively resulting in less-diminished power output while maintaining the same degree of wake steering.

The other purpose of Trial III is to show how the nominal condition affects the sensitivities. Across the entire blade span, the sensitivities for Trial III were reduced, implying that there is an upper limit to how effective increasing the chord is at controlling the wake, which corresponds to reality. Additionally, Trial III suggests that the sensitivities vary smoothly with respect to input conditions, a prerequisite for performing an optimization.

4. Concluding Remarks

Recent developments in WindSE include the ability to simulate unsteady flows with high numerical stability and an adjoint-preserving implementation of actuator line turbine models to enable gradient-based optimization. This capability is used to investigate gradient-based local sensitivities of wake deflection with respect to blade chord at 9 spanwise locations. The wake position seems to be most sensitive to changes in the chord length at $\sim 2/3$ of total blade span. At its most sensitive, $\frac{dJ}{dc_n} \approx 1.2$, indicating that a change in chord of 1 m at a single actuator node results in just over 1 m of wake shift in the $+y$-direction. It is worth noting that in general, changing the chord profile of a blade is likely to involve adjustments to multiple nodes rather than a single point which means these individual sensitives contribute to a cumulative deflection effect. For almost all locations along the blade, the addition of chord length enhances wake steering, with the exception of the blade tip, where $\frac{dJ}{dc_n} < 0$. A more highly refined blade or selective refinement of blade tips with variably sized actuator segments may help pinpoint the wake behavior in this transition region.

The current WindSE actuator line framework predicts a drop in power output of 8% when moving from 0° to 20° yaw. When considering the farm-scale system, the design goal is that this singular drop in power will be more than offset by downstream turbines due to the creation of more favorable wake interactions through steering. A modified blade shape may be inherently more effective at this sort of steering, suggesting that the upstream turbine could be yawed to a lesser degree and achieve the same steering effect while reducing its own drop in power.

A carefully constructed and robust wake-steering objective function that accounts for power output across multiple turbines and the effects of wake steering—in terms of both deflection magnitude and the effects of enhanced vertical entrainment—is critical for the success of future optimization studies.

One of the difficulties associated with this type of high Reynolds number, time series, sensitivity study is the inherently unsteady nature of the flow physics involved [16]. In order to assess the sensitivity of an individual actuator segment with any accuracy, the objective function must be measured with a large enough temporal window to allow the averaged turbulent flow variables to approach a steady state. This requires a large number of timesteps to be simulated and a potentially prohibitive amount of memory to store the complete sequence of calculations necessary for the optimization framework. Additionally, the effect of changes to the control variable must be large enough to avoid being washed out in the subsequent averaging of the highly variable turbulent flow. A full optimization study further complicates this issue by introducing new temporal elements in the form of the control windows or time horizons used successfully in other unsteady optimizations [15].

Choosing to manipulate the chord length as the sensitivity controls is convenient from the perspective of the current actuator line implementation but may not be the most physically motivated choice if the goal is to enable the design and development of rotor blades based on physical airfoil sections. Linking the effects of chord, lift, and drag through a calculation of angle of attack may allow a single unified control for these values that not only integrates with the actuator line model but also has a direct physical interpretation. For turbine blades with a prescribed progression of airfoil shapes, it may be possible to study the sensitivity to angle of twist, for example, allowing an optimization to be performed to find the ideal characteristics of a rotor for a given objective function.

References

[1] Allen J, Young E and King R 2020 *https://github.com/NREL/WindSE*
[2] King R N, Dykes K, Graf P and Hamlington P E 2017 *Wind Energy Science* **2** 115–131
[3] Allen J, King R and Barter G 2020 *Journal of Physics: Conference Series* **1452** 012066
[4] Logg A, Mardal K A, Wells G N *et al.* 2012 *Automated Solution of Differential Equations by the Finite Element Method* (Springer)
[5] Logg A, Wells G N and Hake J 2012 *DOLFIN: a C++/Python Finite Element Library* (Springer) chap 10
[6] Chorin A 1967 *Journal of Computational Physics* **2** 12–26
[7] Kim J and Moin P 1985 *Journal of Computational Physics* **59** 308–323
[8] Sorensen J N and Shen W Z 2002 *Journal of Fluids Engineering* **124** 393–399
[9] Churchfield M, Lee S, Moriarty P, Martinez L, Leonardi S, Vijayakumar G and Brasseur J Nashville, Tennessee, 2012 *Proceedings of the 50th AIAA Aerospace Sciences Meeting*
[10] Funke S W and Farrell P E 2013 *arXiv preprint arXiv:1302.3894*
[11] Farrell P E, Ham D A, Funke S W and Rognes M E 2013 *SIAM Journal on Scientific Computing* **35** C369–C393
[12] Bortolotti P, Tarres H C, Dykes K, Merz K, Sethuraman L, Verelst D and Zahle F 2019 *National Renewable Energy Lab Report NREL/TP-73492*
[13] Ning S A 2014 *Wind Energy* **17** 1327–1345
[14] Dykes K, Hand M, Stehly T, Veers P, Robinson M and Lantz E 2017 *National Renewable Energy Lab Report NREL/TP-5000-68123*
[15] Goit J P, Munters W and Meyers J 2016 *Energies* **9** 29
[16] Wang Q 2013 *Journal of Computational Physics* **235** 1–13