


RESEARCH ARTICLE

WILEY

Aero-servo-elastic co-optimization of large wind turbine blades with distributed aerodynamic control devices

Nikhar J. Abbas¹  | Pietro Bortolotti²  | Christopher Kelley³ |
Joshua Paquette³ | Lucy Pao⁴ | Nick Johnson³

¹Paul M. Rady Department of Mechanical Engineering, University of Colorado Boulder, Boulder, Colorado, USA

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

³Wind Energy Technologies Department, Sandia National Laboratories, Albuquerque, New Mexico, USA

⁴Department of Electrical, Computer, and Energy Engineering, University of Colorado Boulder, Boulder, Colorado, USA

Correspondence

Nikhar J. Abbas, Paul M. Rady Department of Mechanical Engineering, University of Colorado Boulder, Boulder, CO, USA.
Email: nikhar.abbas@colorado.edu

Funding information

U.S. Department of Energy

Abstract

This work introduces automated wind turbine optimization techniques based on full aero-servo-elastic models and investigates the potential of trailing edge flaps to reduce the levelized cost of energy (LCOE) of wind turbines. The Wind Energy with Integrated Servo-control (WEIS) framework is improved to conduct the presented research. Novel methods for the generic implementation and tuning of trailing edge flap devices and their controller are also introduced. Primary flap and controller parameters are optimized to demonstrate potential maximum blade tip deflection reductions of 21%. Concurrent design optimization (i.e., co-design) of a novel segmented wind turbine blade with trailing edge flaps and its controller is then conducted to demonstrate blade cost savings of 5%. Additionally, rotor diameter co-design optimization is demonstrated to reduce the LCOE by 1.3% without significant load increases to the tower. These results demonstrate the efficacy of control co-design optimization using trailing edge flaps, and the entirety of this work provides a foundation for numerous control co-design-oriented studies for distributed aerodynamic control devices.

KEYWORDS

control co-design, control systems, distributed aerodynamic control, multidisciplinary design, optimization

1 | INTRODUCTION

The continuous growth in the size of wind turbine rotors calls for innovative solutions to limit loads, masses, and costs. Long blades encounter very different inflow conditions in sheared and turbulent wind during each revolution, and turbines might also be operated in yawed conditions for wind plant power maximization.¹ Additionally, slender and flexible blades undergo large out of plane deflections that generate high torsional moments at the blade root, challenging the limits of pitch actuators. Because of this pitch actuator limitation and additional pitch bearing damage caused by increasing blade mass, the dynamics of modern wind turbine blades are constrained in how well they can be controlled. One solution to address these growing challenges consists of moving into “smart blade design” or “smart rotors,” which include active flow control devices along the span of each blade.^{2,3} Distributed aerodynamic control (DAC) devices dynamically change the aerodynamic properties of an airfoil. Multiple different types of DAC have been proposed and investigated, such as microtabs,⁴ active twist control,⁵ boundary layer control,⁶ and trailing edge flaps.^{7,8} The presented work focuses on trailing edge flaps, which are among the most popular technologies thanks to their common use in aeronautical and aerospace

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Authors. *Wind Energy* published by John Wiley & Sons Ltd.

applications. We aim to quantify the potential influence of flaps on wind turbine blades in terms of levelized cost of energy (LCOE) by formulating novel control co-design techniques and implementing and exercising them on a modern land-based wind turbine design.

Broadly, control co-design is a fundamental integration of the controller into the design process of a wind turbine. A key aspect of a co-design study is, naturally, the wind turbine control logic, which presents additional challenges when DAC devices are present. A number of researchers have investigated different methods of controlling DAC devices. An intuitive approach consists of changing the aerodynamics of the blade as a function of the blade's azimuthal position.^{8–10} This is, perhaps, the most straightforward method of controlling DAC devices, and it tends to target fatigue loads. It also lends itself to a direct comparison to individual pitch control (IPC) because both methods of load alleviation directly target the once-per-revolution (1P) cyclic loading of a wind turbine blade. Some researchers have investigated the use of more modern, state-space-based control methods. A linear quadratic regulator is used by Bergami et al¹¹ in an attempt to reduce maximum loads on the blade root, and a model predictive control technique is used by Bartholomay et al¹² for fatigue load reductions. Similarly, Ng et al¹³ uses a novel aeroelastic formulation to employ a linear-quadratic-Gaussian controller to reduce the root-mean-square values of the blade root bending moment and maximum blade tip deflections. A number of methods that specifically account for limitations on trailing edge flap actuators are also explored by Unguran et al,^{14,15} where robust control methods, such as H_∞ control, are used to avoid excessive loads due to actuator saturation. More “classical” control methods are employed by Wilson et al¹⁶ and Zhang et al,¹⁷ where proportional-integral (PI) and proportional-integral-derivative (PID) controllers are used, respectively. In both studies, reductions of blade root bending moment and maximum tip deflections are achieved, but neither publication gives a substantially thorough analytical method for tuning the controller gains.

The availability of readily implementable open-source controllers for trailing edge flaps is quite limited in the literature.¹⁸ This work expands the Reference OpenSource Controller (ROSCO)¹⁹ to include generic controller tuning algorithms for active flow control devices. In a similar methodology to the controller tuning and implementations presented in Abbas et al,²⁰ generic algorithms have been developed to tune a (PI) controller to minimize the blade root bending moments. This controller aims to achieve two goals:

1. Reduce fatigue damage to the blade by reducing the cyclic loads on the blade due to the rotor rotation and inherent periodic sampling of the inflow wind conditions (with possible turbulence, shear, veer, etc.).
2. Reduce maximum blade tip deflections that are typically design-driving and enable rotor designs that can decrease the cost of wind energy.

Because the developed controller tuning and implementation methods are purely analytical and can be automatically implemented given a blade's airfoil characteristics, these updates to ROSCO can be included in system design optimization loops. Though optimization studies for active flow control have been done in the past,^{21–23} none of the published studies have considered combined co-optimization of active flow control devices, their controller, and rotor parameters in large-scale design studies. The work by Chen et al²¹ presents the most comprehensive design studies; however, the controller gains were periodically updated during the optimization process due to limitations imposed by computational expense. In an attempt to better understand the efficacy of trailing edge flaps, this work leverages high-performance computing (HPC) resources to conduct comprehensive design studies to specifically target blade mass and rotor diameter as parameters to decrease the cost of wind energy while using trailing edge flaps.

The paper is structured as follows. Section 2 introduces the co-design environment developed for this study. The wind turbine used in this work is presented in Section 3, and the controller logic used on the turbine is presented in Section 4, focusing on the control of the trailing edge flaps. Parametric studies used to characterize the solution space are presented in Section 5, then results from the automated co-design framework are presented in Section 6, including a comparison of loads from standard design load cases (DLCs) in Section 6.6. The paper is closed in Section 7, which discusses the challenges faced during the study, summarizes key takeaways of this work, and recommends several future investigations.

2 | CO-DESIGN ENVIRONMENT

2.1 | Overview of control co-design

The term “control co-design” has become a popular topic in recent years.²⁴ The combined development of wind turbine simulation optimization tools has brought design capabilities to a point in which automated co-design is possible.^{25,26} Historically, an iterative sequential wind turbine design process, where the controller is considered in a separate step than the aeroelastic design, has been the primary practice.^{27–29} Figure 1 shows a flowchart of a simple wind turbine multidisciplinary analysis and optimization (MDAO) design process without a controller optimization. Either a gradient- or nongradient-based optimization solver is employed to iterate and optimize a number of wind turbine design variables to minimize a prescribed objective function. For each iteration, if an aeroelastic tool is implemented in the workflow, a controller is implemented, time-domain simulations are run, and a cost (or similar objective function) analysis is conducted. If the global objective function only changes within a prescribed tolerance, the design is considered converged. Otherwise, the optimization solver iterates the process. In this case, the controller design parameters are frozen during the optimization of the aeroelastic design, and then are updated in the subsequent step before the objective function is analyzed.

In a control co-design process, controller optimization happens concurrently with the wind turbine design optimization, such that the controller parameters (e.g., bandwidth) are included as additional optimization variables. This extends the optimization process to the workflow shown in Figure 2, where the controller and wind turbine parameter optimizations are coupled. Here, the optimization solver will simultaneously iterate and modify both the wind turbine and controller parameters in an attempt to minimize the global objective function.

This, of course, presents an added layer of difficulty in the development of the system design optimization procedures. The Wind Energy with Integrated Servo-control (WEIS)³⁰ framework has been developed at the National Renewable Energy Laboratory (NREL) to specifically address these challenges. There are a number of wind turbine optimization tools that employ fully nonlinear simulations in-the-loop.^{31,32} To the authors' knowledge, however, none of these existing open-source and openly published optimization tools have included the ability to optimize the controller parameters in a completely automated way, let alone run fully coupled control co-design optimization routines with distributed aerodynamic control capabilities.

2.2 | Trailing edge flaps

In this study, we adopt co-design methods to focus on the optimal design of a trailing edge flap device. A conceptual rendering of the type of device used in this study is shown in Figure 3. When deflected, the trailing edge flap modifies the aerodynamic performance along a blade section. We use the flap to achieve two primary goals: reduction of blade root bending moments and reduction of blade tip deflections.

By reducing the blade root bending moments and blade tip deflections, the cost of energy can be reduced in a number of ways. Perhaps most obviously, by reducing blade root bending moments and blade tip deflections, mass can be removed from the blade to reduce blade costs without violating operational design constraints. There are other ways the turbine can be modified as well—for example, the rotor size can be increased for a given load envelope, or the rotor coning angle can be decreased without violating the blade tower clearance constraints. Both solutions offer a pathway to increase the rotor effective area, and subsequently the annual energy capture.

In this work, the popular airfoil solver XFOIL³³ is used to calculate the aerodynamic polar tables for the blade with and without trailing edge flaps. XFOIL is run twice, first to compute the airfoil shape given a flap extension and deflection, and a second time to obtain the polars at multiple Reynolds numbers given the perturbed airfoil shape. XFOIL settings are kept at the default values.

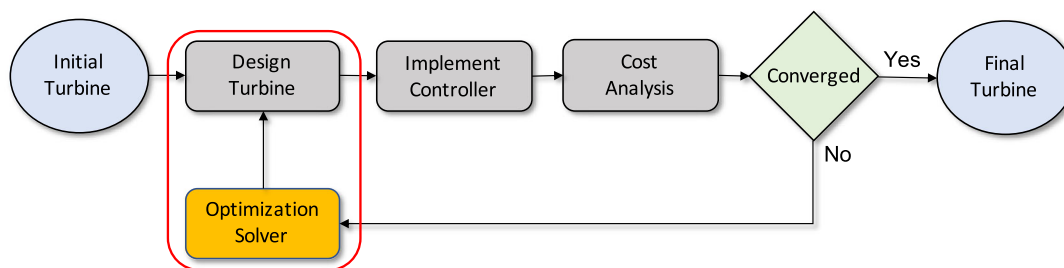


FIGURE 1 Block diagram showing a simple, standard wind turbine design optimization without a controller optimization. Specific turbine parameters (e.g., blade twist and spar cap thickness) might be optimized with respect to some global objective function. Only the wind turbine design parameters are modified to minimize the objective function, as denoted by the red outline.

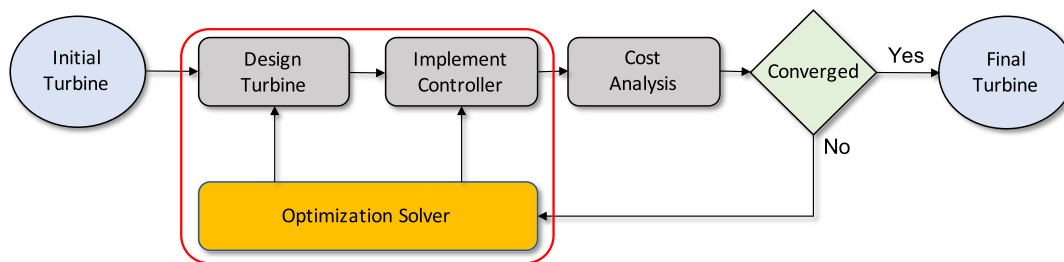


FIGURE 2 Block diagram showing a simple control co-design process workflow. In addition to wind turbine aero-structural design optimization, the controller design is co-optimized as well. The modification of the aero-structural and controller design parameters is coupled, as denoted by the red outline. The optimization solver modifies the design parameters to minimize a global objective function, until convergence.

The aerodynamic polar tables are fed to the aero-servo-elastic solver OpenFAST.³⁴ For the blade sections that contain a trailing edge flap, OpenFAST linearly interpolates between aerodynamic polar tables based on the flap angle to determine the blade's aerodynamic characteristic at any moment in time. OpenFAST includes unsteady airfoil aerodynamics effects adopting the Minnema Pierce variant of the Beddoes Leishman model.³⁵ Section 5 discusses the flap geometries adopted in this study.

2.3 | Optimization framework

The WEIS framework has been developed to perform multifidelity co-design optimization of floating wind turbines. This framework combines many other NREL-developed tools, such as OpenFAST, Wind-Plant Integrated System Design & Engineering Model (WISDEM[®]),³⁶ and ROSCO, to facilitate wind turbine optimization. A very high-level overview of the WEIS workflow is shown in Figure 4. The workflow can be separated into three primary steps, starting and finishing with a standardized wind turbine definition. During the execution of these three steps, numerous other tools are used to comprehensively define, simulate, and analyze the wind turbine. We note that the modules shown in Figure 4 are the primary modules used in this work, not a comprehensive list of the modules used in the WEIS tool set.

In the first step of the WEIS workflow, a preprocessing step initializes the wind turbine design. The user-defined initial parameters of the turbine are translated to the quantities used by the wind turbine simulation and analysis tools. In this work, the blade airfoil types and trailing edge flap characteristics are translated to the airfoil polar tables that are used by OpenFAST during simulation. Additionally, the controller tuning parameters are translated to the PI gains that are used by the ROSCO controller. There are numerous other design parameters that are translated by the individual modules within WEIS as well.




FIGURE 3 A conceptual render of an active trailing edge flap to reduce loads on a wind turbine blade.

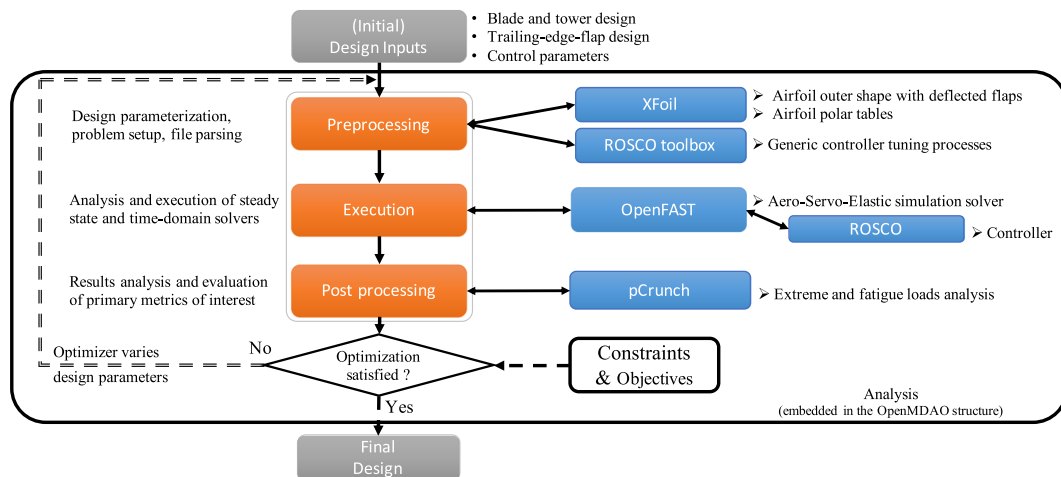


FIGURE 4 A diagram describing the WEIS workflow as relevant to the work presented in this paper.

In the second step of the WEIS workflow, simulation and analysis tools are executed. This includes a number of loads and aerodynamic analysis tools. In particular, OpenFAST is a complete aero-servo-elastic wind turbine simulation tool that is used to simulate the wind turbine. During simulation, the ROSCO controller is called to control not only the standard blade pitch and generator torque actuators but also the trailing edge flaps. Tools such as CCBlade³⁷ and Frame3DD³⁸ are also used for aerodynamic and structural analysis, respectively.

In a final step of the WEIS workflow, post-processing and analysis is conducted. Here, OpenFAST simulation results are parsed, and quantities such as annual energy production (AEP), damage equivalent loads (DELs), and the LCOE are calculated. Another tool dubbed “pCrunch,”³⁹ a Python-based evolution of NREL’s mCrunch and mExtremes⁴⁰ tool sets, is used to calculate some of these values. Additionally, the cost models in WISDEM are used for the financial analysis leading to the levelized cost of wind energy.

These three primary steps are collected and wrapped within OpenMDAO,⁴¹ a Python-based optimization framework. This enables the optimization of many wind turbine input parameters to satisfy numerous objective functions and constraints. A number of gradient and nongradient-based optimization methods can be leveraged too. Finally, parameter studies or “design of experiments” can be conducted to better understand the optimization space or generate surrogate models.

3 | THE BIG ADAPTIVE ROTOR (BAR)

This research has been conducted under the BAR project, a U.S. Department of Energy-funded collaboration among NREL, Sandia National Laboratories, Oak Ridge National Laboratory, and Lawrence Berkeley National Laboratory. The primary purpose of this project is to develop novel designs for a next generation of land-based wind turbines that are three-bladed, have a nameplate power of 5 MW, and a rotor diameter of 206 m. These quantities yield a specific power of 150 W m^{-2} , which is lower than today’s average, but it is expected to represent the future generation of turbines tailored for mature markets.

The work presented in Bortolotti et al.⁴² highlights five designs of the BAR turbine. In this work, the BAR upwind-segmented-carbon fiber blade (BAR-USC) design is chosen because it is considered the closest to a conventional rotor design. This is an upwind rotor design with blades that mount a spanwise joint of 2000 kg at 70% of the blade span and has spar caps made of pultruded carbon fiber. The characteristics of the carbon fiber laminate are discussed in Camarena et al.⁴³ The rotor is designed for a tip-speed ratio of 10.5, has a rated wind speed of 8.1 m s^{-1} , and a rated rotor speed of 7.9 rpm. The blades adopt the airfoils listed in Table 1. In between the actual airfoils, coordinates and polars along the span are generated using a piecewise shape-preserving spline interpolation.

As discussed in Bortolotti et al.,⁴² the design of the BAR-USC blade is primarily driven by ultimate tip deflection (and related tower clearance). This constraint is seen as representative of modern land-based wind turbine rotors. Maximum blade tip deflections are generated during standard operation in extreme turbulence. Bortolotti et al. also note that large loads are often seen because of the low pitch rates that are associated with such large and flexible wind turbine blades. Not only are blades on the order of 100-m long subject to large 1P loading but they are also subject to additional “extreme” loads that could be induced by wind gust events. Slow actuator rates and a desire to minimize pitch actuator duty cycles make it challenging to address both types of load events. The adoption of trailing edge flaps represents an opportunity to address both challenges at the same time.

4 | CONTROLLER FORMULATION

ROSCO, an open-source controller that is actively developed at NREL for research purposes, is used extensively throughout the BAR project. The baseline ROSCO controller consists of a below-rated generator torque controller and an above-rated blade pitch controller. In this study, the generator torque controller is unmodified from that used in the study by Bortolotti et al.⁴² In an initial step for this study, the blade pitch controller is

TABLE 1 Airfoils used along the span of the blades of the BAR-USC wind turbine.

Airfoil	Thickness (%)	Spanwise location (%)
Cylinder	100	0-2
SNL-FFA-W3-500	50	18
FFA-W3-360	36	29.6
FFA-W3-301	30.1	48.2
FFA-W3-241	24.1	76.5
FFA-W3-211	21.1	87.2
NACA63-618	18	100

retuned using the WEIS optimization framework. The pitch controller proportional and integral gains are optimized to minimize tower base fore-aft DELs without violating a 20% rotor overspeed constraint. This design objective, as suggested by Zalkind et al.,⁴⁴ was chosen to ensure sufficient rotor speed tracking without imposing unnecessary loading on the wind turbine tower. The resulting tuning parameters from this optimization are shown in Table 2, and the complete DLC results from the BAR-USC rotor with this controller tuning are used as the baseline in Section 6.6. Note that ROSCO's variable-speed torque controller was not optimized for this work, as the blade dynamics relevant to trailing-edge flap controller tuning are not significantly influenced by the generator torque control settings. A peak shaving percentage (p_s) of 20% used during the development of the BAR-USC rotor is also maintained in this work unless otherwise noted. The peak shaving percentage, p_s , represents the reduction of the peak thrust achieved by imposing a minimum pitch schedule on the collective blade pitch controller. This means that, when $p_s = 20\%$, the steady state maximum thrust is 20% less than the maximum thrust without peak shaving enabled. Further detail on this logic is provided by Abbas et al.²⁰

4.1 | Trailing edge flap controller

A primary contribution of this work is the introduction of a standardized trailing edge flap controller and corresponding generic tuning method to facilitate the co-design optimization of wind turbine blades with flaps. A PI trailing edge flap controller is implemented to control the distributed aerodynamic control devices, where the blade root bending moment measurement signal ΔM_r is fed back in an attempt to mitigate blade tip deflections. This primary PI controller can be described by

$$\Delta\delta_n(t) = k_{p,f}\Delta M_{r,n} + \frac{k_{i2,f}}{\tau_f} \int_0^t \Delta M_{r,n} dt, \quad (1)$$

where $k_{p,f}$ is a feedback gain and τ_f is a time constant. The subscript n denotes each individual blade for $n \in \{1,2,3\}$. For the remainder of this paper, this subscript is removed for simplicity and the feedback control is assumed to act on each blade individually. In addition to the PI control loop, an integrator with gain $k_{i2,f}$ is included to encourage the trailing edge flap to return back to zero. Figure 5 provides a block diagram of the employed controller. The primary feedback control loop also employs a filter (*Filt* in Figure 5) that is described in more detail at the end of this section.

The automated controller tuning formulation is started by defining the angle of the blade chord line relative to the rotor plane as

$$\theta_p(r,\beta) = \beta + \phi(r), \quad (2)$$

TABLE 2 Blade pitch controller tuning parameters and DEL reductions from an optimization of tower base DELs.

Damage equivalent load reduction	ω_{pc} [rad s ⁻¹]	ζ_{pc} [-]
5.6%	0.11	2.42

Note: $\zeta_{pc} > 1$ means that the linearized second-order closed-loop system has two poles on the real axis and is common for large rotors.⁴⁵

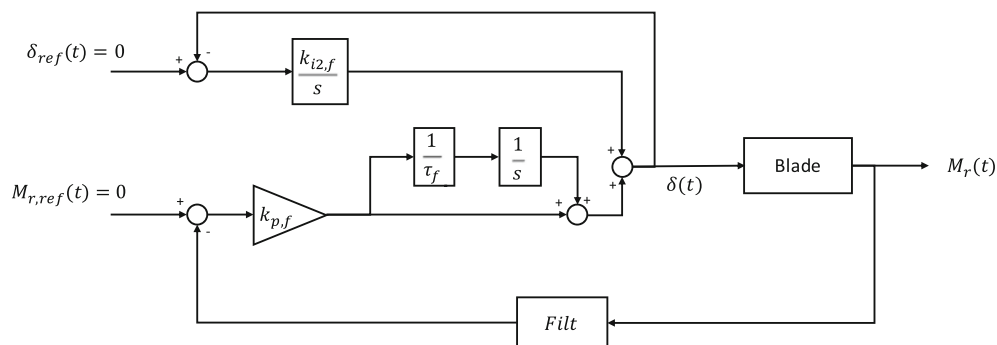


FIGURE 5 A block diagram showing the structure of the trailing edge flap controller for each blade. The proportional and integral gain terms are parameterized by a feedback gain, $k_{p,f}$, and a time constant, τ_f . *Filt* is a filter used to improve the feedback signal and is described more in Figure 6. The input and output of the controller are M_r and δ , which are the blade root bending moment and trailing edge flap angle, respectively. Finally, the subscript *ref* refers to a reference value.

where β is the blade pitch angle, $\phi(r)$ is the blade twist angle, and r is the spanwise distance from the blade root. We then establish the induction factor, a , and define the coefficient of power:

$$C_p = 4a(1 - a)^2, \tag{3}$$

and the tangential flow induction factor:

$$a' = \frac{a(1 - a)}{\lambda^2}. \tag{4}$$

λ is the tip-speed ratio, which is the ratio of the speed of the blade tip to the wind speed, defined as

$$\lambda = \frac{\omega_r R}{U_\infty}, \tag{5}$$

where ω_r is the rotor speed, R is the rotor radius, and U_∞ is the freestream wind speed. We can then calculate the relative wind speed on the blade span as⁴⁶

$$v(r) = \sqrt{U_\infty^2(1 - a)^2 + \omega_r^2 r^2(1 - a')^2}. \tag{6}$$

The normal force on the rotor (i.e., the force in the flap-wise direction) is

$$F_n(r, \beta) = F_l(r) \cos(\theta_p(r, \beta)) + F_d(r) \sin(\theta_p(r, \beta)), \tag{7}$$

where $F_l(r)$ and $F_d(r)$ are

$$F_l(r) = \frac{1}{2} \rho c(r) v(r)^2 C_l(r, \theta_{aaa}(r)), \quad F_d(r) = \frac{1}{2} \rho c(r) v(r)^2 C_d(r, \theta_{aaa}(r)), \tag{8}$$

and correspond to the lift and drag forces, respectively. In (8), ρ is the air density, c is the blade chord, $\theta_{aaa}(r)$ is the angle of attack along the blade span, and C_l and C_d are the lift and drag coefficients, respectively. Note that in (6) and (8), the wind speed is a function of the blade radius, r , because it is relative to the blade's reference frame, not the rotor. We then define the nominal flapwise blade root bending moment as

$$M_r^*(\beta) = \int_0^R F_n(r, \beta) r dr = \int_0^R (F_l(r) \cos(\theta_p(r, \beta)) r + F_d(r) \sin(\theta_p(r, \beta)) r) dr \tag{9}$$

which, upon substitution of the expressions from (8), becomes

$$M_r^*(\beta) = \int_0^R \left(\frac{1}{2} \rho c(r) v(r)^2 r C_l(r, \theta_{aaa}(r)) \cos(\theta_p(r, \beta)) + \frac{1}{2} \rho c(r) v(r)^2 r C_d(r, \theta_{aaa}(r)) \sin(\theta_p(r, \beta)) \right) dr. \tag{10}$$

Additionally, when considering the effect of the blade flap angle, δ , on the lift and drag coefficients of the blade, we assume a consistent linear relationship across operational angles of attack for small flap angles such that

$$\Delta C_l(r) = K_{C_l}(r) \Delta \delta \quad \Delta C_d(r) = K_{C_d}(r) \Delta \delta. \tag{11}$$

Although this assumption is not always valid, we claim that it is reasonable for thin airfoils located in the outer portion of the blade and for angles of attack corresponding to normal operation. For these two conditions, C_l is fairly linear and C_d is fairly flat. The assumption allows us to then define $M_r(\beta)$ with a flap deflection $\Delta \delta$ as

$$M_r(\beta) = \int_0^R \left(\frac{1}{2} \rho c(r) v(r)^2 r (C_l(r, \theta_{\text{aod}}(r)) + K_{C_l}(r) \Delta \delta) \cos(\theta_p(r, \beta)) + \frac{1}{2} \rho c(r) v(r)^2 r (C_d(r, \theta_{\text{aod}}(r)) + K_{C_d}(r) \Delta \delta) \sin(\theta_p(r, \beta)) \right) dr. \quad (12)$$

We do not necessarily want to consider the absolute blade root bending moment; rather, we want to consider the deviation from the expected blade root bending moment for a given flap angle. For control purposes, we would like to understand the blade root bending moment's sensitivity to a change in flap angle. We are therefore interested in a difference from nominal $\Delta M_r(\beta) = M_r(\beta) - M_r^*(\beta)$, so

$$\Delta M_r(\beta) = \int_0^R \left(\frac{1}{2} \rho c(r) v(r)^2 r K_{C_l}(r) \Delta \delta \cos(\theta_p(r, \beta)) + \frac{1}{2} \rho c(r) v(r)^2 r K_{C_d}(r) \Delta \delta \sin(\theta_p(r, \beta)) \right) dr. \quad (13)$$

If we define

$$C_1(\beta) = \int_0^R \left(\frac{1}{2} \rho c(r) v(r)^2 r K_{C_l}(r) \cos(\theta_p(r, \beta)) \right) dr, \quad (14)$$

$$C_2(\beta) = \int_0^R \left(\frac{1}{2} \rho c(r) v(r)^2 r K_{C_d}(r) \sin(\theta_p(r, \beta)) \right) dr, \quad (15)$$

and combine (14) and (15) so that

$$\kappa(\beta) = C_1(\beta) + C_2(\beta), \quad (16)$$

we establish a linear relationship between a perturbation of the trailing edge flap angle, $\Delta \delta$, to a perturbation of the blade root bending moment, $\Delta M_r(\beta)$:

$$\Delta M_r(\beta) = \kappa(\beta) \Delta \delta. \quad (17)$$

This derivation of a so-called ‘‘flap efficacy’’ factor $\kappa(\beta)$ provides us with the necessary term to apply the generalized tuning process such that

$$k_{p,f}(\beta) = \frac{\alpha_f}{\kappa(\beta)}. \quad (18)$$

In (18), α_f is a unitless normalized gain tuning parameter. By defining the proportional feedback gain, $k_{p,f}(\beta)$, as the ratio of a normalized tuning parameter over $\kappa(\beta)$, the single parameter, α_f , can be tuned regardless of flap size, location, and perturbation limits. Moreover, when tuning the controller, the same α_f and τ_f can be prescribed to achieve the same dynamic response for different trailing edge flap configurations. Finally, this derivation modifies the flap controller equation (1) to be

$$\Delta \delta(t, \beta) = \frac{\alpha_f}{\kappa(\beta)} \Delta M_r + \frac{\alpha_f}{\kappa(\beta) \tau_f} \int_0^t \Delta M_r dt. \quad (19)$$

Though this formulation is clearly rooted in highly simplified blade element momentum (BEM) formulations and may not be completely representative of a more realistic higher order model, we suggest that this formulation is sufficient for consistent controller tuning and parameterization purposes. Section 4.2 demonstrates the consistency of this tuning method for varied flap configurations in nonlinear simulations; such consistency is, ultimately, the goal of this tuning procedure. Additionally, in practice, it is found that the changes in $\kappa(\beta)$ across blade pitch angles have a minimal effect on controller performance. For the controller used in this work, (18) is evaluated with $k_{p,f}(\beta = 0)$. This simplifies the controller implementation as well because a pitch-dependent gain schedule is not needed. In future work, the efficacy of a pitch-dependent gain schedule should be considered to better understand the impact across the operating region.

In addition to the PI controller for the trailing edge flaps, a feedback filter is implemented to smooth the flap response signal and ensure a zero mean for the measurement error. A Bode diagram of this filter is shown in Figure 6. The feedback filter is the combination of a first-order high-pass filter, a notch filter, and a second-order low-pass filter, as defined in Appendix A. The high-pass filter cutoff frequency is generically placed at 0.1 rad s^{-1} and serves the purpose of removing the expected DC gain (i.e., steady-state offset) of the blade root bending moment signal and allows us to define $M_{r,ref}(t) = 0$. The notch filter is used to prevent the excitation of the blade flapwise natural frequency, ω_{flp} . The low-pass filter has a damping ratio of 0.7 and a cutoff frequency generically placed at three times the blade first flapwise natural frequency to remove high-frequency noise from the feedback signal.

4.2 | Controller verification

As shown in the derivation of (18), the automatic controller tuning is designed such that the controller tuning parameters will provide the same closed-loop blade response for any change in flap size or placements, assuming that the response is achievable. An extreme operating gust is used to demonstrate the automated flap controller tuning methods used in this study, as shown in Figure 7.

In Figure 7, the blade responses for three different flap sizes are shown: 5%, 10%, and 20% of the total blade span. Each flap is centered at 75% of the blade length. The PI controller for each flap, (19), is tuned with $\alpha_f = 0.1$ and $\tau_f = 10 \text{ s}$. This tuning is chosen for this simple demonstration, and it is not assumed that this is the best performance that the flap controller can achieve. Additionally, the simulation environment is greatly

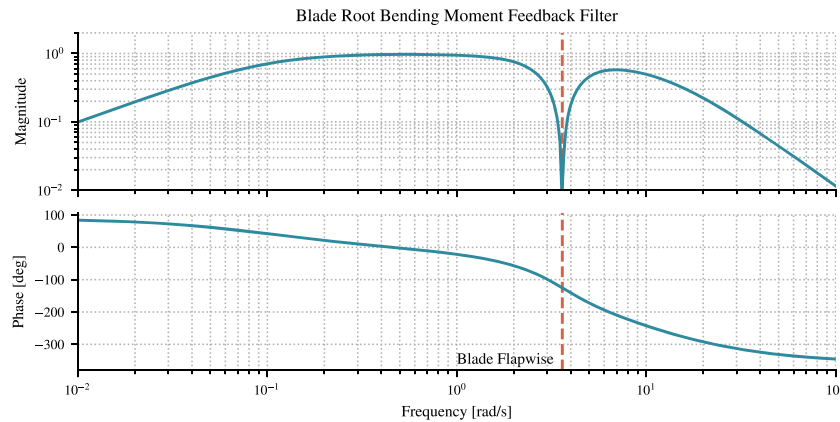


FIGURE 6 A Bode diagram of the feedback filter on the measured blade root bending moment as shown in Figure 5. The filter is a combination of the high-pass, notch, and low-pass filters described in Appendix A. The high-pass filter has a cutoff frequency of 0.1 rad s^{-1} , and the low-pass filter has a cutoff frequency of 10.82 rad s^{-1} , or three times the blade flapwise natural frequency. The dashed red line denotes the blade flapwise natural frequency of 3.61 rad s^{-1} and is the frequency where the notch filter is located.

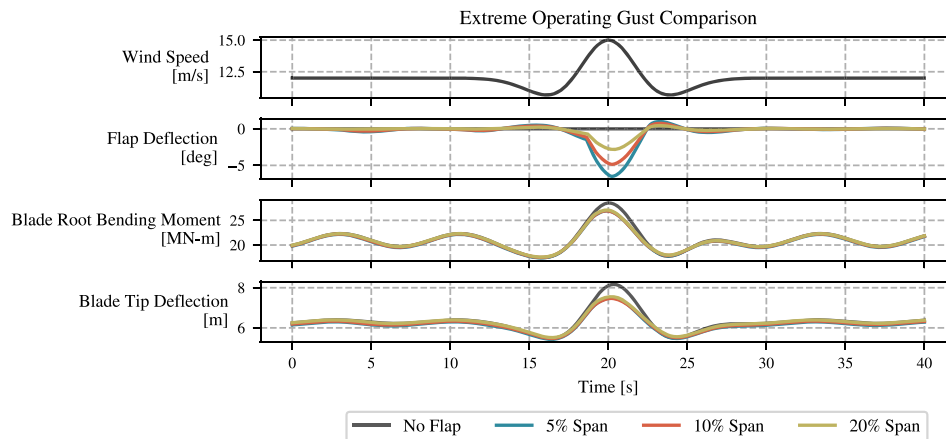


FIGURE 7 A comparison of blade responses of the BAR-USC wind turbine simulated in OpenFAST with different flap sizes and the same controller input tuning parameters, $\alpha_f = 0.1$ and $\tau_f = 10 \text{ s}$. Each trailing edge flap is centered at 75% of the blade span.

simplified to isolate the effect of the trailing edge flap on the blade response. No wind shear or tower influences are modeled. The rotor speed, blade pitch, and generator torque are all fixed at their respective above-rated steady-state values such that they do not impose additional dynamics on the system. Similarly, the tower is modeled as rigid, and only the blade degrees of freedom are enabled. Due to gravitational effects and the 4° of pre-cone on the rotor, a once-per-revolution loading on the blade is seen and is especially obvious in the blade root bending moment response signal. Figure 7 shows consistency in the blade root bending moment and blade tip deflection responses across three different trailing edge flap sizes, though the flap deflections differ significantly. As expected, the smaller the flap is, the more it has to deflect to achieve the blade root bending moment and tip deflection reductions. By using the PI controller tuning formulation in (19) and keeping the parameters α_f and τ_f identical for all three flap sizes, consistency is seen in the blade root bending moment responses. By tuning the controller in this way, the flap size and location can be decoupled from the flap controller tuning parameters during an optimization procedure to help “smooth” the optimization space.

4.3 | A note on individual pitch control

In other studies regarding the efficacy of active flow control devices, such as the work by Lackner and van Kiuk,⁸ IPC is used as a comparison or baseline case. Active flow control is considered to reduce cyclic (i.e., fatigue) loads on the turbine in these studies. For such cases, it is natural to compare the novel technology to IPC, a well-established method of reducing fatigue loads.⁴⁷ Indeed, it has been shown that active flow control can have an effect on fatigue loads; however, for the purpose of this work, the *extreme* loads are of particular interest. This is because, for the BAR-USC rotor, the largest loads on the system correspond to the largest blade tip deflections and have proven to be primary design constraints.⁴³ The control methods presented in Section 4.1 can be used to reduce these extreme loads at frequencies that might not specifically be the once- or three-times-per revolution harmonics.

To demonstrate this difference, IPC and trailing edge flap controller methods are compared in this section. A block diagram of the IPC schema used is shown in Figure 8. This IPC routine is used to specifically target the 1P loads that are often seen on wind turbine rotors due to turbulence, wind shear, yaw misalignment, and tower shadow.^{47,48}

The control algorithm consists of three primary steps. First, the blade root bending moments ($M_n(t)$, where n represents the blade number) are filtered and passed through the forward multiblade coordinate (MBC) transformation,⁴⁹ $T_1(\psi)$, such that

$$\begin{bmatrix} M_t(t) \\ M_y(t) \end{bmatrix} = \frac{2}{3} \underbrace{\begin{bmatrix} \cos(\psi_1) & \cos(\psi_2) & \cos(\psi_3) \\ \sin(\psi_1) & \sin(\psi_2) & \sin(\psi_3) \end{bmatrix}}_{T_1(\psi)} \begin{bmatrix} M_1(t) \\ M_2(t) \\ M_3(t) \end{bmatrix}. \quad (20)$$

In (20), $M_t(t)$ and $M_y(t)$ are the fixed-frame tilt and yaw moments on the rotor, and ψ_n is the azimuth angle of the blade, $n \in \{1,2,3\}$. Next, the fixed-frame tilt and yaw moments are individually passed to the PI controllers with proportional gain, $k_{p,ipc}$, and integral gain, $k_{i,ipc}$. The PI gains are the same values for both the tilt and yaw components of the rotor. The outputs from the PI controllers, $\theta_t(t)$ and $\theta_y(t)$, are then passed to the inverse MBC transform, $T_1^{-1}(\psi)$, such that the IPC contribution to each blade pitch angle is calculated as

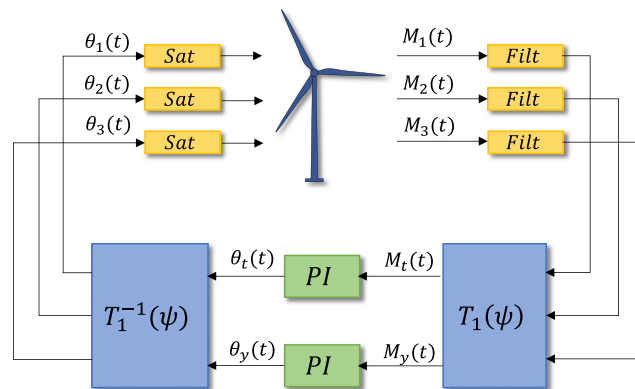


FIGURE 8 A block diagram of the IPC algorithm used in this study. Each filter, denoted by *Filt*, is a band-pass filter, as defined in Appendix A, centered around the once-per-revolution frequency of the rotor with a damping ratio of 0.7. The forward and inverse MBC transformations are denoted by $T_1(\psi)$ and $T_1^{-1}(\psi)$, respectively. The PI controllers that regulate the fixed-frame tilt and yaw moments are represented by the *PI* blocks and share the same gains, and the final pitch signal $\theta_n(t)$ for $n \in \{1,2,3\}$ is passed through a saturation block denoted by *Sat*, which limits the absolute IPC controller contribution.

$$\begin{bmatrix} \theta_1(t) \\ \theta_2(t) \\ \theta_3(t) \end{bmatrix} = \underbrace{\begin{bmatrix} \cos\psi_1 & \sin\psi_1 \\ \cos\psi_2 & \sin\psi_2 \\ \cos\psi_3 & \sin\psi_3 \end{bmatrix}}_{T_1^{-1}(\psi)} \begin{bmatrix} \theta_t(t) \\ \theta_y(t) \end{bmatrix}. \tag{21}$$

The individual blade pitch contributions to the pitch command, $\theta_n(t)$ for $n \in \{1,2,3\}$, are limited by a maximum allowable amount through a saturation block and added to the collective pitch command.

As noted earlier, this type of pitch controller is designed to specifically target the cyclic loading of the wind turbine rotor. It does not, however, address extreme loading events that could result from a turbulent gust. The DAC controller methodology presented in this work (as shown in Figure 5) is designed to regulate the bending moment of a blade, so it can help reduce both cyclic and noncyclic loading of the blades.

Figure 9 shows an example time history for three different simulation cases: collective pitch control (CPC), CPC with IPC, and CPC with DAC devices using the proposed PI controller logic. A 10-min simulation in extreme turbulence with an average inflow wind speed of 12 m s^{-1} is simulated, and a 10-second interval of this is plotted. The IPC PI gains were chosen based on an optimization with the objective of minimizing the DEL of the blade root bending moments. The trailing edge flap controller was optimized to minimize the maximum blade tip deflection. For this demonstration, the trailing edge flap spans from 70% to 90% of the total blade span. The CPC and CPC with IPC see a relatively large blade deflection at $\sim 4 \text{ s}$, whereas the trailing edge flap is able to mitigate the maximum deflection by 15% (at $\sim 6 \text{ s}$) compared to the CPC case and 8.5% compared to the CPC with IPC case. Further, the IPC commands a pitch that violates the pitch actuator's rate limit in the initial pitch maneuver (from 2 to $\sim 4 \text{ s}$), highlighting the challenges introduced with such large rotors.

Figure 10 shows a visualization of the differences in the power-spectral density of the blade tip deflections from the same 10-min simulation. The trailing edge flap is seen to reduce the blade tip deflections in lower frequencies, including the 1P frequency and the blade first flapwise frequency, whereas the IPC only reduces 1P excitation. Notably, blade tip deflections are excited at $\sim 1 \text{ Hz}$ when using trailing edge flaps, which is attributed to the first blade edgewise whirling mode.⁵⁰ Additionally, the reduction in 1P excitation when using the IPC and trailing edge flaps is not as significant as might be expected. This is most strongly attributed to the rate limitations for the IPC and control authority of the trailing edge flap.

5 | PARAMETER STUDIES

To visualize the optimization space for trailing edge flaps, controller parameter studies are provided for three different trailing edge flap locations on the baseline BAR-USC wind turbine rotor. For each parameter study, 215 controller tuning parameter combinations are selected using a Latin Hypercube sampling⁵¹ method. For all cases, the trailing edge flap spans 20% of the blade chord and length and is capable of actuating $\pm 10^\circ$. These two values were adopted as a compromise between control authority and size for the flaps. Future studies should look into flaps that extend differently along the chord and deflect more. Note however that both numerical and manufacturing challenges will arise. On the numerical

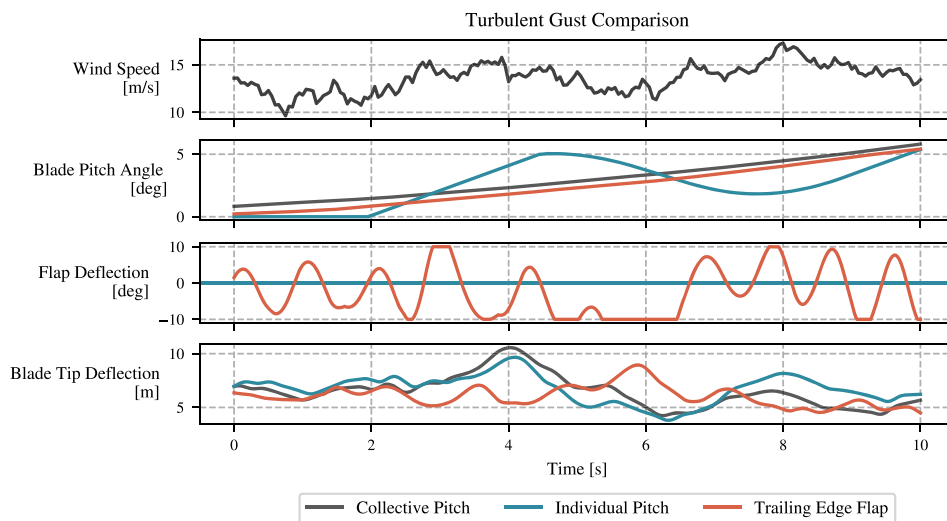


FIGURE 9 Turbulent inflow comparison for the BAR-USC rotor using CPC, CPC with IPC, and CPC with trailing edge flaps. Note that the blade pitch angle for the CPC and CPC with trailing edge flaps cases is nonzero at the beginning of the interval because of slight dynamics differences earlier in the simulation.

side, XFOIL will struggle to converge for both wider flaps and flaps that deflect more than $\pm 10^\circ$. On the manufacturing side, the chordwise extent of flaps will need to compromise with the loss of edgewise stiffness caused by the presence of a nonstructural flap.

The ends of the flaps are placed at 85%, 90%, and 95% of the blade span, effectively perturbing the three outer-most airfoils FFA-W3-241, FFA-W3-211, and NACA63-618 (see Table 1). For each flap setup, polars are computed in XFOIL estimating the Reynolds number at rated wind speed along the span of the blade. Next, for each combination of α_f and τ_f , OpenFAST is run for the wind conditions shown in Table 3. The environmental conditions in DLC 1.3, as defined by the International Electrotechnical Commission,⁵² are used because the design-driving maximum tip deflections presented by Bortolotti et al⁴² were shown to happen in these environmental conditions.

Figure 11 displays the maximum value of the tip deflection ratio seen from all inflow wind conditions in Table 3 across the flap controller design parameters. The tip deflection ratio is defined as the ratio of the measured tip deflection, $\delta_{tip}(t)$, with respect to the undeflected blade tower clearance, Δ_{blt} , times a safety factor, γ_s , or

$$\text{TipDeflectionRatio} = \frac{\delta_{tip}(t)\gamma_s}{\Delta_{blt}}. \quad (22)$$

The safety factor, γ_s , used in this work is 1.4175, which accounts for potential errors in the modeling of loads ($\gamma = 1.35$) and material stiffness ($\gamma = 1.05$). By this definition, a maximum tip deflection ratio greater than one implies that the blade tip has exceeded its maximum allowable deflection and needs to be made stiffer.

As shown in Figure 11, the controller tuning values and trailing edge flap location have a clear effect on the tip deflection ratio. For example, with the flap end at 95% of the blade span, a good choice of controller tuning parameter can provide a tip deflection that is approximately 20% less than the tip deflection when no trailing edge flap controller is used.

Here, we note the large amount of computational power needed to generate such results. As shown in Table 3, to evaluate a single point on the surfaces in Figure 11, 14 individual OpenFAST simulations must be run. These OpenFAST simulations are each 10-min simulations, with an additional initial 2 min of simulation time to remove transient behavior. Using NREL's HPC system, Eagle, each 12-min simulation can take approximately 5–10 min of elapsed clock time, primarily depending on the time step used, the number of active degrees of freedom, and the settings by which OpenFAST is compiled. This means that simply running the OpenFAST simulations for a single point on these surfaces can take nearly 1 h of CPU time. In addition to running OpenFAST, further overhead is necessary to generate the wind turbine model, run XFOIL, tune ROSCO, and post-process the results; thus, for a 10×10 surface, the simulations can take more than 100 h of computational time—and the number of OpenFAST simulations scales exponentially as degrees of freedom are added. For a simple parameter study such as that presented in this section, the OpenFAST simulations can be parallelized, and the surfaces can be generated within a reasonable amount of time. If one wishes, however, to find the minimum value within a higher-dimensional design variable space, a parameter sweep is simply impractical.

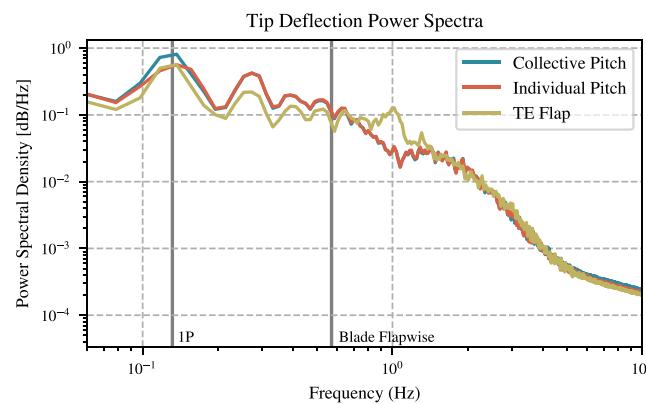


FIGURE 10 The power spectral density of the blade tip deflection for a single 10-min, 12 m s^{-1} turbulent wind simulation with three different controller configurations. The 1P and blade flap-wise frequencies are also marked by solid vertical lines and labeled for reference.

TABLE 3 Wind conditions for each evaluation of the data points used in the controller parameter studies shown in Figure 11.

Representative DLC	Wind speeds [m s^{-1}]	Number of turbulence seeds
1.3	[8, 10, 12, 14, 16, 18, 20]	2

Note: The representative DLC, as defined by the IEC, is used to define the environmental conditions (e.g., turbulence intensity and wind shear). Two turbulence seeds are chosen per wind speed in order to balance the need to capture representative turbulence across design-driving wind conditions with computational efficiency.

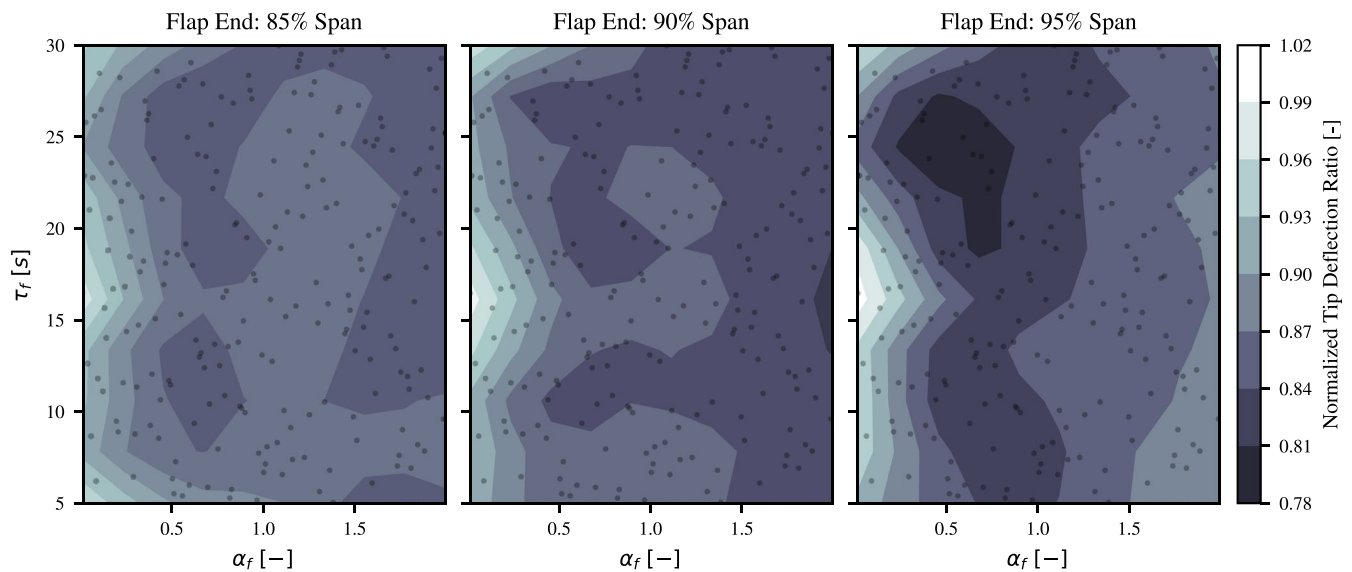


FIGURE 11 Contour plots showing the effect of the trailing edge flap controller tuning parameters on tip deflection. The lowest tip deflections are seen with the trailing edge flap end at 95% of the blade span. The small gray dots portray the data points used to generate the corresponding surfaces. The tip deflection ratio is normalized by the tip deflection ratio of the blade without trailing edge flaps.

6 | DESIGN OPTIMIZATION

A valid alternative to parameter sweeps is automatic design optimization. In this section, a series of design optimization studies is conducted to demonstrate the WEIS optimization framework and the effect of implementing trailing edge flaps and their controller. For each optimization, the following are kept constant:

1. The flap covers 20% of the blade span and chord, regardless of location on the blade. This size flap is chosen to be consistent with similar studies seen in the literature and because it is shown that they are large enough to significantly reduce blade root bending moments and blade tip deflections.
2. The flap is allowed to actuate up to $\pm 10^\circ$ from neutral.
3. No actuator model is used for the flaps, and their actuation is not rate-limit constrained. Though this might be a physically unrealizable assumption, the actuator is left unconstrained to provide a high-level understanding of what active aerodynamic control devices could possibly provide to the wind turbine rotor design if the requisite technology is developed.
4. The maximum blade tip deflection is constrained to be 13.21 m from the blade's neutral position. The proximity to this constraint is the so-called tip deflection ratio, as defined by Equation (22), where a tip deflection ratio of one signifies the maximum allowable tip deflection.

The work by Bortolotti et al.⁴² combined a number of tool sets to finalize the BAR-USC rotor. In that work, the structural dynamics submodule BeamDyn⁵³ was used in OpenFAST to calculate the loads and maximum blade tip deflections, and Numerical Manufacturing and Design (NuMAD)⁵⁴ was used to determine the blade's final structural design. Though BeamDyn is higher fidelity and more accurate than the simpler structural dynamics submodule ElastoDyn (which does not include elastic torsion), the use of BeamDyn during large-scale optimizations is currently impractical due to the significantly increased computational cost. Similarly, the automated implementation of NuMAD within the WEIS framework is not practical either. The most extreme blade tip deflection seen on the BAR-USC rotor during DLC 1.3 when using ElastoDyn does not reach the maximum allowable deflection. For this reason, an initial blade mass optimization is completed to obtain the “best” blade without trailing edge flaps using the WEIS design framework.

6.1 | Optimization algorithm

Two primary optimization methods are considered in this work. For optimizations that have objective functions that *are not* explicitly defined by the outputs of OpenFAST simulations and are analytically calculated within the WEIS framework (e.g., blade mass), the NLOpt implementation⁵⁵

of the gradient-based, globally convergent method of moving asymptotes (MMA)⁵⁶ is found to perform well. When this is the case, the finite difference evaluations are trusted to provide reliable and accurate gradient approximations that can result in quick and consistent convergence. From a candidate point, p_i , MMA forms an approximation of a gradient using finite-difference evaluations along with a quadratic “penalty” on the approximation. This penalty is applied to the approximation to find a new candidate point, p_{i+1} . For an unconstrained problem, this process is repeated for i iterations until a tolerance is satisfied such that $|p_{n_i} - p_{n_i-1}| < tol$. For a problem with inequality constraints, similar tolerances are placed on the constraint function evaluations. The MMA algorithm guarantees convergence to a minima, though it is not guaranteed that the minima is a global minima.

For optimizations that have objective functions that are explicitly defined by the outputs of OpenFAST simulations (e.g., blade tip deflection), the NLOpt implementation of the Constrained Optimization BY Linear Approximation (COBYLA) algorithm⁵⁷ is used. COBYLA is a derivative-free optimization method for problems with nonlinear inequality and equality constraints. Through a simplex of $n + 1$ points, where n is the dimension of the problem, the algorithm constructs a linear approximation of the objective and constraints. At each iteration, these approximations are optimized within a given trust region of decreasing size. In doing so, COBYLA can converge upon a minimum (however, there is still no guarantee that it is a global minima). By constructing a simplex of decreasing size, COBYLA is less sensitive to the inherent noise that comes with using OpenFAST outputs to evaluate the objective function. This inherent noise is attributed to the extremely complex nature of the OpenFAST simulations, which are dependent on numerous nonlinear influences from the turbine, controller, environmental conditions, and even the computational setup.

Other gradient-based optimization routines, such as SLSQP⁵⁸ and SNOPT,⁵⁹ were explored, but they were found to struggle with nonlinear constraints that are defined by OpenFAST outputs. The use of global optimization methods (such as a random search method or genetic algorithm) were not considered because they generally require significantly more computational time.

For consistency, all optimization results presented in this work are found using the derivative-free COBYLA algorithm. Though MMA was found to converge more quickly in some instances, the final results were of negligible difference to those found using COBYLA. For this paper, the use of a single optimization algorithm makes for more consistent comparisons across designs.

6.2 | Optimization formulation

The generic form of the nonlinear, constrained, optimization problem solved in this work can be represented by

$$\begin{aligned} & \underset{x}{\text{minimize}} && f(x) \\ & \text{subject to:} && c(x) \leq b \\ & && x_{min} \leq x \leq x_{max}, \end{aligned} \quad (23)$$

where x is a vector of the design variables that is used to calculate the objective function, $f(x)$, and constraint function, $c(x)$. An upper bound, b , is imposed on the inequality constraint, and the design variables are constrained by minimum and maximum bounds represented by x_{min} and x_{max} , respectively.

In the series of optimizations conducted for this work, first, a blade mass optimization is done to demonstrate a sequential design approach and because the wind turbine design methods within WEIS are not identical to those used by Bortolotti et al.⁴² Then, trailing edge flaps are added to the optimized blades to understand how much they can reduce maximum blade tip deflections. Next, a co-optimization of blade spar cap thickness and trailing edge flaps is conducted to demonstrate the mass savings that are possible when flaps are introduced. Finally, a simple aerostructural optimization is conducted by adding the rotor diameter among the design variables in an attempt to increase the AEP and reduce the LCOE of the wind turbine. The optimization does not have the ability to do a high-fidelity structural redesign of the blades, but it is here assumed sufficient to capture relative trends.

During most optimizations completed for this study, the set of inflow conditions in Table 3 is simulated during each function evaluation. For the optimization including the rotor diameter, a larger set of inflow conditions is necessary, as is discussed in Section 6.5. An overview of the studies discussed in this work is shown in Table 4, and the following sections describe them in more detail.

For each optimization, a set of DLCs is used to calculate the relevant OpenFAST-related optimization parameters. Note that we use “DLCs” to describe the representative environmental conditions based on the full design load cases defined by the IEC.⁵² In this analysis, we simulate a subset of the required wind speeds and seeds needed for a complete DLC analysis as defined by the IEC. Results from DLC 1.2 are used to estimate the AEP and DELs. Results from DLC 1.3 are used to calculate the maximum load events and blade tip deflections. If the optimization objectives and constraints are not dependent on any values from DLC 1.2, only DLC 1.3 is run. Table 5 shows the specific wind speeds used for each DLC.

TABLE 4 A matrix of the various optimization studies done using WEIS and presented in this work.

Study type	Name	Design variables	Constraints	Objective	# Design variables	Baseline turbine				
Blade optimization	BAR-m ₀	X	- 0 0	Blade mass	4	BAR-USC				
Flap optimization	BAR-m ₀ f ₀	- X	- 0 0	Max tip deflection	2	BAR-m ₀				
	BAR-m ₀ f ₁	- X	- 0 0	Max tip deflection	3	BAR-m ₀				
Rotor optimization	BAR-m ₁ f ₂	X X	- 0 0	LCOE	7	BAR-m ₀				
	BAR-m ₁ d ₀	X	X 0 0	LCOE	6	BAR-m ₀				
	BAR-m ₁ d ₀ f ₃	X X	X 0 0	LCOE	9	BAR-m ₀				
		Spar cap thickness	α_f and τ_f	Flap position	Rotor diameter	p_s	Blade strain	Tip deflection ratio	Flap position	Tower loads

Note: Active design variables are denoted by an “X” and active constraints are denoted by an “o.” For the design variables, the spar cap thickness is optimized at four equally spaced spanwise locations along the blade. The flap controller tuning parameters are α_f and τ_f , and p_s is the peak shaving percentage. For the constraints, the tower loads include both fatigue (i.e., damage equivalent loads) and extreme (i.e., maximum tower fore-aft bending moments) loads.

6.3 | Blade mass optimization—BAR- m_0

For the blade mass optimization, the optimization objective $f(x)$ to be minimized is the blade mass. To do so, the thickness of the spar caps along the span of the blade is parameterized by six equally spaced nodes. The central four nodes (excluding the blade root and tip) are allowed to be changed by the optimizer. This, in turn, has a direct effect on the mass and stiffness parameters. WEIS is used to recalculate the blade's structural makeup, and OpenFAST is used to run full aero-servo-elastic simulations. The blade tip deflection and maximum strains are calculated using the OpenFAST simulations for a subsequent load analysis of the blade. The tip deflection ratio is not allowed to exceed 1, and the maximum strains along the blade are not allowed to exceed 3500 microstrains ($\mu\epsilon$). The results from this initial blade mass optimization are shown in Figure 12.

In this optimization, the blade mass is reduced from 51.49 metric tons to 49.52—a reduction of 3.8%. Using the blade cost model in WEIS, this equates to an overall blade cost reduction of 6.6%. The larger cost saving is explained by the high unit cost of the pultruded carbon fiber adopted in the spar caps. For the remainder of the optimizations discussed in this work, the blade from this optimization is used as the initial blade and is referred to as BAR- m_0 .

6.4 | Flap controller optimizations

Trailing edge flaps are then added to the BAR- m_0 blades. In an initial study, the trailing edge flaps span from 70% to 90% of the blade, and the controller parameters, α_f and τ_f , are optimized. As a reminder, the trailing edge flap always covers 20% of the blade chord and is allowed to actuate $\pm 10^\circ$ from neutral. In this flap controller optimization, the trailing edge flaps can reduce the overall tip deflection by 16% compared to the rotor operating in the same conditions without trailing edge flaps. The convergence of this optimization, named BAR- m_0f_0 , is shown by the blue line in Figure 13.

Next, the position of the trailing edge flap is additionally allowed to be varied by the optimizer for the BAR- m_0f_1 optimization. By including the flap's position in the optimization problem, we can reduce the maximum tip deflection by nearly 7% more than if we only optimize the controller. The convergence of the BAR- m_0f_1 optimization is shown by the red line in Figure 13. As expected from the parameter studies shown in Figure 11, the trailing edge flap is most effective near the blade tip, though the optimizer chooses an end location of the flap of 97.5%, just shy of the end of the blade. With a co-optimization of the trailing edge flap position and controller parameters, the overall blade tip deflection can be reduced by 21.5%. The combination of BAR- m_0f_0 and BAR- m_0f_1 shows that the inclusion of the trailing edge flaps can reduce the maximum blade tip deflections seen during turbulent wind conditions and suggests that the blade could be redesigned accordingly. Note that the erratic nature of the BAR- m_0f_1 optimization convergence is a poignant graphical representation of the noise characteristic of optimization using dynamic simulation.

TABLE 5 Inflow wind conditions used for the design optimizations.

Representative DLC	Description	Wind speeds	Number of seeds
1.2	Normal turbulence model	[4, 6, 8, 10, 12, 14, 16, 18, 20, 22, 24]	3
1.3	Extreme turbulence model	[4, 6, 8, 10, 12, 14, 16, 18, 20, 22, 24]	6

Note: The IEC DLC with representative environmental conditions (e.g., turbulence intensity and wind shear) are provided. Both DLC 1.2 and 1.3 are considered power production situations.




FIGURE 12 Convergence results from optimizing the spar cap thickness to minimize the blade mass.

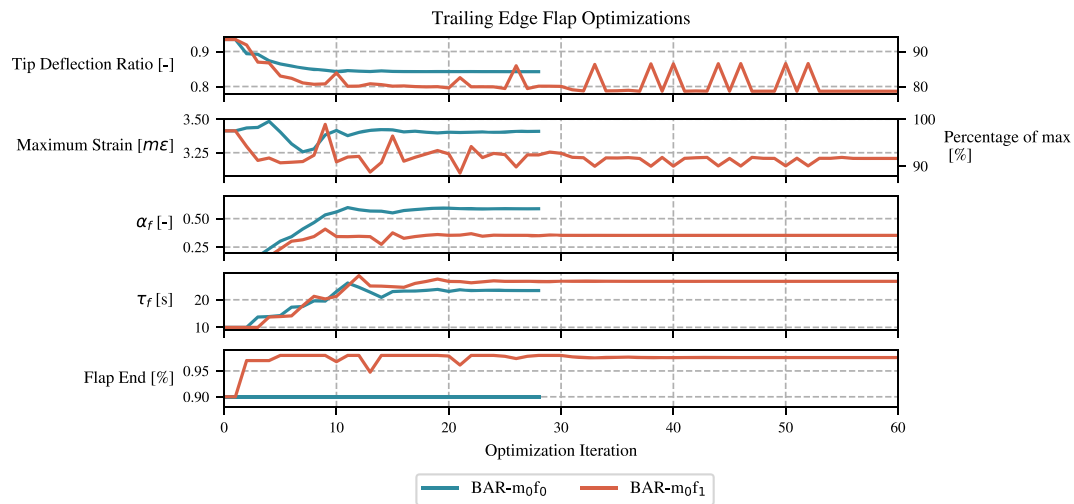


FIGURE 13 Iterative convergence from optimizing the trailing edge flap controller only (BAR-m₀f₀) and co-optimizing the flap controller and position (BAR-m₀f₁)

6.5 | Rotor and controller co-optimization

In Section 6.4, it is shown that the addition of trailing edge flaps can reduce the maximum tip deflections seen during normal wind turbine operation in extreme turbulence. For this reason, optimization studies are run where the trailing edge flap is co-optimized with the blades in an attempt to further reduce blade cost and the overall cost of energy.

First, the spar cap thickness, flap position, and the flap controller are allowed to be varied by the optimizer in optimization BAR-m₁f₂. The tip deflection and blade strains are constrained in the same way that they were for the BAR-m₀ optimization. Convergence results from the BAR-m₁f₂ optimization are shown in Figure 14.

Though the trailing edge flaps reduced the tip deflections on the BAR-m₀ by slightly more than 21%, when the blade mass and flaps are co-optimized, only 3.1% blade mass savings is shown. Using the rotor cost model in WISDEM, this corresponds to a 5.2% blade cost saving compared to BAR-m₀. Because the wind turbine blade is such a small portion of the overall LCOE, and the addition of trailing edge flaps and more flexible blades can reduce power production in near-rated operation, the addition of flaps during blade design equates a mere 0.2% decrease in LCOE. Though it might be possible to slightly increase the cost savings through larger flaps, these initial results suggest that significant blade tip deflection reductions might not directly translate to large LCOE savings. An additional optimization was run with a strain limit constraint of 5000 $\mu\epsilon$ to further explore the design space, and it was found that the blade mass could be reduced by an additional 2%. This, however, still does not translate to significant LCOE reductions.

Similarly to what was done in Bortolotti et al.,²⁵ the rotor diameter is also explored as a design parameter. By allowing the rotor diameter to increase, the AEP and LCOE can be directly influenced by the optimization process. An increased rotor size does, however, result in larger tower base loads. For this study, the maximum tower base fore-aft bending moment and DEL is constrained so that it would be no larger than the loads seen during BAR-m₀. The ROSCO controller tuning for the BAR-USC rotor includes a peak shaving percentage of 20% to reduce rotor thrusts. Because the rotor diameter and trailing edge flaps directly affect the rotor thrust, the peak shaving percentage is also included as a design variable in the rotor diameter optimizations.

To have a substantive comparison, two rotor diameter optimizations are run. In the first optimization, BAR-m₁d₀, the spar cap thickness, rotor diameter, and peak shaving percentage, are allowed to be modified with the objective of minimizing the LCOE. In the second optimization, BAR-m₁d₀f₃, trailing edge flaps are added, and the flap position and controller are also allowed to be optimized. Results from these optimizations are shown in Figure 15.

The mass, stiffness, and spar cap thickness distributions along the blade are shown in Figure 16. Notably, the mass distribution has a clear outlier in all designs that is attributed to the joint used in the segmented blade. There are also noticeable decreases in blade mass and stiffness along the blade, especially ranging from 20%–40% and 70%–90% span. As shown in the optimized blade with a trailing edge flap result (BAR-m₁f₂), the addition of a flap affords the opportunity to reduce the blade mass and stiffness, especially near the center of the blade, without violating the design constraints. There is also a greater thickness farther outboard, as might generally be expected. This can be attributed to the added mass of the joint and its influence on the load analysis in WEIS. There are not significant changes in the blade characteristics when the rotor diameter is added as a design variable (BAR-m₁d₀f₃) compared to the initial BAR-m₀ case because the optimization objective favors increases in the rotor diameter (and subsequently AEP) and, in fact, tends to slightly increase blade mass.

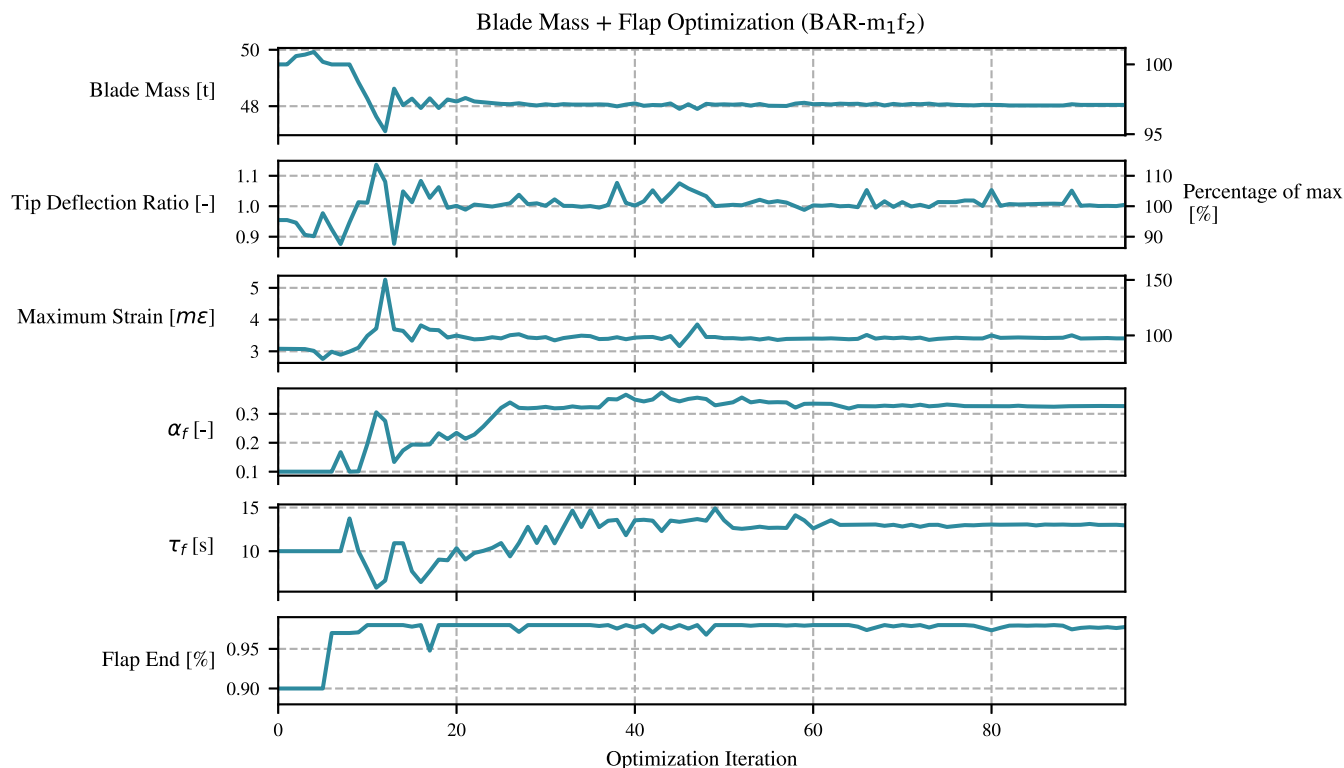


FIGURE 14 Iterative convergence from optimizing the blade mass, trailing edge flap position, and flap control parameters.

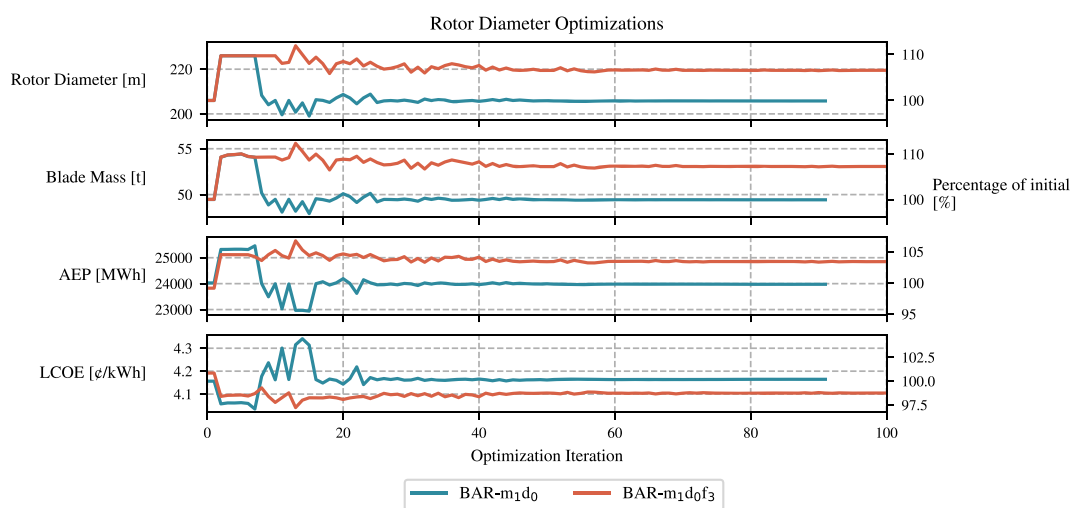


FIGURE 15 Iterative convergence from optimizing the blade mass, trailing edge flap position and controller, rotor diameter, and peak shaving percentage. For clarity, only relevant parameters to this optimization are shown. For both optimizations, the constraints converge to satisfaction, and for the BAR- $m_1d_0f_3$ optimization, the flap controller design variables converge similarly to the BAR- m_1f_2 convergence. Though optimized, the peak shaving percentage changes by only $\sim 2\%$ in both rotor diameter optimizations.

6.6 | Load analysis and discussions

DLCs 1.2 and 1.3 are run for 12 different wind turbulent seeds to provide a final comparison of the resulting rotor models from the optimizations run. Figure 17 provides a comparison of relevant parameters from the blade and rotor optimizations with respect to the results from BAR- m_0 . Note that this comparison uses a different set of wind turbulent seeds than the DLCs used during the optimization process, so slight differences in the output measures are expected. In each comparison, the fatigue loads and energy-related results are calculated using DLC 1.2, and all maximum loads and extreme events (e.g., max tip deflection) are calculated using DLC 1.3.

The results from adding a trailing edge flap to the initial optimized blade and optimizing the spar cap thickness and flap controller to arrive at BAR- m_1f_2 are shown in the blue (left) bars in Figure 17. Though the addition of the flaps was found to reduce both the maximum tip deflections and the blade root DELs, this comes at the cost of a slightly decreased AEP and a very minimal decrease in LCOE. This is unsurprising because the

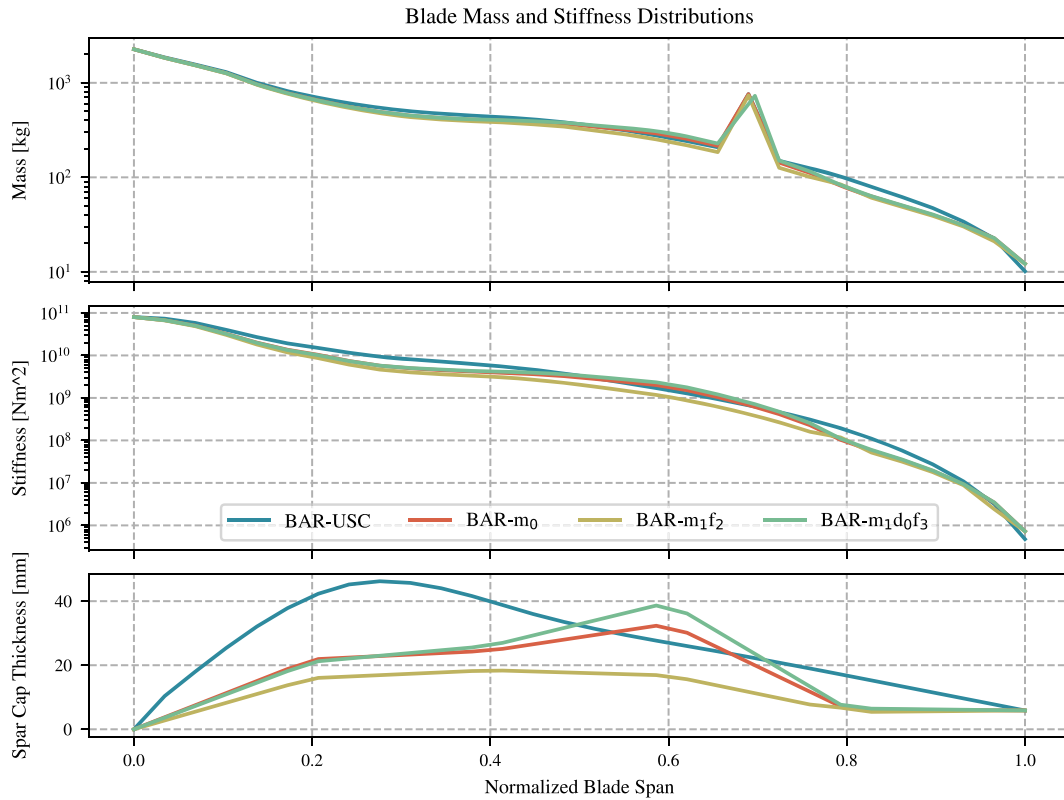


FIGURE 16 Mass and stiffness distributions along the blade span. The vertical grid lines are located at the spanwise optimization nodes, and a spline interpolation is used to determine the complete mass and stiffness distributions. Note that the outlier in the mass distribution is due to the added mass of the joint in the segmented blade.

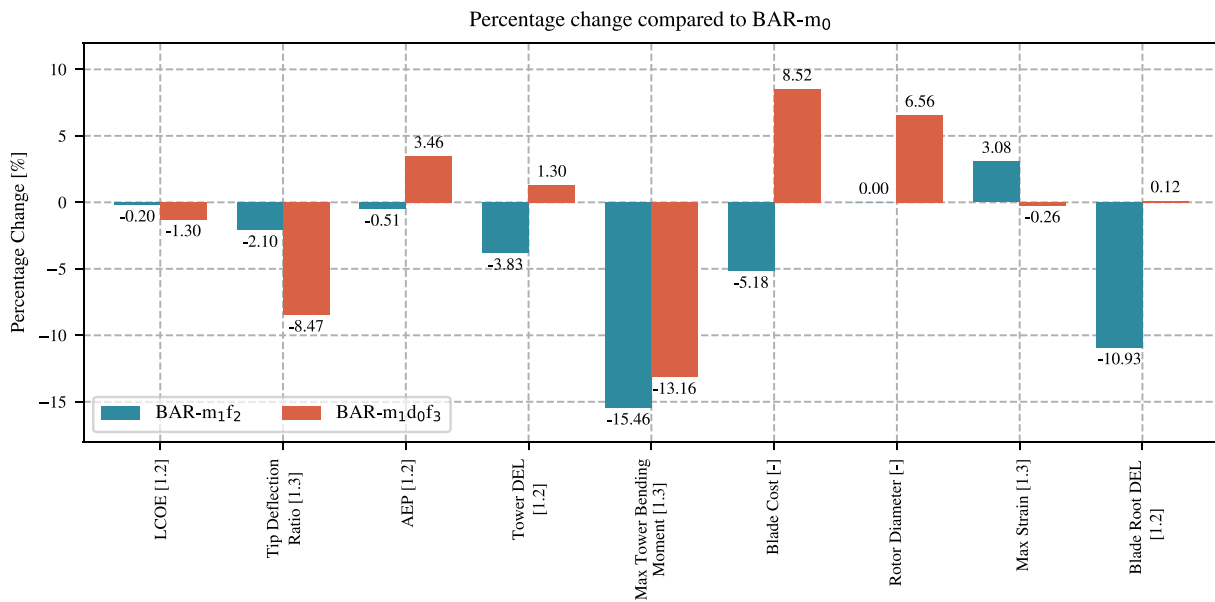


FIGURE 17 DLC results as computed from BAR- m_1f_2 and BAR- $m_1d_0f_3$. All percentages are plotted compared to BAR- m_0 other than the tip deflection ratio, which is relative to the maximum allowable value of 1. The relevant DLC to each category is denoted by [1.2] or [1.3], unless unrelated (e.g., for rotor diameter)

flaps must reduce the thrust on the blade to reduce its tip deflection, which then negatively affects the power production. This power loss cannot be completely recovered through the blade mass savings.

With a concurrent rotor diameter, spar cap thickness, and trailing edge flap optimization (BAR- $m_1d_0f_3$), LCOE is reduced further (1.3%), as shown by the red (right) bars in Figure 17. There is a clear influence on tower loading when changing the rotor size, which should be studied further; however, despite an 8.5% increase in blade cost from the increased rotor size, the associated AEP increase provides a net benefit. Note that the tower DEL constraint is seen as violated in Figure 17—this is due to different turbulence seeds used for the final loads check, and demonstrates the need for a high number of turbulent simulations to effectively capture the load envelope during an optimization process.

Though this high-level analysis is by no means a holistic evaluation of the two resulting wind turbine models, it demonstrates the potential positive and negative influences of including trailing edge flaps during a co-design process.

7 | CONCLUSIONS AND FUTURE WORK

This paper presents control co-design studies using the WEIS optimization software for distributed aerodynamic control devices. It has been shown that trailing edge flaps reduce design-driving blade tip deflections by 21.5% as compared to a rotor optimized using a sequential design approach (BAR- m_0). By conducting a control co-design optimization of the wind turbine blade with trailing edge flaps, blade mass and cost are reduced by 3.1% and 5.2%, respectively, through an optimization of spar cap thickness. The presented optimizations show that both tip deflection and material strain constraints are active in the optimized blade during the DLC 1.3 simulations. Ultimately, the 5.2% blade cost reduction results in only a 0.2% LCOE reduction for a final set of DLC simulations.

To further explore the design space and attempt to directly reduce the LCOE, the rotor diameter was included in the optimization process. The BAR-USC rotor diameter is increased by 6.5% to approximately 220 m, which results in an LCOE reduction of 1.3%, without significantly affecting the loads at the tower base. Though changing the rotor diameter has the potential to require significant modifications to the rest of the turbine, this initial study shows the potential for another way in which the co-optimization of wind turbine rotors with DAC devices can reduce the cost of wind energy.

7.1 | Possibilities of the optimization framework

This work presented a framework by which aero-servo-elastic control co-optimization of distributed aerodynamic control design can be conducted. A multitude of aerodynamic, structural, and control-oriented modeling tools have been combined into the WEIS framework. This framework uses the popular wind turbine simulation tool OpenFAST to simulate turbulent fully coupled wind turbine dynamics in varied turbulent inflow conditions. This allows for extensive, albeit computationally expensive, evaluation of the wind turbine design during each optimization iteration. Additionally, the ROSCO framework has been further developed to include controller implementation and automated tuning methods for distributed aerodynamic control devices. The controller tuning methods are developed generically and analytically such that they can be applied consistently across varied distributed aerodynamic control configurations. This development of ROSCO provides the research community with the first publicly available DAC tuning and implementation framework. Additionally, the integration of ROSCO's DAC capabilities into WEIS provides the first fully open-source tool set for the co-design optimization of wind turbines with trailing edge flaps.

7.2 | Lessons learned and future work

A number of important lessons were learned and demonstrated through this work. First and foremost, co-optimization using coupled nonlinear simulation can be computationally expensive. Using the COBYLA optimizer, each design iteration can take approximately 8–10 min using $n_{of} + 1$ computer cores, where n_{of} is the total number of OpenFAST runs at each iteration. This resulted in optimizations taking 12–24 h using up to 100 cores on NREL's HPC systems. Gradient-based optimization algorithms can converge faster than a gradient-free method such as COBYLA, but $n_{DV} \times n_{of} \times 2$ (where n_{DV} is the number of design variables) cores must be available at each design iteration for each central finite difference evaluation. Though we are fortunate enough to have access to powerful HPC systems, methods of problem simplification should certainly be explored for users without such capabilities. Further, to enable these parallelization methods, significant work was put into the multithreading capabilities for all aspects of WEIS.

By nature of using complex modeling tools within the WEIS framework, the sensitivity of large-scale problems to minor computational changes was thoroughly explored in the development of the methods presented in this work. For example, both XFOIL and OpenFAST can easily fail if they are not properly initialized. Methods for redefining the finite element mesh and constraining the trailing deflection amounts were developed to guarantee the convergence of XFOIL over the thousands of calls needed during an optimization. Similarly, the problem formulation, active

constraints, and low-order analysis within WEIS were made robust enough to provide OpenFAST with a viable turbine model, to avoid there being an OpenFAST failure 12 h into an optimization. Finally, the use of nonlinear simulation results can lead to nonsmooth optimization spaces. This not only provides additional challenges for gradient-based optimization algorithms, but even gradient-free methods can struggle to consistently converge to global minima. As shown in the convergence history for $\text{BAR-}m_0f_1$ (in Figure 13), even small design changes can result in large perturbations of the optimization objectives or constraints.

From an optimization perspective, further work can be done to shorten the necessary time for convergence and to better understand the optimization space. WEIS is constructed within the OpenMDAO framework, which facilitates running design studies, constructing surrogate models, and exploring other design parameters. Also, the inclusion of more design variables could be useful for better understanding the potential influence of distributed aerodynamic control devices. For example, a co-design optimization of a trailing edge flap, rotor size, blade spar cap thickness, tower diameter, and blade pitch controller could provide an interesting analysis into the potential influence of trailing edge flaps on tower design.

The modeling of the trailing edge flaps also deserves refinement in the future. The analysis in this work does not include trailing edge flap actuator or cost models, which would certainly influence the overall benefit of using such devices. Although this work focused on developing co-design methods and implementations themselves, future refinement of the actuator and cost models will improve the quality and feasibility of the co-design results. Similarly, optimization studies concerning the size of the trailing edge flap could more reasonably be conducted. Additionally, it is recommended that future work considers methods of generating aerodynamic polar tables other than XFOIL. XFOIL is a very useful and powerful tool; however, it can be challenging to guarantee convergence within a completely automated framework and to generate feasible polar tables for cases with larger flaps or greater flap deflection angles. Further, this work has primarily focused on the spar cap thickness during the blade optimization process, though other blade design aspects, such as outer shell and shear web layers or aerodynamic design, could allow for further reductions in blade cost or improvements in the AEP.

The trailing edge flap controller methods developed in this work lay a strong foundation for the inclusion of DAC devices within an optimization framework. Future work could improve the overall performance of this controller, particularly through more extensive investigations into the aero-servo-structural couplings that are introduced. This could simply hold the form of including scheduling parameters to activate the DAC devices only when they are needed, or it could include more advanced methods such as adaptive or predictive control. This could also include an investigation into the influences of such controllers on the aeroelastic stability of the blade.

Collectively, this paper presents a thorough and comprehensive groundwork for which future distributed aerodynamic control studies can be conducted within the ever-growing open-source software frameworks for wind energy. This work has opened the door to numerous avenues of future work to further explore the design advantages of smart rotor control.

ACKNOWLEDGMENTS

The 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. This work was authored in part 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 the US Department of Energy's Office of Energy Efficiency and Renewable Energy Wind Energy Technologies Office. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the US Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

PEER REVIEW

The peer review history for this article is available at <https://www.webofscience.com/api/gateway/wos/peer-review/10.1002/we.2840>.

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/wisdem/weis>.

ORCID

Nikhar J. Abbas  <https://orcid.org/0000-0002-1838-3463>

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

REFERENCES

1. Fleming P, Gebraad PMO, Lee S, van Wingerden J-W, Johnson K, Churchfield M, Michalakes J, Spalart P, Moriarty P. Simulation comparison of wake mitigation control strategies for a two-turbine case. *Wind Energy*. 2015;18(12):2135-2143.
2. Barlas TK, van Kuik GAM. Review of state of the art in smart rotor control research for wind turbines. *Prog Aerosp Sci*. 2010;46(1):1-27.

3. Johnson SJ, Baker JP, van Dam CP, Berg D. An overview of active load control techniques for wind turbines with an emphasis on microtabs. *Wind Energy: An Int J Prog Appl Wind Power Convers Technol*. 2010;13(2-3):239-253.
4. Yen Nakafuji DT, van Dam CP, Michel J, Morrison P. Load control for turbine blades: a non-traditional microtab approach. In: *Wind Energy Symposium*. ASME; 2002:321-330. doi:10.1115/WIND2002-54
5. Chopra I. Review of state of art of smart structures and integrated systems. *AIAA J*. 2002;40(11):2145-2187.
6. Tongchitpakdee C, Benjanirat S, Sankar LN. Numerical studies of the effects of active and passive circulation enhancement concepts on wind turbine performance. *J Solar Energy Eng*. 2006;128:4320-444.
7. Joncas S, Bergsma O, Beukers A. Power regulation and optimization of offshore wind turbines through trailing edge flap control. In: 43rd AIAA Aerospace Sciences Meeting and Exhibit. AIAA; 2005. doi:10.2514/6.2005-394
8. Lackner MA, van Kuik G. A comparison of smart rotor control approaches using trailing edge flaps and individual pitch control. *Wind Energy: An Int J Prog Appl Wind Power Convers Technol*. 2010;13(2-3):117-134.
9. van Wingerden J-W, Hulskamp AW, Barlas T, Marrant B, van Kuik GAM, Molenaar D-P, Verhaegen M. On the proof of concept of a smart wind turbine rotor blade for load alleviation. *Wind Energy: An Int J Prog Appl Wind Power Convers Technol*. 2008;11(3):265-280.
10. Bernhammer LO, van Kuik GAM, De Breuker R. Fatigue and extreme load reduction of wind turbine components using smart rotors. *J Wind Eng Ind Aerodyn*. 2016;154:84-95.
11. Bergami L, Poulsen NK. A smart rotor configuration with linear quadratic control of adaptive trailing edge flaps for active load alleviation. *Wind Energy*. 2015;18(4):625-641.
12. Bartholomay S, Krumbein S, Deichmann V, Gentsch M, Perez-Becker S, Soto-Valle R, Holst D, Nayeri CN, Paschereit CO, Oberleithner K. Repetitive model predictive control for load alleviation on a research wind turbine using trailing edge flaps. *Wind Energy*. 2022;25:1290-1308.
13. Ng BF, Hesse H, Palacios R, Graham JMR, Kerrigan EC. Aeroservoelastic state-space vortex lattice modeling and load alleviation of wind turbine blades. *Wind Energy*. 2015;18(7):1317-1331.
14. Ungurán R, Kühn M. Combined individual pitch and trailing edge flap control for structural load alleviation of wind turbines. In: *Proc. American Control Conference*. ACC; 2016:2307-2313.
15. Ungurán R, Petrović V, Pao LY, Kühn M. Smart rotor control of wind turbines under actuator limitations. In: *Proc. American Control Conference*. ACC; 2019:3474-3481.
16. Wilson DG, Berg DE, Barone MF, Berg JC, Resor BR, Lobitz DW. Active aerodynamic blade control design for load reduction on large wind turbines. In: *Proc. European Wind Energy Conference*, Marseille, France. EWEA; 2009:643-678.
17. Zhang M, Tan B, Xu J. Parameter study of sizing and placement of deformable trailing edge flap on blade fatigue load reduction. *Renew Energy*. 2015; 77:217-226.
18. Meng F, Lio WH, Barlas T. DTUWEC: An open-source DTU wind energy controller with advanced industrial features. In: *Journal of Physics: Conference Series*. IOP Publishing; 2020:022009.
19. NREL. ROSCO. Version 2.0.0. <https://github.com/NREL/rosco>; 2020.
20. Abbas NJ, Zalkind DS, Pao L, Wright A. A reference open-source controller for fixed and floating offshore wind turbines. *Wind Energy Sci*. 2022;7(1): 53-73.
21. Chen ZJ, Stol KA, Mace BR. Wind turbine blade optimisation with individual pitch and trailing edge flap control. *Renew Energy*. 2017;103:750-765.
22. Barlas T, Pettas V, Gertz D, Madsen HA. Extreme load alleviation using industrial implementation of active trailing edge flaps in a full design load basis. In: *Journal of Physics: Conference Series*. IOP Publishing; 2016:042001.
23. McWilliam MK, Barlas TK, Madsen HA, Zahle F. Aero-elastic wind turbine design with active flaps for AEP maximization. *Wind Energy Sci*. 2018;3(1): 231-241.
24. Garcia-Sanz M. Control co-design: an engineering game changer. *Advanced Control for Applications: Engineering and Industrial Systems*; 2019.
25. Bortolotti P, Bottasso CL, Croce A. Combined preliminary-detailed design of wind turbines. *Wind Energy Sci*. 2016;1(1):71.
26. Herber DR. Advances in combined architecture, plant, and control design. *Ph.D. Thesis*: Champaign, IL, USA; 2017.
27. Deshmukh AP, Allison JT. Multidisciplinary dynamic optimization of horizontal axis wind turbine design. *Struct Multidiscip Optim*. 2016;53(1):15-27.
28. Herber DR, Allison JT. Nested and simultaneous solution strategies for general combined plant and control design problems. *J Mech Design*. 2019; 141(1):011402.
29. Pao LY, Zalkind DS, Griffith DT, Chetan M, Selig MS, Ananda GK, Bay CJ, Stehly T, Loth E. Control co-design of 13 MW downwind two-bladed rotors to achieve 25% reduction in levelized cost of wind energy. *Ann Rev Control*. 2021;51:331-343.
30. NREL. WEIS. Version 0.0.1. GitHub repository; 2021.
31. Zahle F, Tibaldi C, Verelst DR, Bitsche R, Bak C. Aero-elastic optimization of a 10 MW wind turbine. In: *Proc. Wind Energy Symposium*. AIAA; 2015: 0491.
32. Bottasso CL, Croce A. Cp-Lambda: User's Manual. Dipartimento di Ingegneria Aerospaziale, Politecnico di Milano, Italy; 2009.
33. Drela M. Xfoil: An analysis and design system for low Reynolds number airfoils. In: *Low Reynolds Number Aerodynamics*. Springer; 1989:1-12.
34. NREL. OpenFAST. Version 3.0.0. GitHub repository, GitHub, <https://github.com/openfast/openfast>; 2021.
35. Damiani RR, Hayman G. The unsteady aerodynamics module for fast8; 2019.
36. NREL. Wisdem. GitHub, <https://github.com/wisdem/wisdem>; 2022.
37. Ning SA. Ccblade documentation: National Renewable Energy Laboratory; 2013.
38. Gavin HP. Frame3DD. Frame3DD Structural Analysis Code, Duke University; 2010.
39. NREL. Pcrunch. Version 0.0.1. GitHub repository, <https://github.com/NREL/pcrunch>; 2021.
40. Buhl ML. *MCrunch User's Guide for Version 1.00*: National Renewable Energy Laboratory Golden, CO; 2008.
41. Gray JS, Hwang JT, Martins JRR, Moore KT, Naylor BA. OpenMDAO: an open-source framework for multidisciplinary design, analysis, and optimization. *Struct Multidiscip Optim*. 2019;59(4):1075-1104.
42. Bortolotti P, Johnson N, Abbas NJ, Anderson E, Camarena E, Paquette J. Land-based wind turbines with flexible rail-transportable blades—part 1: conceptual design and aeroservoelastic performance. *Wind Energy Sci*. 2021;6(5):1277-1290.
43. Camarena E, Anderson E, Paquette J, Bortolotti P, Feil R, Johnson N. Land-based wind turbines with flexible rail-transportable blades—part 2: 3D finite element design optimization of the rotor blades. *Wind Energy Sci*. 2022;7(1):19-35.

44. Zalkind DS, Nicotra MM, Pao LY. Constrained power reference control for wind turbines. *Wind Energy*. 2022;25(5):914-934.
45. Gaertner E, Rinker J, Sethuraman L, Zahle F, Anderson B, Barter GE, Abbas NJ, Meng F, Bortolotti P, Skrzypinski W, Scott GN, Feil R, Bredmose H, Dykes K, Shields M, Allen C, Viselli A. IEA wind TCP task 37: definition of the IEA 15-megawatt offshore reference wind turbine. National Renewable Energy Laboratory, Golden, CO (United States); 2020.
46. Burton T, Jenkins N, Sharpe D, Bossanyi E. *Wind Energy Handbook*: John Wiley & Sons, Ltd; 2011.
47. Bossanyi EA. Wind turbine control for load reduction. *Wind Energy*. 2003;6(3):229-244.
48. Mulders SP, Pamososuryo AK, Disario GE, van Wingerden J-W. Analysis and optimal individual pitch control decoupling by inclusion of an azimuth offset in the multiblade coordinate transformation. *Wind Energy*. 2019;22(3):341-359.
49. Bir G. Multi-blade coordinate transformation and its application to wind turbine analysis. In: 46th AIAA aerospace sciences meeting and exhibit. AIAA; 2008:1300.
50. Hansen M. *Aerodynamics of Wind Turbines*: Routledge; 2015.
51. Helton JC, Davis FJ. Latin hypercube sampling and the propagation of uncertainty in analyses of complex systems. *Reliabil Eng Syst Saf*. 2003;81(1):23-69.
52. IEC. Wind Turbines-Part 1: Design Requirements. IEC 614001 Ed 3; 2006.
53. Wang Q, Sprague MA, Jonkman J, Johnson N, Jonkman B. BeamDyn: a high-fidelity wind turbine blade solver in the fast modular framework. *Wind Energy*. 2017;20(8):1439-1462.
54. Berg JC, Resor BR. Numerical Manufacturing and Design Tool (NuMAD v2. 0) for Wind Turbine Blades: User's Guide., Albuquerque, NM, and Livermore, CA, Sandia National Laboratories (SNL); 2012.
55. Johnson SG. The NLOpt Nonlinear-Optimization Package. Github repository; 2021.
56. Svanberg K. A class of globally convergent optimization methods based on conservative convex separable approximations. *SIAM J Optim*. 2002;12(2):555-573.
57. Powell MJD. A direct search optimization method that models the objective and constraint functions by linear interpolation. *Advances in Optimization and Numerical Analysis*, Springer; 1994:51-67.
58. Kraft D. A software package for sequential quadratic programming; 1988.
59. Gill PE, Murray W, Saunders MA. SNOPT: an SQP algorithm for large-scale constrained optimization. *SIAM Rev*. 2005;47(1):99-131.

How to cite this article: Abbas NJ, Bortolotti P, Kelley C, Paquette J, Pao L, Johnson N. Aero-servo-elastic co-optimization of large wind turbine blades with distributed aerodynamic control devices. *Wind Energy*. 2023;26(8):763-785. doi:10.1002/we.2840

APPENDIX A: FILTERS

A combination of standard filters is used to filter the blade root bending moment signal. In the following filters, ω_f is a cornering frequency in rad/s, and ζ_f is a unitless damping ratio.

Second-order low-pass filter:

$$F_{L2}(s) = \frac{\omega_f^2}{s^2 + 2\zeta_f\omega_f s + \omega_f^2}. \tag{A1}$$

First-order high-pass filter:

$$F_{H1} = \frac{s}{s + \omega_f}. \tag{A2}$$

Notch filter:

$$F_N = \frac{s^2 + 2\omega_f\zeta_{f1}s + \omega_f^2}{s^2 + 2\omega_f\zeta_{f2}s + \omega_f^2}. \tag{A3}$$

Band-pass filter:

$$G = \frac{\omega_f s}{\zeta_f s^2 + \omega_f s + \zeta_f \omega_f^2}. \tag{A4}$$