Control Co-Design Studies for a 22 MW Semisubmersible Floating Wind Turbine Platform

Journal of Physics: Conference Series **2767** (2024) 082020 doi:10.1088/1742-6596/2767/8/082020

Control Co-Design Studies for a 22 MW Semisubmersible Floating Wind Turbine Platform

Daniel Zalkind, Pietro Bortolotti

National Renewable Energy Laboratory, Golden, CO 80401, USA

E-mail: daniel.zalkind@nrel.gov

Abstract. We present a control co-design software framework that can be used to optimize floating wind turbines and their controllers. Because this framework has many options for design variables, constraints, and merit figures, along with modeling fidelity levels, we seek to demonstrate best practices for using the tool while designing a floating platform for the new 22 MW offshore reference wind turbine developed within the International Energy Agency Wind Technology Commercialization Programme 55 on Reference Wind Turbines and Farms. During these studies, we evaluate the use of different simulation fidelity levels, the effect of using different load cases for controller tuning, and the difference between sequential and simultaneous control co-design solutions. Based on these efforts, we suggest using an algorithm that performs an initial search of the design space before optimization. We find that solving smaller optimization problems, in a sequential manner, leads to more reliable outcomes in fewer iterations than larger, simultaneous control co-design solutions. However a simultaneous CCD solution produces a platform with a 2% lower mass than the sequential CCD outcome.

1 Introduction to Control Co-design Design for Floating Wind Turbines

We use an optimization framework to design a semisubmersible platform and controller for a 22 MW floating offshore wind turbine (FOWT) developed within the International Energy Agency (IEA) Wind Technology Commercialization Programme (TCP) 55 on Reference Wind Turbines and Farms. The outcome of this design process can be used as a reference for future designs and research processes while highlighting the challenges facing the next generation of FOWTs. While performing the optimizations, we evaluate how choices about the design process affect the outcome and then recommend our current best practices for the control co-design (CCD) of wind turbine systems.

In CCD, the controller and the system are designed together. When using a sequential CCD problem formulation, the system is designed first, followed by the controller, and this process can be iterative. In simultaneous CCD, the system and the controller are optimized together in the same problem [1]. Many CCD problems are posed as an optimization problem, where some merit figure is to be minimized, subject to constraints on the design like load limits or maximum displacements. Merit figures and constraints are typically the output of some analysis performed on the system. Turbine and controller design variables are input parameters to be optimized, ideally impacting the merit figure and constraints. An optimization-based design process is desirable because it can use a solver to automatically balance competing objectives between the merit figure and constraints, and inherently account for the coupling between the design variables and these outputs.

CCD of wind energy systems requires the coupling of multidisciplinary modeling tools [2]. Several design-oriented FOWT models have been developed in recent years, with varying levels of fidelity [3; 4]. Alternative modeling solutions have been developed to directly solve for the optimal controller and reduce computational costs [5]. The solutions presented in this article run simulations with a fidelity level

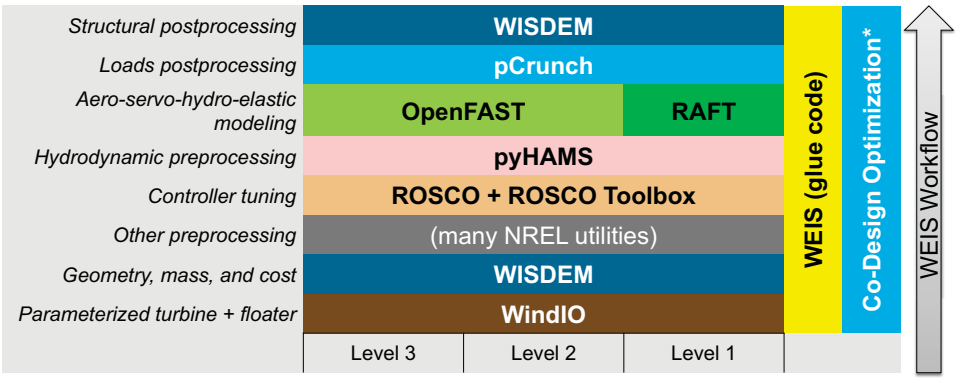

Figure 1: The WEIS software stack, where turbines are first defined by the WindIO parameterization. The WISDEM system engineering model and other preprocessing tools convert the model into an aeroelastic definition for simulations in RAFT and OpenFAST. pyHAMS is a boundary element solver for potential flow modeling. A wind turbine controller can be tuned using ROSCO, and simulations are postprocessed using pCrunch. WISDEM can use load simulations to compute stresses and strains on various components.

adequate for wind turbine certification, along with a realistic controller, using the framework originally presented in [6]. We expand the previous work by evaluating more features across more fidelity levels. We also investigate the design methodologies available in the tool by comparing sequential CCD to simultaneous CCD.

In Section 2, we provide an overview of the Wind Energy with Integrated Servo Controls (WEIS) toolset. In Section 3, we present a baseline CCD problem that we solve using WEIS. We evaluate the outcome from several configurations of the CCD problem in Section 4. Based on the results, we summarize our findings in Section 5 and propose ideas for future work.

2 Using the Wind Energy and Integrated Servo Controls Toolset for Floating Platform Control Co-design

To perform CCD studies, we use WEIS, which links the tools shown in Fig. 1 that have been developed by the National Renewable Energy Laboratory (NREL) and others to model floating wind turbine systems, controllers, and cost estimates.

2.1 WISDEM and WindIO

The Wind Plant Integrated System Design and Engineering Model (WISDEM® [7]) is a systems engineering framework devoted to the design optimization of land-based and offshore wind turbines. WISDEM is composed of modules that are integrated into a single workflow but that can also be run separately. This project used WISDEM to convert a WindIO [8] parameterized geometry input of the turbine into a floating platform with mass and cost estimates for each component. WISDEM enforces several geometric and hydrostatic constraints, like a constraint that ensures that the bottom of the platform does not leave the water (or the top of the platform does not enter the water) during an extreme heel offset of 10°. WISDEM also provides an automated ballast calculation: based on the system mass, displacement, and mooring forces, it computes the mass and volume of water that will balance the system to a zero mean heave. WISDEM can also be used to compute hydrostatic parameters, like the metacentric height and the natural periods of each floating degree of freedom.

2.2 Aeroelastic FOWT Models: RAFT and OpenFAST

We evaluate the use of two different FOWT modeling tools. A lower-fidelity model, RAFT [9], provides a linear, frequency domain representation of a FOWT. RAFT includes a model of the rotor aerodynamics and the effect of the controller on rotor motion. RAFT is used to compute mean offsets, response amplitude operators, and the power spectral densities resulting from disturbances that can be modeled in the frequency domain, including wind turbulence and wave spectra. These power spectral densities are then converted into motion standard deviations. Maximum and minimum displacements are estimated

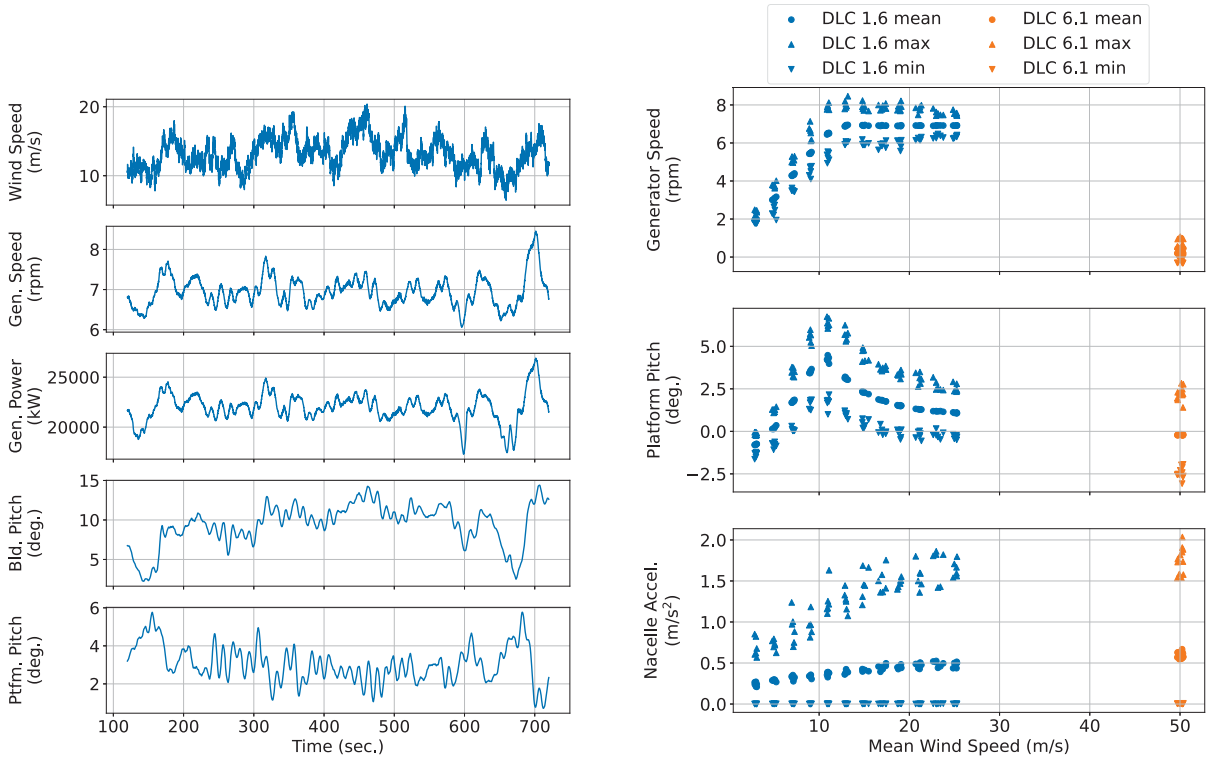

(a) Example OpenFAST time series, depicting the case in Fig. 2b with the greatest platform pitch.

(b) Example statistics from the first iteration of the OpenFAST-based optimization in Fig. 4.

Figure 2: Simulation results generated in a single OpenFAST-based optimization iteration.

as three standard deviations from the mean. Hydrostatic properties and natural periods can also be computed in RAFT. In this study, RAFT models 20 frequencies from 0.159 to 0.3138 Hz. Because the RAFT model is formulated strictly in the frequency domain and transfers the wind and wave input spectra to the six floating degrees of freedom, it can run a single load case in a few seconds, resulting in rapid design evaluations and a good starting point for higher-fidelity models.

WEIS also runs the multiphysics framework OpenFAST, a nonlinear aero-servo-hydro-elastic solver [10] that runs time domain design load case simulations. WEIS generates OpenFAST inputs to represent the blade, nacelle, tower, and platform models from WISDEM geometry inputs. In this platform design study, platform hydrodynamics is simulated using only Morison elements with standard coefficients for pressure, added mass, and drag coefficients. In future studies, these coefficients could be tuned based on field experiments or high-fidelity fluid dynamic simulations. Hybrid modeling using both potential flow solutions and Morison elements is possible in OpenFAST, but requires more sophisticated automated meshing of the platform geometry than is currently available in the WEIS toolset. The final, optimized platform has been analyzed using both hydrodynamic modeling methods with good agreement. Thus, we believe that using Morison elements is sufficient for an early-stage design optimization like the one we present. Understanding how the optimal platform design changes with fidelity levels is an avenue for future work. We show an example OpenFAST simulation in Fig. 2a and the statistics from a set of load cases in Fig. 2b.

2.3 NREL's ROSCO

The OpenFAST model of the FOWT interfaces with the nonlinear Reference OpenSource Controller (ROSCO) [11], which was used in this work and several of its tuning parameters were optimized in this study. ROSCO is designed to mimic the operational functions of standard industry wind turbine controllers. ROSCO also includes features helpful for controlling floating wind turbines. The ROSCO control architecture is fixed, but inputs, including gains and set points, can be easily adapted to different turbine models and design goals. With a given wind turbine model (e.g., from WEIS) and a few tuning

parameters, the detailed controller input can be derived. These tuning parameters can also be used as design variables in WEIS to automatically tune ROSCO given a prescribed list of objectives and constraints.

In below-rated operation, ROSCO's torque controller tracks the rotor speed for the optimal tip speed ratio, where an extended Kalman filter, standard in the ROSCO controller, estimates the wind speed using generator speed, blade pitch, and generator torque measurements. The torque controller was initially tuned for stability to ensure that the generator speed and torque vary on similar time scales. Generally, if the torque controller is stable, its effect on platform motion is expected to be minimal.

On the other hand, the pitch controller affects the platform motion significantly. Rotor thrust affects platform motion and is very sensitive to blade pitch. If not properly controlled, the blade pitch/platform motion coupling can lead to instabilities between the pitch controller, generator speed, and platform motion [12] A floating feedback loop uses information about the rotor motion, determined using an inertial measurement unit in the nacelle, to stabilize the platform motion. Recent research has shown that this feedback can be used to decouple the generator speed dynamics from the platform motion and also increase the damping on the platform motion [13; 14]. In WEIS, a data-driven approach is used to account for the coupling, where the simulation outputs determine which parameters are best for reducing platform motion and generator speed transients [15]. Another control parameter in ROSCO worth noting is the peak shaving parameter, which limits maximum thrust on the turbine rotor, affecting platform motion and transients. However, reducing thrust also reduces power capture, so more sophisticated measures, like levelized cost of energy, must be used to identify the optimal trade-off.

2.4 OpenMDAO and COBYLA

The various software programs in WEIS are linked using the Python library OpenMDAO [16], which organizes each tool into components and groups, defines their interfaces, and determines the order in which each should be executed. The inputs and outputs of OpenMDAO components can be used to assign design variables, merit figures, and constraints. OpenMDAO also interfaces with an optimization driver, which determines how the design variables are updated between iterations.

For all the optimizations presented in this article, we use the Constrained Optimization by Linear Approximation (COBYLA) [17] solver, implemented by NLOpt [18], to drive the design variables toward their optimal values. COBYLA is an iterative, derivative-free method that estimates and updates linear approximations of the merit figure and constraints over a shrinking trust region. We favor COBYLA over other optimization solvers because of its gradient-free algorithm. Exact gradients cannot be calculated in many of the WEIS components. When estimated, gradients are often noisy because stochastic simulations with turbulent inflow are used to determine constraints or merit figures. COBYLA reliably estimates these noisy constraints, even when instabilities occur, while taking reasonable step sizes based on the design variable's bounds. In our experience, we found that other solvers outperform COBYLA in specific optimization setups, but COBYLA allows more consistent convergence across optimization problems.

3 A Baseline Floating Wind Turbine Control Co-Design Optimization Problem

In the following optimizations, we seek to design a floating platform for the IEA 22 MW reference wind turbine, whose macro properties are listed in Table 1. The description of rotor, tower, and initial platform design can be found online [19]. Neither the rotor nor tower were changed from the fixed-bottom version of this turbine, which results in a soft stiff tower that might need additional controller features in the future to avoid the rotor speed interacting with the natural frequencies of the first tower fore-aft and side-to-side modes.

We use metocean conditions based on the Gulf of Maine [20] to determine the sea state in normal and extreme environments. We use design load cases (DLCs) 1.6 and 6.1 based on the IEC standards [21] for a Class IB turbine. In DLC 1.6, normal turbulence wind and severe sea states, we use six turbulence seeds for each mean wind speed from cut-in (3 m/s) to cut-out (25 m/s) in 2 m/s increments. The rated wind speed of this rotor is 11.5 m/s. For DLC 6.1, we use six turbulence seeds for each yaw offset (8° in each direction), which results in a total of 84 simulations for each design iteration. Later in this article, we will investigate the use of DLC 1.1 (normal sea state and turbulence) instead of DLC 1.6 to evaluate the effect of load case on optimal controller parameters. We found that six seeds results in a good trade-off between computational cost and convergence in both the load statistics and optimizations.

We optimize the semisubmersible platform geometry with three outer columns and a central main column, like the one shown in Fig. 3. The lower pontoon diameter is fixed at 10 m and the wall thickness is fixed at 5 cm, to avoid performing structural analysis at this design stage.

Table 1: General system properties of the IEA 22 MW reference wind turbine.

Property	Unit	Value
Turbine Rating	kW	22,000.0
Hub Height	m	170.0
Rated Rotor Speed	rpm	7.06
Platform Type	-	Semisubmersible
Tower Mass	t	1,574
RNA Mass	t	1,216
Water Depth	m	200.0
Mooring System	-	Three-line chain catenary

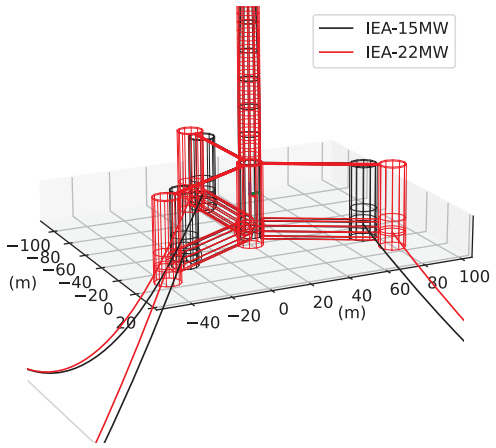

Figure 3: Optimized IEA 22 MW platform design, compared with the IEA 15 MW reference platform.

Turbine Rating (MW)	15	22
Draft (m)	20	24.2
Column Spacing (m)	51.75	64.9
Column Diameter (m)	12.0	12.4
Hull Mass (t)	4,014	4,769
Total Mass (t)	14,993	20,802

Table 2: Design variable and output comparison for an optimized 22 MW floating platform, compared with the 15 MW reference platform. The total mass of the platform includes the ballast mass. Note that the merit figure of the optimizations is the total structural mass, which includes the turbine.

3.1 Design Variables

The design variables of the optimization problem are summarized in Table 3, along with the constraints on the platform and controller. We optimize the column spacing, measured from the center of the middle column to the center of the outer columns. The outer column diameter lower bound is limited by the lower pontoon diameter. The ROSCO pitch controller is optimized across the above-rated wind speeds to provide a different response based on the natural frequency (ω_{PC}) and damping (ζ_{PC}) at three different wind speeds: 12, 18, and 24 m/s. The floating feedback controller gain (k_{float}) and phase (ω_{float}), are also optimized. All of these parameters affect platform motion and generator speed regulation. We use a peak shaving percent of 80%, matching the value used for rotor and tower design. The percentage is relative to the maximum value of aerodynamic thrust experienced by the rotor during power extraction.

3.2 Constraints

For the platform optimization, we seek a maximum pitch of 6° or less. The other dynamic constraint is the maximum nacelle acceleration. Both constraints can be evaluated in both RAFT and in OpenFAST across all the considered DLCs. An example of these measures for a single FOWT design iteration is shown in Fig. 2a and 2b. We bound the lower limit of the pitch and heave periods to avoid wave periods.

The controller affects the generator speed dynamics, which we constrain so that no simulations exceed 20% of the rated generator speed; values greater than this could trigger a shutdown procedure and reduce power capture. We can also enforce a pitch travel constraint, which measures the average distance the pitch actuator moves per second, over all simulations.

Table 3: Design variables and constraints for the baseline platform and controller design optimizations.

Design Variables	Lower Bound	Upper Bound
Column spacing	35 m	75 m
Draft	20 m	50 m
Outer column diameter	10 m	16 m
Pitch control natural frequency (ω_{PC})	0.025 rad/s	0.5 rad/s
Pitch control damping ratio (ζPC)	0.1	3.0
Floating feedback gain (k_{float})	-40 s	0 s
Floating feedback filter bandwidth (ω_{float})		1 rad/s
Constraints	Lower Bound	Upper Bound
Pitch period	20 s	
Heave period	16 s	
Survival heel		10 deg
Maximum platform pitch		6 deg
Maximum nacelle acceleration		2 m/s^2
Maximum generator speed		20% above rated
Average pitch travel		0.085 deg/s

3.3 Merit Figure

For the controller optimizations, we use the tower fore-aft moment damage equivalent load (DEL) as the merit figure to be minimized. The tower DEL provides a good measure of platform pitch motion as well as blade pitch actuation. While DLC 1.6 is not used for computing DELs in the analysis prescribed by the standards, we can still use it as an optimization objective.

For the platform design and simultaneous CCD studies, we use the system structural mass as the objective, which includes the platform hull mass and tower mass. The tower mass is unchanged in this study. In this design, only water ballast is used, which is considered to have negligible cost. In the design studies that follow, we check the effect of platform design changes on the annual energy production, but have found the differences to be at most 0.3%.

The sequential CCD method first performs a platform optimization, then a controller optimization. They do not share a merit figure because the controller parameters will not affect platform mass. In theory, the controller optimization could have used the platform mass as an objective function. In that case, the control parameters would be optimized to satisfy the shared design constraints. However, we believe that providing a merit figure like the tower DEL also serves a useful purpose during design.

4 Analysis of the Control Co-Design Solutions

Using the previously described framework and optimization problem, we evaluate the solutions of a few configurations to inform the setup of future CCD problems. We look at the optimization outcome of using different modeling fidelity levels, different load cases for optimization, and different CCD problem formulations: sequential versus simultaneous. The three studies were run on a high-performance computing cluster. One iteration of an OpenFAST run, which performs DLC simulations in parallel, takes around 20 minutes on NREL's high-performance computing cluster, compared to RAFT taking less than 30 seconds for the same analysis.

4.1 Platform Design Using RAFT Versus OpenFAST

First, we study how using different fidelity levels of WEIS affect the design outcome. We optimized a platform using the problem summarized in Table 3 without controller design variables and constraints. In one optimization we used RAFT for the dynamic constraints, which we refer to as RAFT-optimized. In the next, we start with the same initial design and use OpenFAST (OpenFAST-optimized). In the last, we use OpenFAST, but start from the outcome of the RAFT-optimized design. The OpenFAST optimization took around 11 hours on one node (with 100 cores), whereas the RAFT optimization took less than 1 hour, despite three times as many iterations.

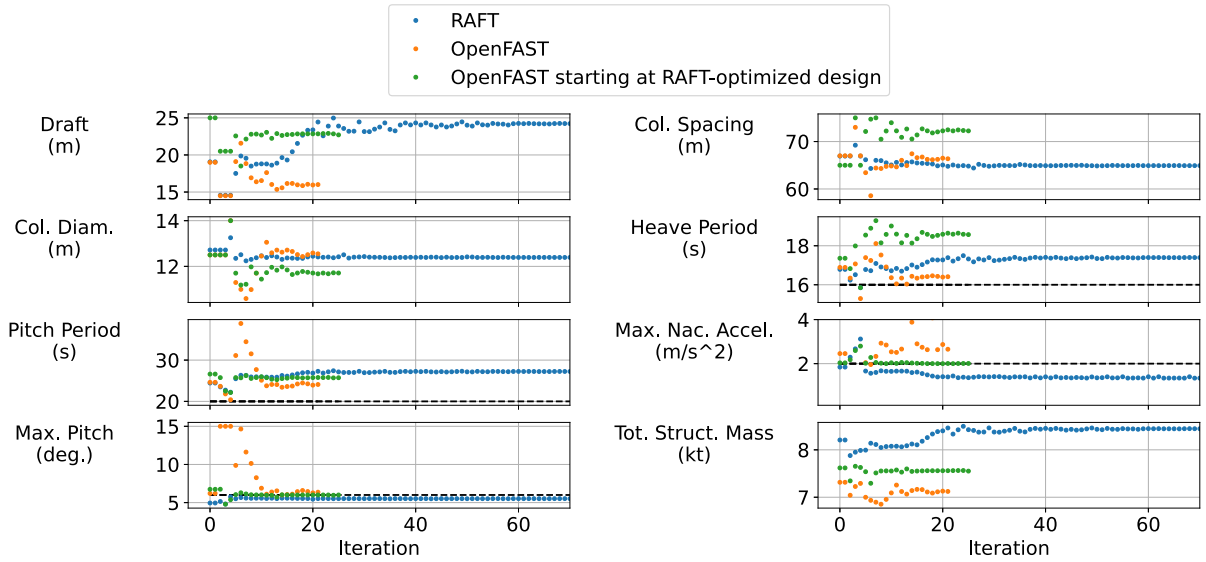

Figure 4: Platform optimization convergence trends of a RAFT-based optimization, an OpenFAST-based optimization, and an OpenFAST-based optimization starting from the RAFT-optimized design.

The convergence trends are shown in Fig. 4. The RAFT-optimized design satisfies all the constraints according to RAFT, but when used as a starting point for OpenFAST, it violates the maximum pitch constraint, even though we use a more conservative (5.5° versus 6°) constraint on the maximum platform pitch. The OpenFAST solution that starts at the same point as the RAFT optimization does not satisfy the nacelle acceleration constraint. However, the OpenFAST solution starting from the RAFT optimization does reach a feasible but different solution with a reduced draft and mass but greater column spacing in order to satisfy the platform pitch constraint.

The lesson learned in this study is that it is necessary to start with a feasible solution. One way to accomplish that is by using a multiple start algorithm or by performing an initial search of the design space. Multi-fidelity optimization procedures [22] could also be useful to formally couple two or more levels of fidelity that do not necessarily agree, whether constraints are violated or not.

4.2 Controller Design in Different DLCs

The next step in a sequential CCD process is to reoptimize the controller. Using the OpenFAST-optimized platform of the previous section, we optimize the controller design variables based on the constraints and merit figures discussed in Section 3. We typically design a controller for normal operation (DLC 1.1) and then detect and adapt the controller for severe sea states, which occur less frequently. However, the previous platform optimization used DLC 1.6 simulations. In this study, we evaluate how the optimal controller design depends on the load cases used for analysis. In previous work, we have seen that floating feedback can increase loading in some extreme wave environments [23].

To evaluate this question, we set up optimizations using both DLC 1.1 and 1.6. The convergence trends of each are shown in Fig 5. Without adding an additional constraint on the blade pitch travel, the DLC 1.6 optimization does not converge to a feasible solution. The biggest difference between the optimized controllers is in the floating feedback gain k_{float}, which is reduced in magnitude from the DLC 1.1 optimization. Overall, however, both optimizations converge to a similar set of ROSCO parameters. Thus, we believe it is sufficient to use DLC 1.6 simulations for simultaneous CCD optimizations, including the controller design variables.

4.3 Sequential and Simultaneous Control Co-Design

In the final design study, we evaluate the outcome of optimizing both the controller and platform together in a simultaneous CCD solution. The platform and controller design variables and constraints of Table 3 are combined in the simultaneous CCD optimization. We use the total structural mass as the merit figure. The optimization convergence, shown in Fig. 6, is compared to the OpenFAST-optimized platform-only optimization from Fig. 4.

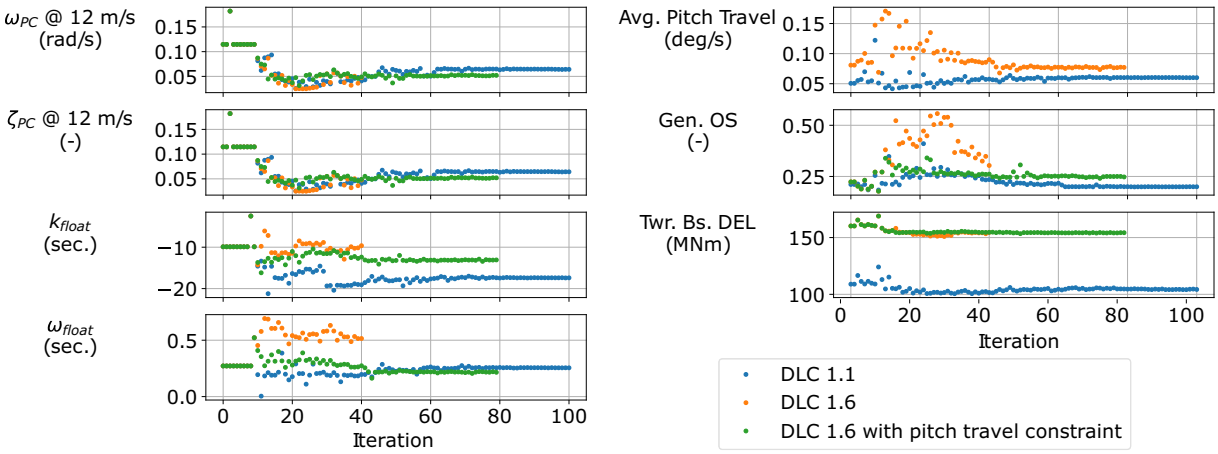

Figure 5: ROSCO controller optimization, based on DLC 1.1 or 1.6 simulations in OpenFAST. In a third optimization, a blade pitch travel constraint has been added to guide the DLC 1.6 solution to a feasible solution. The DLC 1.6 optimization was also allowed a greater generator overspeed (Gen. OS) value.

Neither optimal design, the final iterations in the convergence of Fig. 6, differ significantly from the initial point, the RAFT-optimized design. However, the simultaneous CCD-optimized design reduces draft further, possibly because the controller can also contribute to reducing platform motion. It is interesting that the generator overspeed is not an active constraint, while it was for the controller optimization. Overall, the optimized controller does not differ much from the initial controller used in the platform-only optimization. The simultaneous CCD takes about four times longer to converge to an optimal solution than the sequential method, but achieves a 2% lower platform mass. In general, the controller design variables take longer to converge, possibly because their effect on the constraints and merit figure is more stochastic compared to the platform design properties. While preparing this article, we more reliably and quickly reached a solution when using a sequential CCD framework, where the platform and controller are separately optimized. The platform optimized using the sequential method is shown in Figure 2.

5 Conclusions and Future Work

In this article, we have demonstrated the WEIS tool's ability to optimize floating platforms, controllers, and both simultaneously. A baseline CCD problem was proposed and used to analyze several choices when setting up optimization problems. We found that using lower-fidelity models can be very efficient, but if the outputs are not in agreement with higher-fidelity tools, they can only provide good starting points for future studies. We found that the optimal controller parameters do not depend significantly

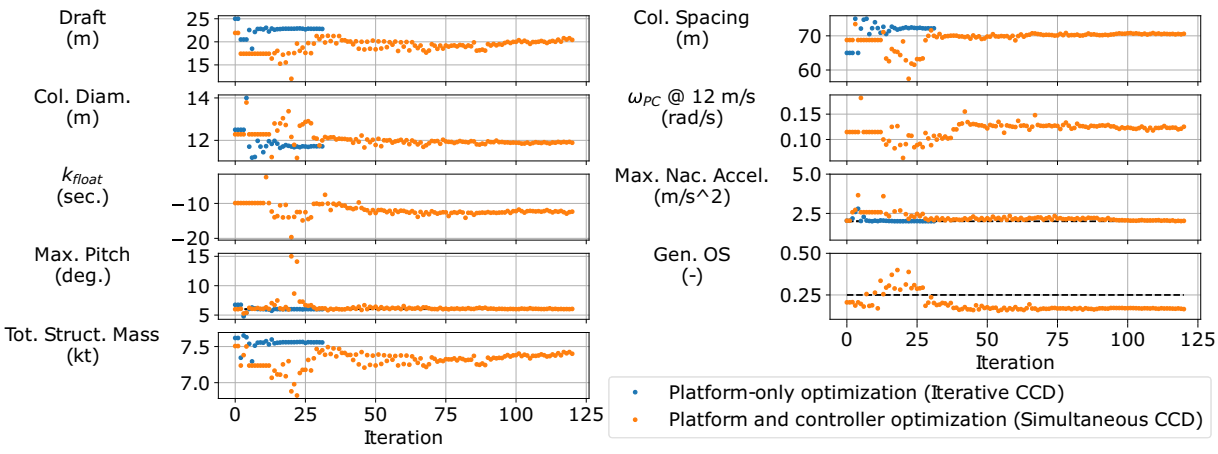

Figure 6: Platform-only optimization compared with a simultaneous CCD solution.

on the sea state and design load case for the models used in this study. We found that the sequential CCD method, where the platform and controller are designed separately, leads to more reliable outcomes than a simultaneous method. Simultaneous CCD can provide modest performance improvements with a larger computational cost.

More research is planned to evaluate different optimization solvers that can handle a complex model with computed, noisy gradients. One potential method to smooth noisy gradients is to train surrogate models of the design space. Using absolute maxima as constraints can be replaced with smooth max functions or by using the frequency of exceedance. More load cases should be analyzed to bring the systems designed using this tool closer to certification. We are currently seeking feedback from industry to provide more realistic constraints on floating designs and understand future modeling requirements. For example, the maximum mean platform pitch is commonly used in platform design rather than the maximum overall platform pitch. Eventually, design optimizations should account for the structural design of the floating platform and its distributed loads. We hope this reference model will be used in future studies of extreme-scale FOWTs, and this study can inform design methods of future wind energy systems, leading to improved cost, reliability, and performance.

Acknowledgements
This research was performed using computational resources sponsored by the Department of Energy's Office of Energy Efficiency and Renewable Energy and located at the National Renewable Energy Laboratory. This work was authored by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by the U.S. Department of Energy Office of Energy Efficiency and Renewable Energy, Wind Energy Technologies Office. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

References

[1] Allison JT, Herber DR. Multidisciplinary Design Optimization of Dynamic Engineering Systems. AIAA Journal. 2014;52(4):691-710.

[2] Pao LY, Pusch M, Zalkind DS. Control Co-Design of Wind Turbines. Annual Review of Control, Robotics, and Autonomous Systems. 2024;7(1).

[3] Lemmer F, Yu W, Luhmann B, Schlipf D, Cheng PW. Multibody modeling for concept-level floating offshore wind turbine design. Multibody System Dynamics. 2020 06;49.

[4] Leimeister M, Collu M, Kolios A. A fully integrated optimization framework for designing a complex geometry offshore wind turbine spar-type floating support structure. Wind Energy Science. 2022;7(1):259-81.

[5] Sundarrajan AK, Hoon Lee Y, Allison JT, Zalkind DS, Herber DR. Open-Loop Control Co-Design of Semisubmersible Floating Offshore Wind Turbines Using Linear Parameter-Varying Models. Journal of Mechanical Design. 2023 11;146(4):041704.

[6] Abbas NJ, Jasa J, Zalkind DS, Wright A, Pao L. Control co-design of a floating offshore wind turbine. Applied Energy. 2024;353:122036.

[7] NREL. WISDEM. GitHub; 2023. Accessed Jan 24, 2024. `https://github.com/WISDEM/WISDEM`.

[8] IEA Wind Task 37 Team. WindIO; 2020. Accessed on Jan 24, 2024. `https://windio.readthedocs.io/en/latest/index.html`.

[9] Hall M, Housner S, Ogden D, Zalkind D, Barter G, Bortolotti P. RAFT (Response Amplitudes of Floating Turbines). US National Renewable Energy Laboratory; 2022.

[10] NREL. OpenFAST;. Accessed on Jan 24, 2024. `https://github.com/OpenFAST/openfast`.

[11] Abbas N, Zalkind D, Pao LY, Wright A. A reference open-source controller for fixed and floating offshore wind turbines. Wind Energy Science. 2022;7:53-73.

[12] van der Veen GJ, Couchman IJ, Bowyer RO. Control of floating wind turbines. In: Proc. American Control Conf.; 2012. p. 3148-53.

[13] Stockhouse D, Pao L. Trade-offs in the design of multi-loop controllers for floating wind turbines. In: Proc. American Control Conf.; 2023. p. 2530-5.

[14] Capaldo M, Mella P. Damping analysis of floating offshore wind turbines (FOWTs): a new control strategy reducing the platform vibrations. Wind Energy Science. 2023;8(8):1319-39.

[15] Zalkind D, Abbas NJ, Jasa J, Wright A, Fleming P; IOP Publishing. Floating wind turbine control optimization. J Phys: Conf Ser. 2022;2265(4).

[16] Gray JS, Hwang JT, Martins JRRA, Moore KT, Naylor BA. OpenMDAO: An open-source framework for multidisciplinary design, analysis, and optimization. Structural and Multidisciplinary Optimization. 2019;59:1075-104.

[17] Powell MJ. A direct search optimization method that models the objective and constraint functions by linear interpolation. In: Advances in optimization and numerical analysis. Springer; 1994. p. 51-67.

[18] Johnson SG. The NLopt nonlinear-optimization package; 2007. `https://github.com/stevengj/nlopt`.

[19] IEA Task 55 REFWIND. IEA-22-280-RWT;. Accessed on Jan 24, 2024. `https://github.com/IEAWindTask37/IEA-22-280-RWT`.

[20] Viselli AM, Forristall GZ, Pearce BR, Dagher HJ. Estimation of extreme wave and wind design parameters for offshore wind turbines in the Gulf of Maine using a POT method. Ocean Engineering. 2015;104:649-58.

[21] International Electrotechnical Commission. Wind turbines - Part 1: Design requirements; 2005. 61400-1.

[22] Jasa J, Bortolotti P, Zalkind D, Barter G. Effectively using multifidelity optimization for wind turbine design. Wind Energy Science. 2022;7(3):991-1006.

[23] Wang L, Bergua R, Robertson A, Wright A, Zalkind D, Fowler M, et al. Experimental investigation of advanced turbine control strategies and load-mitigation measures with a model-scale floating offshore wind turbine system. Applied Energy. 2024;355:122343.