Design of fatigue-based specifications for the design of controllers in wind turbines

Irene Miquelez-Madariaga[1], Jesús Arellano[2,1], Daniel Lacheta-Lecumberri[3,1], and Jorge Elso[1]

[1]Universidad Pública de Navarra, Pamplona, Spain
[2]Siemens Gamesa Renewable Energy, Sarriguren, Spain
[3]VORS Control, Pamplona, Spain

Correspondence: Irene Miquelez-Madariaga (irene.miquelez@unavarra.es)

Abstract. Pitch controllers are critical components in wind turbines, and their design typically involves lengthy iterative processes of tuning and simulation-based validation. Using well-defined control specifications, rigorously linked to engineering requirements, can significantly reduce the effort needed to achieve optimal performance. This work presents a methodology for generating control specifications directly related to mechanical fatigue caused by driving loads in wind turbine applications. The method is tailored for frequency-domain controller design techniques, such as QFT and H∞, and is based on Dirlik's method for fatigue assessment. It is tested through the design of controllers for a reference 15 MW wind turbine operating above rated conditions. The specification-driven design results in a validated reduction in fatigue, with simulation outcomes showing a mean error of 1.07% between the fatigue value predicted by the specifications and the one obtained from simulations.

1 Introduction

The design of wind turbines is an inherently multi-objective problem. To reduce the Levelized Cost of Energy (LCoE, Bruck (2018)), wind turbines should be designed to maximize the generated power and to minimize design, production and operation expenses. When focusing on the control strategies, these main objectives are translated into maximizing generated power, regulating the generator speed, and reducing the driving mechanical loads.

However, controller design is still somewhat disconnected from the evaluation of its performance. According to the standard (IEC (2020)), mechanical fatigue must be assessed using aeroelastic simulation of turbulent wind fields. The generation of the wind fields and their simulation are time-consuming steps incompatible with a swift workflow. Besides, to get the fatigue indicator, load cycles are quantified using the nonlinear rainflow counting algorithm and the Wöhler curve.

All in all, a gap exists between the controller design environment, which oftentimes uses linear frequency-domain models for the wind turbine and for the wind (Kaimal (1972); Mann (1998); Song (2022); Singh (2016)), and the evaluation environment, which requires nonlinear time-domain simulations. Some work has been done to solve this problem for a general wind turbine design application (Tibaldi (2016); Pao (2024)), but it has not targeted the design of control specifications − a key step in the design of controllers −. The main contribution of this article

is a method for the design of fatigue-based control specifications using linear models. The proposed solution has two main characteristics: (1) it identifies the range of frequencies in which the specification contributes more to the overall fatigue, and (2) it quantifies the expected fatigue decrease when a change in the specification or the controller is introduced.

In Section 2, a context on fatigue evaluation is presented, focusing on frequency domain methods. Then, in Section 3 the main contribution of this work is described, namely a method for obtaining control specifications based on the sensitivity of the damage equivalent load to changes in the signal spectrum. Section 4 validates the method on the design of controllers for a 15MW wind turbine. Lastly, Section 5 summarizes the main results and conclusions of the work.

2 Frequency domain assessment of mechanical fatigue

The IEC 61400 standard for the design of wind turbines (IEC (2020)) states that the mechanical damage caused by fatigue in a specific wind turbine configuration must be evaluated by the simulation of turbulent wind fields at different operating points using an aeroelastic model of the wind turbine. Each of the simulations produces a realistic time series of the mechanical loads. Then, the number and amplitude of the load cycles are computed using a rainflow-counting algorithm. Finally, the total accumulated damage is obtained by using the Wöhler curve (Wöhler (1870)) and a Palmgren–Miner linear damage hypothesis (Manson (1994)).

Although this process is compulsory to ultimately certify a wind turbine configuration, the same is not true during the intermediate step of controller design. Firstly, mechanical damage does not have an explicit relation with most design parameters, which means that the effect of a change in the design is unknown until tested in simulation. Additionally, the generation of wind fields and their simulation in the aeroelastic model are time-consuming steps that prevent the design process from being agile, especially due to its iterative nature.

The literature on fatigue assessment includes a set of methods that approximate fatigue with expressions that depend on the spectral properties of the load. The most basic approximation is the narrow-band method (Wirsching (1980)), an analytic method that provides an estimation of the fatigue damage in structures subjected to random stress processes. The method assumes that the stress process is narrow-band, thus simplifying the analysis by treating the stress cycles as approximately sinusoidal with a constant amplitude and frequency. The method estimates the number of stress cycles and their respective amplitudes using the properties of the narrow-band random process. Then, the damage caused by fatigue ΔD_{NB} is

$$\Delta D_{NB} = \nu_0 N_0^{-1} S_0^{-k} (2m_0)^{\frac{k}{2}} \Gamma\big(1+\frac{k}{2}\big), \tag{1}$$

where ν_0 is the zero-crossing rate, N_0 is a normalization factor related to the number of cycles, S_0 is the spectral width parameter, m_0 is the zeroth spectral moment, $\Gamma()$ is the Gamma function and k is a material exponent of the S-N curve.

Most of the existing methods are variations on the narrow-band method, which offer more precise results for wide-band processes (Benasciutti (2005, 2012)). They often follow an empirical approach and are precise when the analyzed load is similar to the training data. Among these empirical solutions for frequency-domain fatigue assessment, Dirlik's work (Dirlik (1985)) is one of the most accurate and widely accepted ones. It is specifically designed to handle a wide range of frequency content. This method provides an efficient and accurate approach to predicting fatigue life by leveraging the statistical properties of the stress response. Dirlik uses a probabilistic approach to estimate the distribution of stress cycles, deriving an empirical formula for the probability density function (PDF) of stress ranges in a random loading process. Essentially, it uses the spectral properties of the load and the mechanical properties that define the Wöhler-Curve. The formula accounts for the distribution of peaks and valleys in the stress history, providing a detailed statistical description. The empirical formula for the PDF is given by

$$\Delta D_{Dirlik} = \nu_p N_0^{-1} S_0^{-k} (m_0)^{\frac{k}{2}} \left[G_1 Q^k \Gamma(1+k) + 2^{\frac{k}{2}} \Gamma(1+\frac{k}{2})(G_2 |R|^k + G_3) \right], \tag{2}$$

where Q is a parameter related to the peak factor, which is derived from the ratio of higher-order spectral moments and characterizes the non-Gaussianity of the stress process and R is the mean value of the stress process normalized by its standard deviation, describing the relative location of the process mean concerning the stress amplitude. G_1, G_2, and G_3 are empirical expressions derived by Dirlik (Dirlik (1985)) that depend on the spectral moments

$$m_n = \int_0^\infty \omega^n S(\omega)\, d\omega \tag{3}$$

where n is the order of the spectral moment and $S(\omega)$ is the spectrum of the mechanical load.

3 Generation of fatigue-based specifications for the design of controllers

The use of controllers in dynamic systems primarily aims to modify the behavior of the system to achieve certain operational objectives. In this context, the design of the controller is based on a constraint known as the control specification.

In many design methodologies, this specification is expressed as a desired closed-loop frequency response, $W_y(i\omega)$. In the specific case of wind turbines, closed-loop functions allow for the description of the wind's effect on relevant variables, such as the generator's rotational speed or critical mechanical loads that determine the system's lifespan. For example, the effect of wind on the generator speed (Ω_g) is described by the closed-loop transfer function

$$T_{\Omega_g,W}(i\omega) = \frac{P_{\Omega_g,W}(i\omega)}{1 + C(i\omega) P_{\Omega_g,\beta}(i\omega)}, \tag{4}$$

and the effect of wind on the tower base load (MyT) is given by

$$T_{MyT,W}(i\omega) = \frac{P_{MyT,W}(i\omega) + \left(P_{\Omega_g,\beta}(i\omega) \cdot P_{MyT,\beta}(i\omega) + P_{\Omega_g,W}(i\omega) \cdot P_{MyT,W}(i\omega) \right) \cdot C(i\omega)}{1 + C(i\omega) P_{\Omega_g,\beta}(i\omega)}, \tag{5}$$

where P_y, u is the open loop transfer function that relates input $u(t)$ with output $y(t)$.

However, setting specifications directly on mechanical loads is complex. This complexity arises due to the gap between the design environment − typically based on linear models in the frequency domain − and the controller performance evaluation environment, where metrics are based on cycle counting and fatigue damage estimation from time series obtained via simulation. As a result, iterative trial-and-error-based design strategies are often employed. In such cases, fatigue evaluation is conducted *a posteriori* for each controller version, which prolongs the design process.

Methods like Dirlik's (Equations 2 and 3) help accelerate these iterative cycles, as they provide a fatigue damage estimate from the output signal spectrum $S_{MyT}(\omega)$. This spectrum, in turn, can be calculated using the closed-loop transfer function associated with the considered controller as follows

$$S_{MyT}(\omega) = |T_{MyT,W}(i\omega)|^2 S_W(\omega), \tag{6}$$

where $S_W(\omega)$ is the rotor effective wind spectrum.

The use of Dirlik's method, along with Equations 5 and 6, improves the efficiency of the design process by eliminating the need to perform aeroelastic simulations for each iteration. However, it does not remove the iterative nature of the design-validation-redesign cycle. From a control engineering perspective, it is preferable to have *ex ante* specifications that explicitly relate the impact of a new specification to the expected fatigue damage. This way, once the control specifications are selected, the chosen design method would result directly in a satisfactory controller.

A straightforward way to represent the effect of modifying the specification on the expected fatigue is to compute the sensitivity of damage to variations in the output spectrum. Due to the complexity of Equation 2, a numerical approximation of the derivative is used, based on an initial control configuration $C_0(i\omega)$, which results in an output spectrum $S_{yy,0}(\omega)$ and fatigue damage D_0. For greater generality, the initial configuration may also be the open-loop system ($C_0(i\omega) = 0$).

The sensitivity of fatigue damage, $\Delta D_{\%}(\omega)$, is calculated as

$$\Delta D_{\%}(\omega_0) = \frac{D_0 - D_{\omega_0}}{D_0} \cdot 100 \tag{7}$$

by introducing a 1% amplitude reduction in $S_{yy}(\omega)$ at frequency ω_0 for the whole frequency range in which the input spectrum is defined. There, D_{ω_0} is the damage associated with the modified spectrum. This function $\Delta D_{\%}(\omega)$ allows us to identify the frequency range where modifying the specification will have the greatest impact.

Then, the change in fatigue damage caused by a specification modification can be calculated as

$$\Delta D_{W_y} = \sum_{\omega} \Delta D_{\%}(\omega) \cdot \Delta S_{W_y}(\omega), \tag{8}$$

where $\Delta S_{W_y}(\omega)$ is the spectral variation caused by a change in the specification and can be calculated as:

$$\Delta S_{W_y}(\omega) = \frac{\Delta S_{W_{y,0}}(\omega) - \Delta S_{W_y}(\omega)}{\Delta S_{W_{y,0}}(\omega)} \cdot 100. \tag{9}$$

Finally, Equation 8 can be directly linked to the specification through the expression

$$\Delta D_{W_y} = \sum_{\omega} \Delta D_{\%}(\omega) \cdot \frac{|T_{yu}(i\omega)|^2 - |W_y(i\omega)|^2}{|T_{yu}(i\omega)|^2} \cdot 100. \tag{10}$$

In this way, the sensitivity $\Delta D_{\%}(\omega)$ helps designing the specification $W_{MyT}(i\omega)$ in two ways. First, it shows the range of frequencies in which a change in the specification has a greater impact in fatigue. Besides, it allows to quantify the impact of changing the specification on the fatigue by applying Equation 10.

4 Validation for the 15 MW reference wind turbine

The validation of the method proposed in this work is performed by using QFT for the design of linear controllers based on the fatigue specifications. The performance of the controllers is evaluated by their simulation in aeroelastic code (OpenFAST (2024)) and the post-processing of the resulting time series with a rainflow counting algorithm. For that purpose, the 15 MW reference wind turbine with the ROSCO controller (Abbas (2022)) is used.

4.1 System description

The wind turbine model used for this study is the IEA 15MW reference wind turbine (Gaertner (2020)). This turbine is an offshore, monopile model with three blades and a horizontal axis. The most relevant parameters of the model are gathered in Table 1.

Parameter	Value
Rotor diameter	240 m
Hub height	150 m
Cut-in rotor speed	5 rpm
Rated rotor speed	7.56 rpm

Table 1. Value of the main parameters of the 15 MW reference wind turbine.

Initially, this wind turbine operates using the reference Open-Source Controller (ROSCO) as a baseline control, which was developed to offer a modular control structure, with an industry-level performance and compatibility with the OpenFAST design and simulation environment. ROSCO includes the control strategies corresponding to the main operating regions of a wind turbine, from low-speed winds (Region 1) to above-rated wind speed (Region 3). There are two main control strategies, corresponding to pitch and torque controllers. Torque control is especially relevant at lower wind speeds, in which maximal power production is sought. ROSCO offers different strategies for below-rated operation, among which the quadratic control law has been chosen. Pitch control is used at above-rated wind speeds to ensure a constant generator speed and nominal power production.

For the purpose of this work, which is linear fatigue assessment, the analysis focuses on the above-rated operation. There, the control strategy consists of a set of linear controllers with scheduled gains to face the nonlinear dynamics of the wind turbine. Typically, these controllers follow a PI structure in which the integrator rejects the effect of wind in the lower frequencies, and the zero increases the phase margin of the system. Figure 1 shows the magnitude plot of the PI controller (lower plot) operating at a 19 m/s wind speed and its effect on the generator speed and tower base load. The main effect of the controller can be appreciated in the generator speed plot. There, the red line corresponding to the ROSCO controller shows the attenuation of the effect of wind on the generator speed at the lower frequencies.

Additionally, ROSCO includes switching, filtering and load reduction strategies, such as the Active Tower Damping (ATD). The ATD strategy reduces tower base fatigue by actively damping the first fore-aft natural frequency at the tower base. The ATD has been deactivated in the baseline control version in order to test the proposed method for the design of control specifications and, more specifically, its accuracy when predicting the quantitative impact of the new specification. As a result, a significant reduction in fatigue is expected regardless of the chosen controller design methodology.

The open-loop linear models of the wind turbine have been obtained with the aid of OpenFASTs linearization tool and the closed-loop models have been calculated following the structure represented in Figure 2.

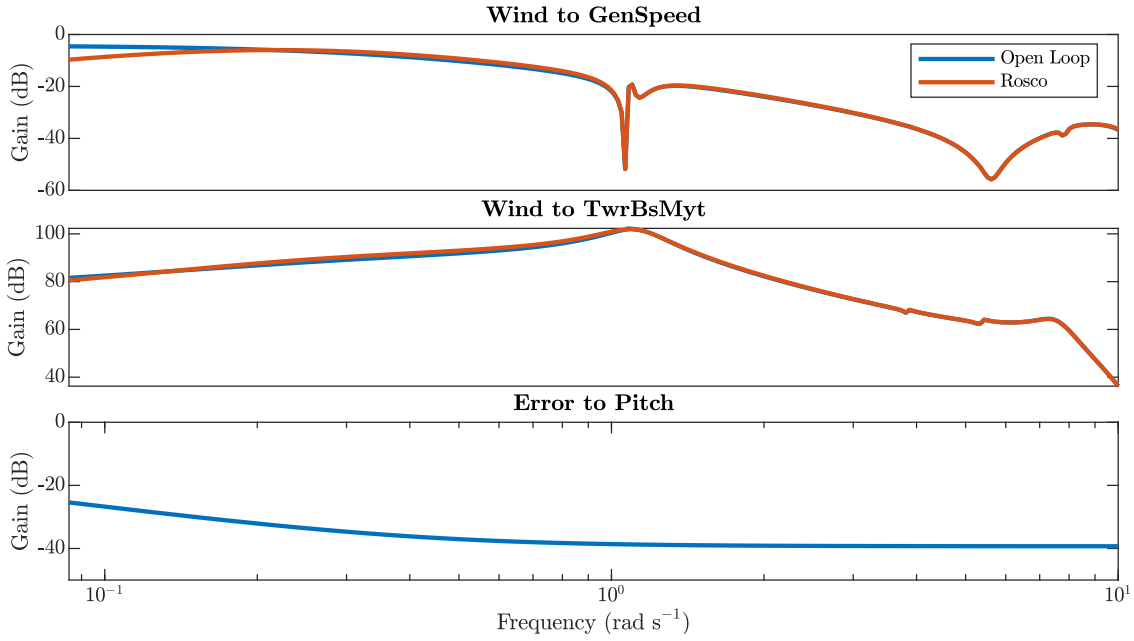

Figure 1. Effect of the ROSCO controller on the dynamics of the system. The upper plot shows the effect of the PI controller on the generator speed, which is more relevant at the lower frequencies. The middle plot shows the effect of the controller on the tower base bending moment. Lastly, the lower plot shows the magnitude of the feedback controller.

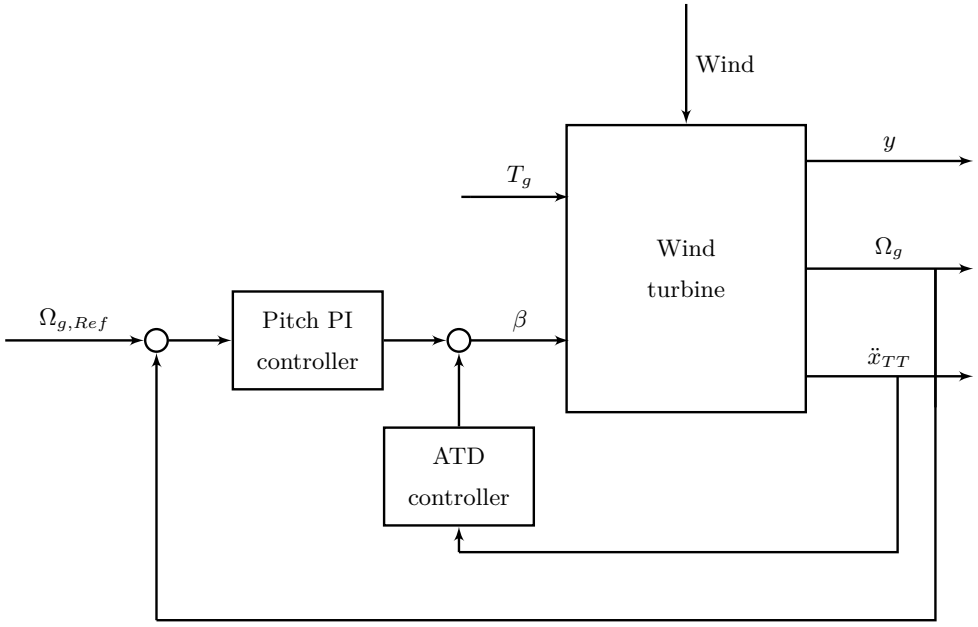

Figure 2. Block diagram of the control structure.

4.2 QFT fundamentals

Quantitative Feedback Theory is a controller design methodology that allows to obtain robust, multivariable, and multiobjective solutions (Elso (2017)). The main steps in the design process are:

1. Obtaining the model of the plant. QFT is based on the use of linear, frequency domain models of the system. Robustness is achieved by using uncertain models during the design.

2. Selecting the design frequencies. QFT design is performed on a discrete set of frequencies as a reference. These frequencies should cover all the relevant dynamics of the system.

3. Choosing the specifications. In the context of QFT design for wind turbine controllers, control specifications are typically upper bounds on the closed-loop disturbance rejection function of the different control objectives. These upper bounds can be constant values or transfer functions and should ideally be linked to performance indicators such as the generator speed standard deviation or the mechanical fatigue of driving components. Besides, a stability specification is imposed as an upper bound on the complementary sensitivity function.

4. Generation of bounds. At each design frequency, the regions of all controllers meeting each specification are found in the complex plane. The boundary of these regions, multiplied by the nominal plant, in the complex plan is called bound. One bound is obtained per specification and design frequency.

5. Controller synthesis, also known as loopshaping, consists of tuning the controller parameters until the open-loop transfer function that lies within the bounds at every design frequency.

6. Checking the specifications. The first validation step uses the uncertain linear model of the plant to ensure that specifications are met at every frequency and not only the design ones.

7. Simulation. If the linear uncertain model of the plan has been derived from a more complex mode (i.e. a nonlinear one), as is the case for wind turbines, a second evaluation of the performance of the controllers is performed via simulation.

The whole process is carried out with the aid of the QFT Toolbox (Yaniv (1997)), which includes graphical tools for the design of specifications and the loop-shaping process. QFT is inherently an iterative method, more so when several controllers are being tuned at the same time, and requires some practice in the loop-shaping stage before a succesful design has been obtained. As a counterpart, it is a versatile and transparent methodology, that grants engineers full control of the design process.

4.3 Controller design

The design of controllers starts with the design of the control specifications based on the procedure presented in Section 3. The first step consists of obtaining a set of simulations of the system with the baseline control. In particular, four wind seeds with length 600 s have been simulated at each operating point. Then, these simulations are used to calculate the output spectra and an initial fatigue evaluation. Lastly, the sensitivity of fatigue to changes in the tower base load spectrum $\Delta D_{\%}$ is calculated based on the simulation results.

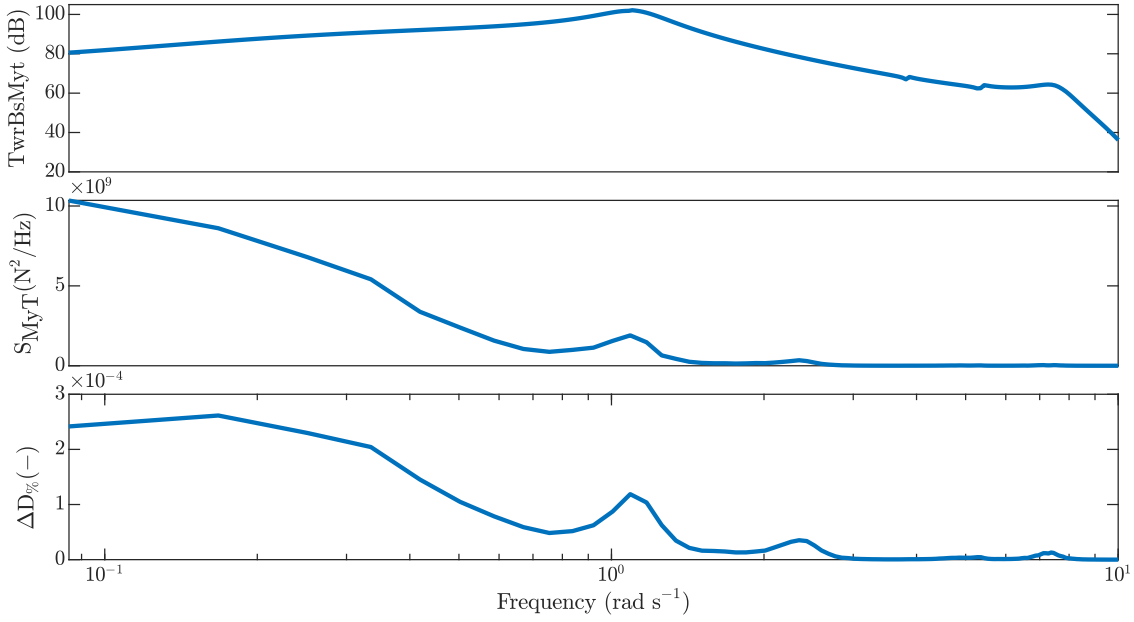

Figure 3. This figure shows different aspects of the relation between wind and tower base load. The upper plot shows the magnitude of the closed loop transfer function between wind and load $|T_{MyT,W}(i\omega)|$. The middle plot, that represents the load power spectrum $S_{MyT}(\omega)$, shows how the biggest amplitudes of the load appear at the lower frequencies. The lower plot represents the fatigue sensitivity $\Delta D_{\%}$.

Figure 3 shows the relation between the magnitude Bode plot, the load spectrum and the variation of fatigue for the simulations corresponding to a 19 m/s mean wind speed. While the Bode plot only holds information on the system response, the spectrum of the load includes information on the disturbance. Besides, as the spectrum has been obtained from simulation data, a peak can be observed at around 2.3 rad/s, which corresponds to the 3P frequency. Lastly, the lower plot shows how the lower frequencies have a greater impact on fatigue, but the peaks corresponding to the first fore-aft mode of the tower and the 3P frequency are also relatively relevant.

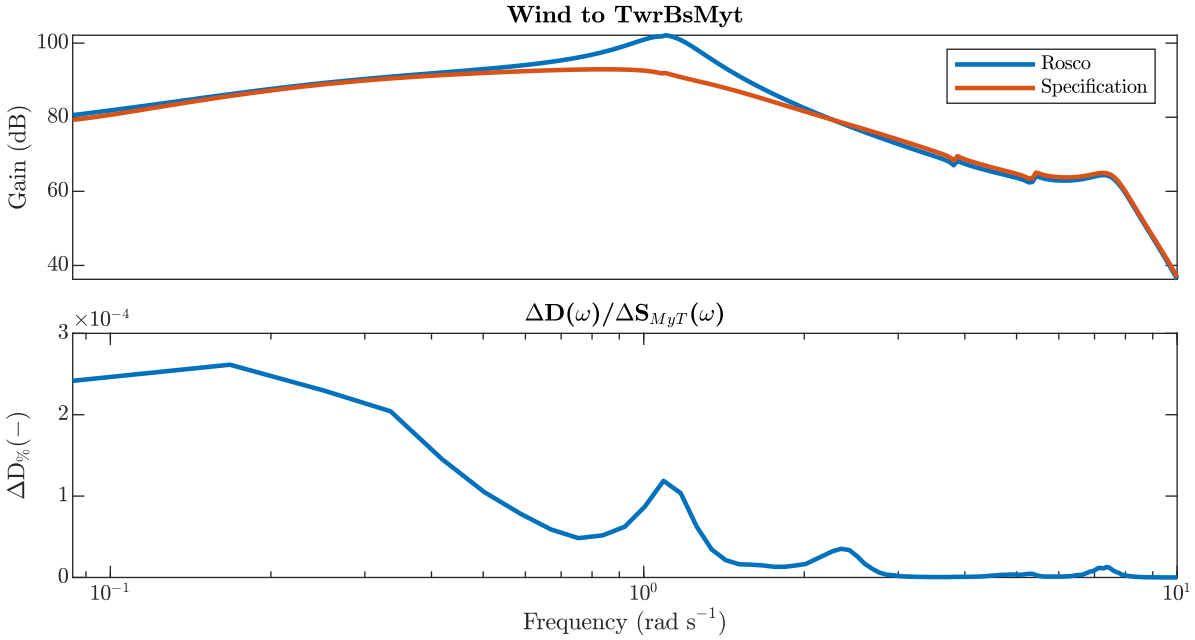

Figure 4. Design of the control specification for an operating point corresponding to a 19 m/s wind speed. The control specification has a similar magnitude as the ROSCO controller at all frequencies but the one surrounding the first fore aft mode of the tower (1.1 rad/s).

The design of the control specification $W_y(s)$, which takes as a starting point the closed loop transfer function produced by the baseline controller, is based on the information provided by function $\Delta D_{\%}(\omega)$. Even though the lower frequencies have the greatest impact on fatigue (see the lower plot in Figure 4), the designer must keep in mind that the main control objective is to have a constant power production, for which the lower frequencies are key. As a consequence, the control effort is located around the first fore-aft mode of the tower, which appears at approximately 1.1 rad/s. The control specification has been obtained with the aid of the *lpshape* function of the

QFT Toolbox targeting a fatigue reduction of 15%. A set of notch filters has been used to reduce the magnitude of the specifications at the chosen frequencies until the desired reduction of fatigue has been obtained or improved. The linear fatigue estimation for the specifications at different operating points appears in Table 2.

As already mentioned, power production and, in turn, generator speed regulation are the most relevant objectives in the design of the pitch controllers. As a consequence, they must also be included in the design specifications. In this case, the closed loop transfer function (wind to generator speed) obtained with the ROSCO controller is used as a specification. This way, a performance similar to the baseline controller is expected in terms of generator speed regulation.

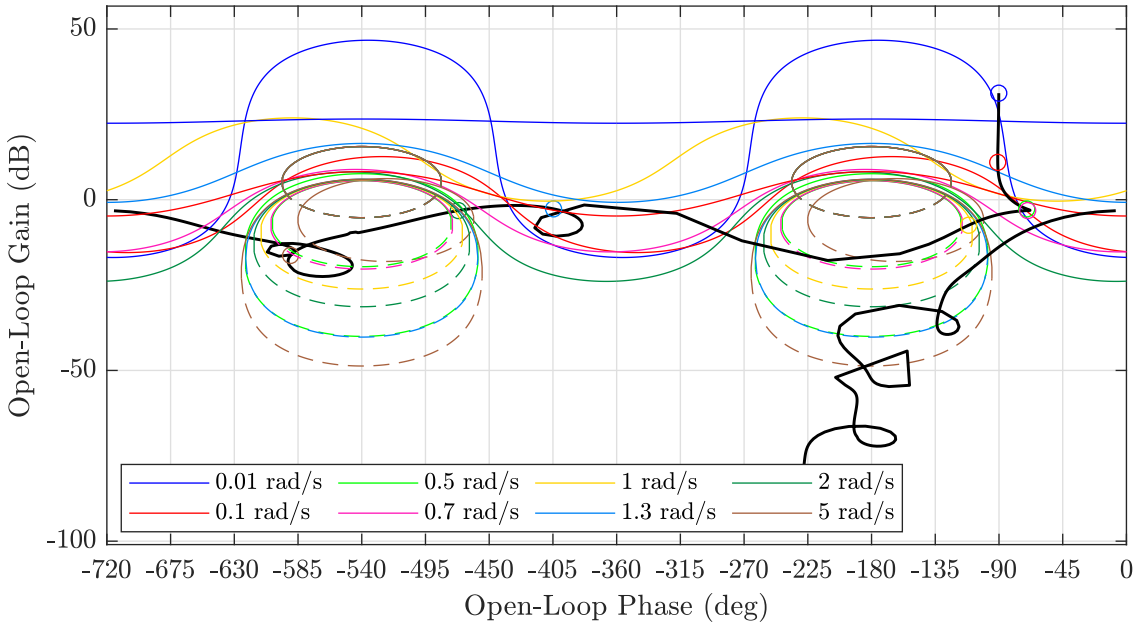

Figure 5. Bounds for the loopshaping of the main feedback pitch controller. Then main characteristic of this controller is the presence of an integrator to eliminate steady-state error and the use of low frequency poles to increase magnitude and phase below 1 rad/s. The open-loop nominal transfer function meets all bound at all frequencies.

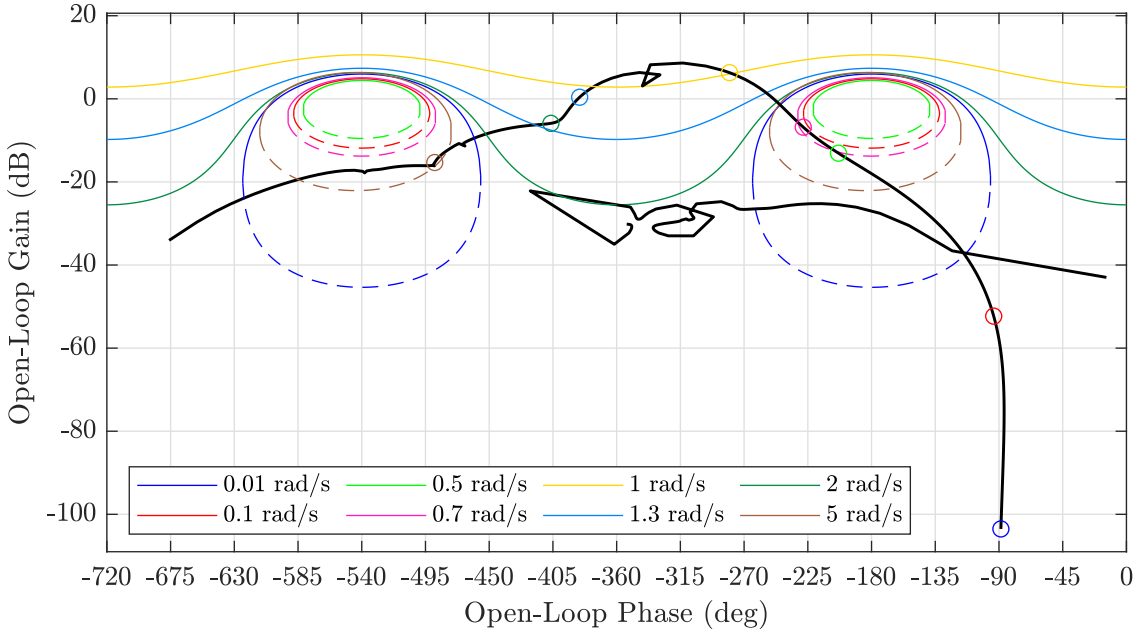

Figure 6. Bounds for the loopshaping of the ATD controller. This controller has a pole 0.3276 rad/s and a zero at 4.357 rad/s. The main control effort of the ATD strategy appears close to 1rad/s, at the first fore-aft mode of the tower.

Once the control specifications for the two control objectives have been designed, the bounds for the controllers have been obtained using the procedure described in Section 4.2. The bounds are obtained using function *genbounds* and the design of controllers is performed with the aid of *lpshape*. *lpshape* is a graphical design tool that represents the bounds and the open-loop nominal transfer functions in the Nichols plot (gain in dB against phase in degrees). The bounds are the limits between the allowed and forbiden values for the open-loop nominal transfer function. They can have different colours, depending on their corresponding frequency, and solid and dashed lines indicating lower and upper limits. The open-loop nominal transfer function is represented by the solid black line and a set of circular markers at the design frequencies, with the same colour as their corresponding bounds. The design is performed by modifying the position of the markers by adding zeros and poles to the controller.

As two different controllers have been designed for each operating point -the feedback pitch controller and the ATD- the design process is iterative. Due to the lack of ATD controller in the baseline configuration, the design process begins with the tuning of a preliminary ATD controller. Then, several iterations are performed until specifications are met. Figures 5 and 6 show the final design of the pitch and ATD controllers for the operating point corresponding to a 19 m/s average wind speed. At this wind speed, the pitch controller is

$$C_\beta(s) = \frac{-4.8167(s+0.6)(s^2+0.7337s+0.2441)}{s(s^2+1.872s+2.306)(s^2+2.297s+26.67)}. \tag{11}$$

The ATD controller at the same operating point is

$$C_{ATD}(s) = \frac{0.014187(s+4.357)}{s+0.3276}. \tag{12}$$

The bounds are met tightly at frequencies surrounding 1 rad/s and with a greater margin for lower frequencies, as seen in Figures 5 and 6. A similar conclusion can be inferred from the Bode plots presented in Figure 7. There, the markers represent the design frequencies, which are used for the calculation of bounds and the loop-shaping process. As already anticipated, at frequencies below 1 rad/s, at least one of the specifications is met with some margin. At higher frequencies, the tower base load is exactly equal to the specification. With a simple enough control structure, a good result is expected in between design frequencies. Figure 7 shows that the QFT controller provides a poorer response than the specification between 0.8 and 1 rad/s and 2 and 5 rad/s for the tower base load. However, the overall response is better than the specification and shows a significant improvement concerning the ROSCO controller.

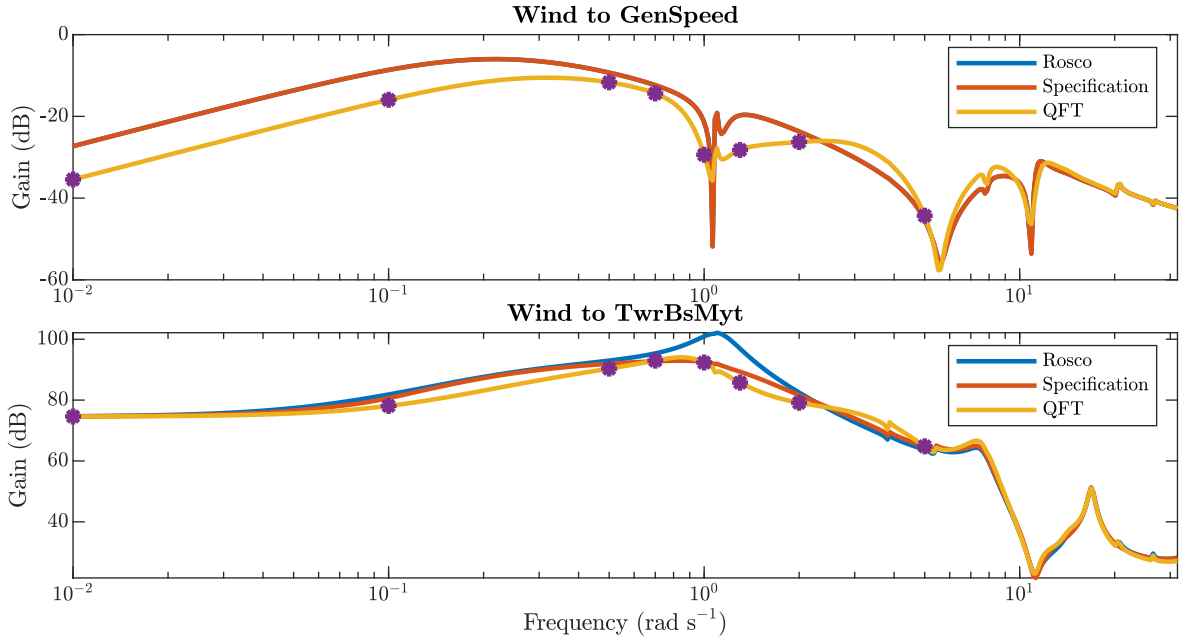

Figure 7. Magnitude Bode plot for ROSCO (blue), the design specifications (red) and the QFT controller (yellow) for the generator speed (upper plot) and the tower base load (lower plot). The purple dots show the design frequencies in which the specifications are always met. A simple control structure guarantees that the response between design frequencies is smooth and that specifications are meth most of the times.

4.4 Result analysis

Once a pitch controller and an ATD controller have been designed for each operating point in the above-rated region (11 to 25 m/s), they have been integrated into the aeroelastic simulator using its Simulink interface. Due to the

different structures of the controllers at different operating points, an output blending strategy has been used for the interpolation between controllers. A filtered pitch signal has been used as an interpolation variable.

At each operating point, four different seeds have been used for the generation of turbulent wind fields, to ensure a reasonable trade-off between results variability due to turbulence and computation time. Each wind file has a 600 s length and a class B turbulent intensity according to the Normal Turbulence Model (NTM, Ishihara (2012)) and the IEC:61400:1-2019 standard. Wind fields have been generated using the tool Turbsim (Jonkman (2006)).

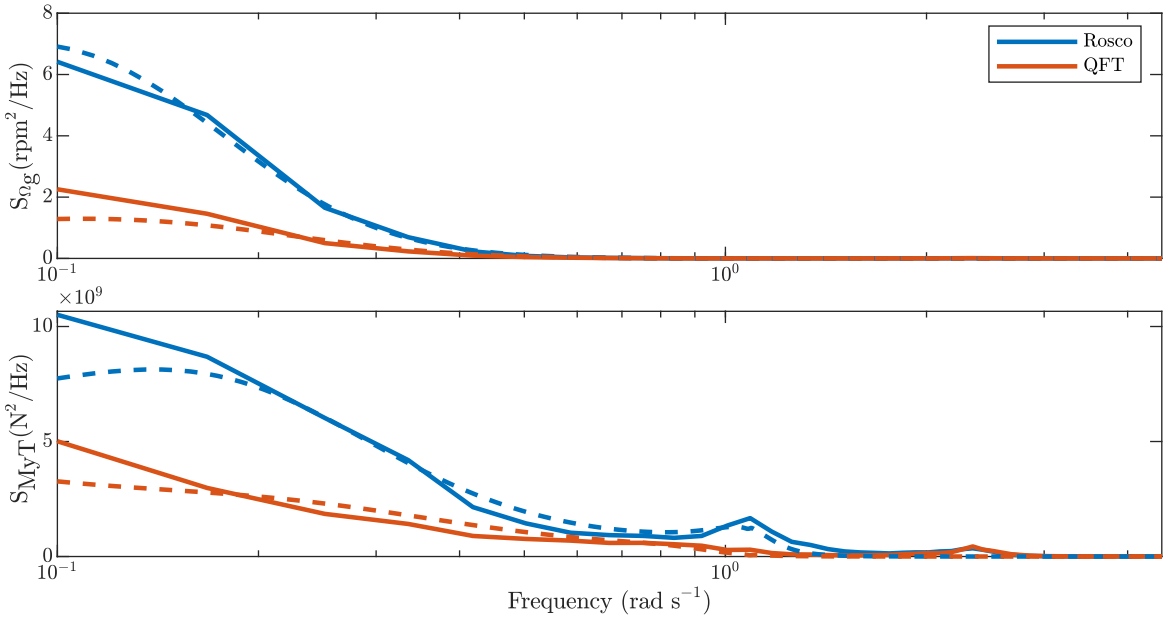

Figure 8. Validation of the correspondence between linear (dashed) and nonlinear (solid) models with the ROSCO (blue) and QFT (red) controllers. Besides a good matching between both models, the effect of the QFT controllers can be observed in the lower magnitude of the spectra, especially in frequencies below 0.5 rad/s and in the peak corresponding to the first fore-aft mode of the tower (close to 1 rad/s).

The first step in the analysis of the results consists of the validation of the relation between the linear and the nonlinear models. Figure 8 shows a comparison between the theoretical and simulation spectra of the generator speed and tower base bending moment with the baseline and the QFT controllers at an operating point of 19 m/s. The theoretical spectra have been obtained using the linear model of the wind turbine and the theoretical spectrum of rotor effective wind speed. The simulation spectra have been obtained by applying the Welch method to the time series obtained from OpenFAST. Two main conclusions can be extracted from this figure: (i) the QFT controller outperforms the ROSCO controller for both outputs and (ii) there exists a good correspondence between the estimation obtained from the linear model and the results of the nonlinear simulation. This second fact is critical for the performance of the proposed fatigue estimation method.

Performance parameters	Operating point (m/s)						
	15	17	19	21	23	25	Mean
Generator speed std	31.54%	-1.91%	-39.48%	-48.96%	-50.07%	-52.18%	-0.14%
Mean power	0.47%	0.37%	0.05%	-0.17%	-0.42%	-0.46%	0.28%
Tower base DEL	-25.91%	-22.24%	-27.75%	-27.28%	-34.71%	-28.84%	-25.69%
Pitch activity	-8.22%	-0.53%	1.13%	6.66%	1.87%	5.43%	-2.77%

Table 2. Comparison of the main performance indicators for the ROSCO and QFT controllers.

Even though improving the performance of ROSCO is not the main objective of this work, the performance of the QFT controller has been studied to understand the possibilities of this design technique and the fatigue-based design of the specifications. The performance of the controllers has been evaluated using four main indicators. The standard deviation of the generator speed is used as a measure of the quality of the generator speed regulation. Table 2 shows how the standard deviation is reduced for all operating points but 15 m/s, which is closer to the transition between regions. The mean generated power shows small variations at different operating points that cancel each other when calculating the total average. As expected from the linear fatigue prediction, the fatigue at the tower base is reduced significantly at all operating points, partly due to the absence of ATD in the baseline controller. Lastly, pitch activity is evaluated using the standard deviation of the collective pitch signal, which is significantly reduced at 15 m/s and then increases for wind speeds 19 to 25 m/s. A more thorough analysis of the cost of pitch control should be based on the actuator fatigue, where less favourable results would be expected. Mean values have been obtained assuming a Weibull distribution in mean wind with shape factor of 2.2 and a scale factor of 11.29.

Similar information can be perceived in Figure 9, which represents the simulation outputs for a single wind seed with a mean speed of 19 m/s. Both the generator speed and the tower base load show smaller deviations from their mean value for the QFT controller, which accounts for a smaller standard deviation and fatigue respectively. The power mean value is close to 15 MW in both cases. The presence of the ATD controller in the QFT configuration can be observed in the ripple that appears in the pitch signal, accounting for an increased pitch standard deviation.

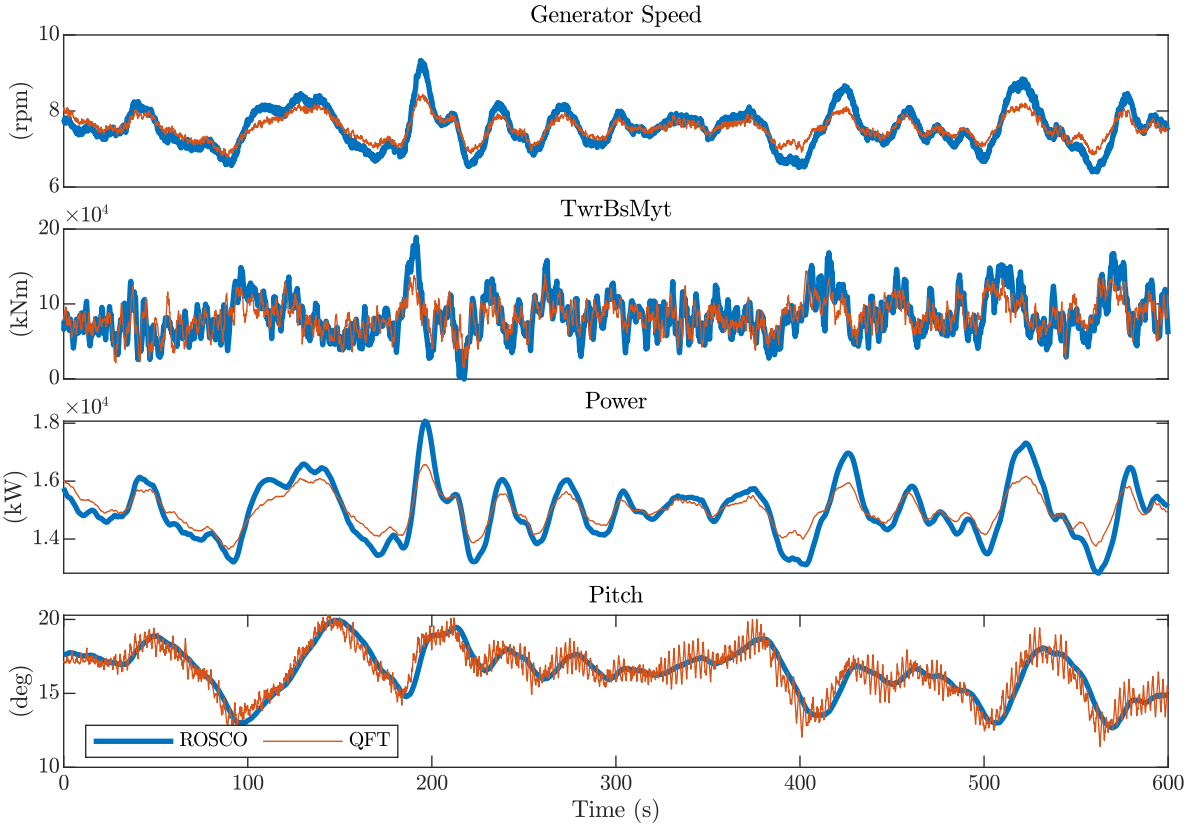

Figure 9. Time series of relevant variables in an OpenFAST simulation with mean wind speed 19 m/s. The standard deviation of generator speed and power is visibly smaller for the QFT controller (red) than for ROSCO (blue). Besides, the fatigue in the tower base load should also be reduced as the ATD controller mitigates the bigger peaks in the blue line, that do not appear for the QFT controller. Lastly, the lower plot shows that the use of an ATD controller results on an increased high frequency pitch activity.

The main result of this work is the design of the specifications on the mechanical loads based on a quantitative indicator of fatigue. Table 3 shows the fatigue reduction promised by the specifications, the linear estimation of the fatigue reduction for the new controllers and the actual fatigue reduction calculated using a rainflow counting algorithm. Two main conclusions can be extracted from the data: (i) the fatigue predicted by the linear model is always smaller than the promised by the specifications except from the 15 m/s operating point and (ii) the deviation between the linear prediction and the rainflow-based fatigue evaluation is smaller than 2% for all operating points, having an average value of 1.07%.

Fatigue estimation	Operating point (m/s)						
	15	17	19	21	23	25	Mean
Specifications	37.37%	10.10%	11.78%	20.55%	35.97%	30.35%	
Linear prediction	24.81%	24.15%	28.60%	27.93%	35.12%	30.36%	
Rainflow evaluation	25.91%	22.24%	27.75%	27.28%	34.71%	28.84%	
Estimation deviation	1.1%	1.91%	0.85%	0.65%	0.41%	1.52%	1.07%

Table 3. Fatigue reduction at different operating points. The table includes information on the fatigue reduction promised by the specification, the reduction estimated with the help of linear models and the actual fatigue reduction as evaluated with a rainflow counting algorithm.

Figure 10 shows a graphical representation of the accuracy of the linear approximation of fatigue by plotting the
285 frequency-domain estimation of fatigue and the time-domain evaluation against the mean wind speed. Even though the achieved fatigue reduction varies significantly in the different operating points, the estimation error remains much smaller than the fatigue variation.

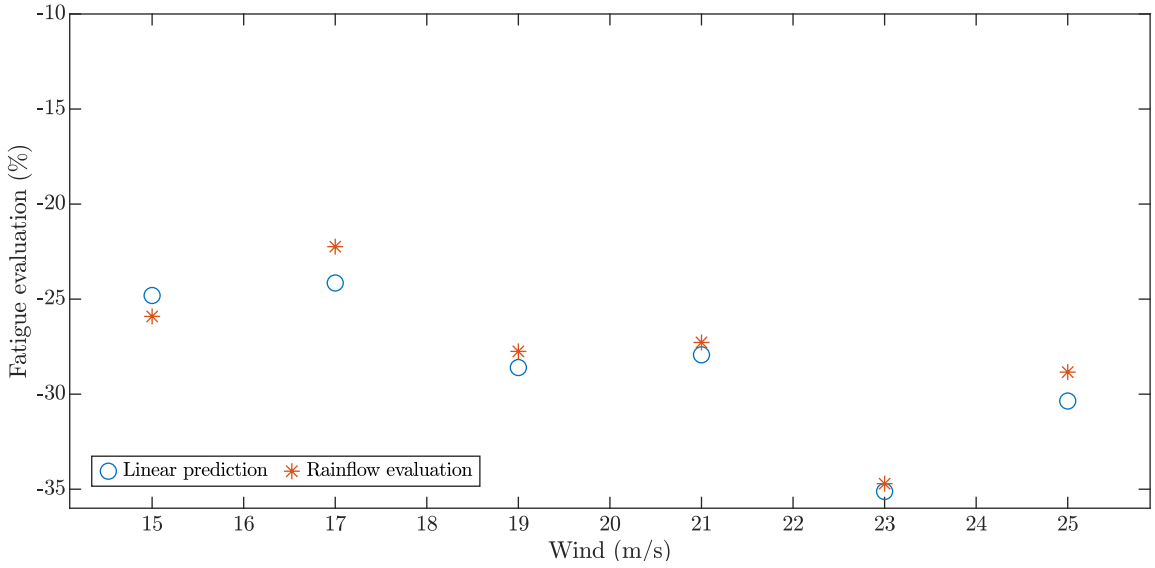

Figure 10. Graphical comparison of the linear fatigue estimation (linear prediction) and the rainflow-based fatigue assessment (Rainflow evaluation).

5 Conclusions

This article presents a method for designing control specifications based on mechanical fatigue. By having a priori information about the expected fatigue, the number of design iterations and the time spent on aeroelastic simulations are reduced, making the control design workflow smoother.

The method has been validated with the design of pitch controllers for the 15 MW reference wind turbine. More specifically, the feedback pitch controller has been redesigned and the active tower damping have been introduced to improve the performance of the ROSCO PI controller. With the aid of the linear fatigue estimation and the design of specifications, fatigue has been reduced by 22 to 36% at operating points ranging from 15 to 25 m/s mean wind speed. Due to the good correspondence between the linear fatigue prediction and the fatigue evaluation obtained with a rainflow algorithm (under a 2%), a single iteration in the non-linear simulation step has been required. Consequently, the design process has been accelerated significantly, by reducing the number of required simulations to two sets.

The error present in the fatigue estimation can be attributed to three different factors. The first one is that the results provided by the linear model do not exactly match the ones provided by the nonlinear simulator. In addition to the numerous non-linearities present in the realistic model of the wind turbine and the complete control structure, other phenomena such as the 3P frequency and the spatial variation of the wind cannot be taken into account by a linear approximation. The second factor is the difference between time-domain and frequency-domain fatigue evaluation methods. The rainflow-counting algorithm has proven to be the best approximation for fatigue estimation and is the reference in the standard. On the other side, due to the fundamental empirical nature of frequency-domain methods, their performance varies significantly depending on the characteristics of the load and the system. While Dirlik's method has proven to provide a good result, the search for a more accurate method remains open. Lastly, the proposed method is based on a linear approximation to fatigue estimation, which can introduce an error.

All in all, the use of linear models for specification design has proven to be useful and their applications can be extended. On the one hand, the use of this frequency domain estimation of fatigue could be included in the iterative design of other wind turbine components, such as structural elements. On the other hand, their use could also be studied for different controller design methodologies, the most obvious one being H_∞. Lastly, the full potential of the methodology could be achieved for multi-objective design, in which the role of the control system in the fatigue of different elements could be linked to the business case of the wind turbine.

Author contributions. I. Miquelez-Madariaga: Conception, simulation data collection, writing. J. Arellano: Conception, writing. D. Lacheta-Lecumberri: Simulation data collection, proofreading. J. Elso: Conception, proofreading.

Competing interests. The contact author has declared that none of the authors has any competing interests.

References

Menezes, E. J. N., Araújo, A. M., Da Silva, N. S. B. (2018). A review on wind turbine control and its associated methods. Journal of cleaner production, 174, 945-953.

Bruck, M., Sandborn, P., and Goudarzi, N. (2018). A Levelized Cost of Energy (LCOE) model for wind farms that include Power Purchase Agreements (PPAs). Renewable Energy, 122, 131-139.

Kaimal, J. C., Wyngaard, J. C. J., Izumi, Y., & Coté, O. R. (1972). Spectral characteristics of surface-layer turbulence. Quarterly Journal of the Royal Meteorological Society, 98(417), 563-589.

Mann, J. (1998). Wind field simulation. Probabilistic engineering mechanics, 13(4), 269-282.

Song, Y., Jeon, T., Paek, I., & Dugarjav, B. (2022). Design and validation of pitch H-infinity controller for a large wind turbine. Energies, 15(22), 8763.

Singh, N., Pratap, B., & Swarup, A. (2016). Design of robust control for wind turbine using quantitative feedback theory. IFAC-PapersOnLine, 49(1), 718-723.

Pegalajar-Jurado, A., Borg, M., Bredmose, H. (2018). An efficient frequency-domain model for quick load analysis of floating offshore wind turbines. Wind Energy Science, 3(2), 693-712.

International Electrotechnical Commission. (2020). IEC 61400: Wind energy generation systems.

Pao, L. Y., Pusch, M., Zalkind, D. S. (2024). Control co-design of wind turbines. Annual Review of Control, Robotics, and Autonomous Systems, 7.

Tibaldi, C., Henriksen, L. C., Hansen, M. H., & Bak, C. (2016). Wind turbine fatigue damage evaluation based on a linear model and a spectral method. Wind Energy, 19(7), 1289-1306.

Wöhler, A. (1870). Über die festigkeitsversuche mit eisen und stahl. Ernst & Korn.

Manson, S. S. and Newman, M. E. J. (1994). Fatigue of Materials. Springer, New York, 2nd edition edition.

Wirsching, P. H., & Light, M. C. (1980). Fatigue under wide band random stresses. Journal of the Structural Division, 106(7), 1593-1607.

Benasciutti, D., & Tovo, R. (2005). Spectral methods for lifetime prediction under wide-band stationary random processes. International Journal of fatigue, 27(8), 867-877.

Benasciutti, D. (2012). Fatigue analysis of random loadings. A frequency-domain approach. LAP Lambert Academic Publishing AG & Co KG.

Dirlik, T. (1985). Application of computers in fatigue analysis (Doctoral dissertation, University of Warwick).

Soltani, M. N., Knudsen, T., Svenstrup, M., Wisniewski, R., Brath, P., Ortega, R., Johnson, K. (2013). Estimation of rotor effective wind speed: A comparison. IEEE Transactions on Control Systems Technology, 21(4), 1155-1167.

OpenFAST. Available at https://github.com/openfast/openfast/.

Abbas, N. J., Zalkind, D. S., Pao, L., & Wright, A. (2022). A reference open-source controller for fixed and floating offshore wind turbines. Wind Energy Science, 7(1), 53-73.

Gaertner, Evan, et al. Definition of the IEA 15-megawatt offshore reference wind turbine. 2020.

Elso, J., Gil-Martinez, M., & Garcia-Sanz, M. (2017). Quantitative feedback control for multivariable model matching and disturbance rejection. International Journal of Robust and Nonlinear Control, 27(1), 121-134.

Yaniv, O., Chait, Y., & Borghesani, C. (1997). The QFT control design toolbox for MATLAB. IFAC Proceedings Volumes, 30(16), 103-108.

Ishihara, T., Yamaguchi, A., Sarwar, M. W. (2012). A study of the normal turbulence model in IEC 61400-1. Wind Engineering, 36(6), 759-765.

Jonkman, B. J. (2006). TurbSim user's guide (No. NREL/TP-500-39797). National Renewable Energy Lab.(NREL), Golden, CO (United States).