Analysis and multi-objective optimisation of model-based wind turbine controllers

Livia Brandetti[1,2], Sebastiaan Paul Mulders[2], Yichao Liu[2], Simon Watson[1], and Jan-Willem van Wingerden[2]

[1]*Flow Physics and Technology, Faculty of Aerospace Engineering, Delft University of Technology, Delft, The Netherlands*

[2]*Delft Center for Systems and Control, Faculty of Mechanical Engineering, Delft University of Technology, Delft, The Netherlands*

Email: l.brandetti@tudelft.nl

The authors appreciate the time and effort the reviewers dedicated to providing feedback on our paper. We are grateful for the insightful comments and valuable improvements to our manuscript. We have incorporated the suggestions made by the reviewers. Please see below, in blue, for a point-by-point response to the reviewers' comments and concerns. All page numbers refer to the revised paper.

Reviewers' comments to the authors:

Reviewer 1

The paper compares different state-of-the-art methodologies for partial load control of wind turbines. The authors conclude that an advanced Tip Speed Ratio-tracking scheme cannot outperform the traditional kw2 controller in terms of power maximization. The advanced controller can, however, help to reduce loads, while reducing the bandwidth of the controller.

The paper is very well written, well-structured and does not require larger revisions. Please find a few comments in the pdf attached and consider them for the revised version of this paper.

Line 39 Make clear why this is a problem for kw2 and not for TSR-tracking.

Author response: Thank you for your comment. This part of the introduction is a literature study dedicated to the disadvantages of the Kw2 controller found in previous studies rather than a comparison with the combined WSE-TSR tracking controller. A similar discussion devoted to the WSE-TSR tracking controller is provided later on in this introduction. We, however, realise and agree with the reviewer that the formulation of the text needed to be clarified, and we have accordingly made improvements in the introduction. Furthermore, we have added a reference to a previous work that discusses this issue.

Line 129 Provide a reference showing that this statement is true.

Author response: Thank you for pointing this out. We have clarified this statement and provided references showing that it is true, as:

Lines 132 to 134:“ Therefore, modern large-scale wind turbines are controlled by more advanced WSE-TSR tracking schemes (Mulders et al., 2023), and wind turbine manufacturers are currently exploring the possibilities of applying model predictive control (MPC) to provide such flexibility (Hovgaard et al., 2015; Pamososuryo et al., 2023).”

Line 146 Explain why this one is used, in terms of rating and state-of-the-art.
Line 377 Is there a summary table with the NREL 5MW main operational parameters?

Author response: We think these are excellent suggestions. Accordingly, we have added a motivation for why the NREL 5MW is used and a table (Table 1) summarising the NREL 5MW main operational parameters. The revised text reads as follows:

Lines 152 to 155:” This study focuses on showing the potential benefits of an advanced controller for large-scale turbines at both onshore and offshore locations. Therefore, for its size and rated power capacity, the National Renewable Energy Laboratory (NREL) 5 MW (Jonkman et al., 2009) wind turbine model is used to strike a balance. The main operational parameters are summarised in Table 1, and the Cp(.) curve covering the operating region of interest is illustrated in Figure 1.”

Section 3 It'd be useful to understand what exactly is linearized (model or controller), which inputs, outputs and states are defined and what exactly is calculated in the frequency-domain (load spectra given a wind spectrum?).

Author response: Thank you for your comment. The full-system dynamics are obtained by first analysing the dynamics of the individual components (i.e. the wind turbine in Section 3.1, the estimator in Section 3.4.1, the tip-speed ratio tracking controller in Section 3.4.2) and then by interconnecting them to find the closed-loop system dynamics (i.e. Section 3.2, 3.3, 3.4.3). The end results are the transfer functions, relating the amplitude and phase of the output as a function of the frequency to constant periodic inputs. We evaluate these transfer functions in the frequency domain; no load spectrum is involved in the analysis. Furthermore, for every transfer function, the input, the output and the state used are defined in the main text, for example:

Lines 211 to 212:” Subsequently, the resulting expression is linearised with respect to the rotor speed state, generator torque control input, and wind speed disturbance input, resulting in”

Lines 228 to 229:” Here, the controllers are generalised as a single block with two inputs and one output, being the reference tip-speed ratio, rotor speed, and generator torque control signals, respectively.”

Lines 267 to 268:” As illustrated in Figure 4, the estimator has the generator torque and the rotor speed as inputs and the estimated tip-speed ratio as output.”

Lines 282 to 283:” According to Figure 3, the TSR tracking controller has two inputs, the tip-speed ratio estimate and set-point, and one output, the generator torque.”

After careful consideration, we believe the existing text effectively conveys the intended information. Additionally, the approach employed in this paper has been verified against a state-space linearisation approach for the same coupled system in (Mulders et al., 2023), providing equal results.

Section 3.3, Equations 19 and 20 Have you verified that this linearization is suitable for the considered load spectra, i.e. have you compared nonlinear against linear response?

Author response: Thank you for the comment. The intent of the paper is to analyse and evaluate the controller and system using open-loop and closed-loop transfer functions. Section 3.2 introduces an analysis framework to evaluate the controller transfer functions (i.e. open-loop) for both the baseline and WSE-TSR tracking controllers. Using this framework, it is possible to generalise the controllers as a single block with two inputs (i.e. the reference tip-speed ratio and the rotor speed) and one output (i.e. the generator torque control). In the linear and frequency-domain formulation, the control scheme is formalised as follows:

$$\mathcal{T}_{\mathrm{g}}(s) = K_{\Omega_{\mathrm{r}} \to \mathcal{T}_{\mathrm{g}}}(s)\Omega_{\mathrm{r}}(s) + K_{\Lambda_{*} \to \mathcal{T}_{\mathrm{g}}}(s)\Lambda_{*}(s)$$

This framework is then used in Section 3.5 to provide an open-loop comparison and analysis by frequency responses of both controllers. One could use this linear analysis framework to capture the nonlinearity to analyse the system and controller characteristics at different operating points.

Moreover, to ensure agreement of the proposed linear with the nonlinear controller expression, we conducted a time-domain evaluation using the mid-fidelity software OpenFAST. The results of this comparison are summarised in the Figure below. The linear model shows a good agreement with the nonlinear response of the controller for deviations of the generator torque around a specific operating point.

To further enhance the clarity of our work, we have revised the manuscript, specifically in Section 3, where we describe the frequency-domain framework (lines 203 to 210). These revisions provide a more comprehensive explanation of our methodology and the verification steps taken.

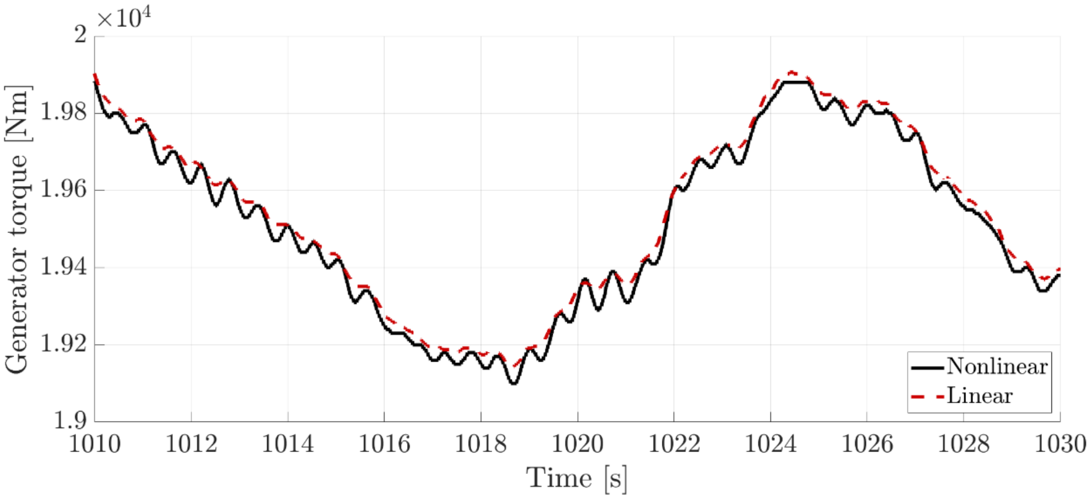

Figure - Verification between nonlinear and linear controller response

Line 304 Introduce the word static
Line 315: Clarify if you mean static or dynamic behavior?

Author response: Thank you for your comments. We fully agree with the reviewer. The revised text reads as follows:

Lines 316 to 317:"Equations (19) and (20) show that the controller transfer functions are merely frequency-independent static gains for the baseline controller."

Lines 327 to 328:" Thus, the two controllers will have the same static behaviour (Aström and Murray, 2010), operating at the same point of power extraction efficiency, Cp,∗(λ∗)."

Section 4 Discuss whether the WSE should be calibrated independent from the controller in order to obtain clear results without estimator-controller coupling effects or not.

Author response: Thank you for your insightful comment. As demonstrated in Appendix A, the considered control strategy shows similarities to a state-feedback controller design. According to the separation principle in control theory, we believe designing an optimal feedback controller for the considered stochastic system is indeed possible by independently designing the estimator and the controller. However, this claim should be further elaborated upon in future work, as in this paper, the estimator and controller are always calibrated in unison. To further clarify this aspect, we have revised the manuscript as follows:

Lines 389 to 391:" Furthermore, as can be recognised from the defined input vectors Γd, the estimator and the controller are consistently and intricately calibrated in unison throughout the entire work."

Line 328 Introduce word combined.

Author response: Thank you for your comment. We have revised the text accordingly.

Line 376 Commonly, high-fidelity is computational fluid dynamics. Suggest to call it "reference simulation".

Author response: Thank you for your suggestion. We have revised the manuscript as follows:

Lines 392 to 393:” Aero-servo-elastic simulations are performed with NREL’s mid-fidelity wind turbine simulation software OpenFAST (NREL,2021) to compute the objective function vector f (Γd).”

Line 401 Revise sentence.

Author response: We fully agree with the reviewer. The revised text reads as follows:

Lines 417 to 419:” It follows that only by adding a proportional control gain (i.e. Kp,c) leads to more flexibility in reaching desired (Pareto) optimal solutions minimising torque fluctuations and corresponding (structural) loads, with a minimal impact on the power extraction performance.”

Caption Figure 8 add word the.

Author response: Thank you for your comment. We have revised the text of the caption accordingly.

Section 4: why is Kp,w of importance -> interpretation missing.

Author response: This is a valid point, and we agree with the reviewer that an interpretation of this variable is missing in the manuscript. Accordingly, the revised text reads as follows:

Lines 384 to 389:” Note that d = 5 refers to the original formulation of the WSE-TSR tracking controller, for which the integral term in the estimator Ki,w was introduced recently in the work of Liu et al. (2022). The integral term ensures that the internal estimated rotor speed state is consistent with the actual rotor speed measurement. Furthermore, combining a proportional and integral term (Kp,w and Ki,w) results in faster estimation convergence by rapidly reducing the estimation error. The input vectors Γd ⊂ Γ5 for d = {2, 3, 4} while Γ1 represents the one-dimensional design space of the Kw2 controller, in which the variation in λ∗ leads to variation in the gain K according to Equation 5. “

Caption Figure 9 State that these results are from the frequency domain model.

Author response: Thank you for your comment. However, these results are not obtained from the frequency domain model but from aero-servo-elastic simulations performed with the NREL’s mid-fidelity wind turbine simulation software OpenFAST. The details of the simulation settings are provided in lines 392 to 399. Specifically, the NREL 5MW reference wind turbine is subject to a realistic turbulent wind profile with a mean wind speed of 9 m/s at hub height and a turbulence intensity of 15%. Each simulation lasts 3600 seconds, of which the first 100 seconds are discarded to exclude the transient start-up effects from the results.

Furthermore, to improve clarity on this aspect, we have modified the Captions of Figure 7 and Figure 9 to explain how the results have been obtained.

Lines 451-452: Wouldn't this be even more the case for the controllers with larger bandwidth?

Author response: Thank you for your comment. The Pareto front obtained for the WSE-TSR tracking control scheme and related to the Γ5 design variables has two extremes: E (the point of minimum torque variance) and A (the point of maximum power extraction). As can be observed from Figure 9, there is a slight increase in power performance between points B and A at the expense of increased torque fluctuations. Therefore, having a controller with a larger bandwidth than A would not make sense since no power gain and increased torque fluctuations will be achieved. The resulting controller will be more aggressive, eventually likely causing instability in the system. We have clarified these aspects in the revised manuscript as:

Lines 509 to 512:" In this context, it is essential to consider that while a slight increase in power performance is observed for case A, it is accompanied by elevated torque fluctuations. Therefore, having a controller with a bandwidth exceeding that of case A would not be advantageous, as it would likely be more aggressive, potentially leading to system instability and yielding no power gain at the expense of increased torque fluctuations."

Section 5.2.2. explain what they mean, is it the complementary sensitivity?

Author response: Thank you for pointing this out. We have modified Section 5.2.2. to explain what the closed-loop transfer functions mean. Accordingly, the revised text reads as follows:

Lines 516 to 518:" For the different cases, Figure 11 illustrates the frequency responses of the transfer functions TΛ∗ →Λ(s) and TV→Λ(s), representing the closed-loop system performance in terms of reference tracking (complementary sensitivity) and disturbance rejection (sensitivity), respectively."

Reviewer 2

Major comments:

- This paper provides a detailed analysis of torque control tuning with a primary focus on the frequency domain. Specific guidance for tuning the torque control is not provided or any insight into secondary effects, like drivetrain loads and the ability to avoid tower frequencies.

Author response: Thank you for your comments on our paper.

We recognise the importance of providing tuning guidelines for the WSE-TSR tracking controller, as the scheme has five variables to tune: Kp,w, Ki,w, Kp,c, Ki,c and λ∗; there is not (yet) an established systematic way described in the literature of tuning the considered advanced WSE-TSR tracking control scheme. To provide the first steps towards such a systematic procedure, this paper presents a frequency-domain framework for the open-loop components and the closed-loop system by (complementary) sensitivity functions for the WSE-TSR tracking scheme. In this paper, the frequency-domain framework is used to evaluate the stability and performance characteristics (in terms of controller bandwidth) of the controlled system using fundamental control theory.

While the linear frequency-domain framework is well-suited for the abovementioned purposes, the performance requirements of present-day wind turbines become ever more complex, and the actual turbine performance cannot be directly related to the linear framework. This is why we, as authors, proposed the multi-objective optimisation approach to find the set of optimal solutions with the presumed conflicting objectives of power maximisation and torque fluctuation. The optimal set of solutions is subsequently assessed in terms of stability and controller bandwidth using the frequency-domain framework.

We also acknowledge the reviewer’s interest in overspeed exclusion zones and the avoidance of tower frequencies and appreciate the suggestion for future research directions. However, the intent of this paper is to fundamentally analyse and evaluate the possible performance benefits of a model-based advanced torque controller for the below-rated region. For this reason, these additional features are left out of the current work.

- Structural loads are mentioned but not presented anywhere. "Torque variation as a proxy for loads" could be made more specific. What structural loads are affected by the torque control tuning?

Author response: We appreciate the feedback, and in response to the reviewer’s request, we have performed additional simulations to comprehensively analyse the structural loads affected by the torque control tuning. The outcomes of these additional simulations have been integrated into the revised version of the manuscript. To offer readers a comprehensive view of our findings, we have introduced two new sections:

- Qualitative Assessments of Optimal Controller Solutions (Section 5.2): This section provides further insights into how turbine loading is influenced by a set of optimal solutions (controller calibration variables) on the Pareto frontier. The outcomes of this analysis are summarised in Table 2 in the revised manuscript.
- Sensitivity Analysis of Optimal Calibration Variables (Section 5.3): In this section, we conduct a sensitivity analysis of the optimal calibration variables on the turbine loading and the considered objectives. This analysis is centred around point C of the Pareto front and is selected as a representative trade-off between minimising generator torque fluctuations and maximising power production. Each of the five calibration variables is individually assessed for its positive or negative impact on turbine performance metrics. The results of this analysis are summarised in Table 3, providing readers with insights into the specific effects of varying these variables.

By introducing these new sections and tables, we aim to offer a detailed and meaningful understanding of the relationship between torque control calibration and structural loads.

We are confident that these additions will significantly enhance the quality and depth of our paper.

- The main conclusions of the authors are already known: both controllers track the optimal TSR, and the TSR tracking controller gives more flexibility; these are the design objectives of each controller. The conclusion that higher torque control (and WSE) gains lead to tighter control, with more torque variation and slightly more power output seems obvious. I'm not sure so much analysis is needed to prove that.

Author response: Thank you for your thoughtful comments, and we appreciate the opportunity to clarify certain aspects of our paper in response to your concerns.

Wind turbine technology has rapidly advanced, resulting in larger and more dynamically complex turbine designs encompassing a broader spectrum of performance requirements. The evolving dynamics of modern wind turbines have introduced complexities in the calibration of controllers, posing challenges when relying solely on traditional control theory. In light of this evolution, revisiting established methodologies to ensure their applicability to modern wind turbines is essential. The context within which existing conclusions were initially drawn, i.e. on smaller turbine systems, may no longer fully capture the intricated performance optimisation for today's larger turbines, such as the NREL 5 MW.

The <u>first contribution</u> of this paper is, therefore, to assess whether previous conclusions on power gains using the considered advanced torque control scheme still hold for modern multi-MW turbines. Contrary to conventional expectations, our findings significantly deviate from earlier literature (e.g., Bossanyi, 2000). While it was previously suggested that a manually calibrated WSE-TSR tracking controller could yield energy capture benefits of 1 to 3%, our study demonstrates that an optimally calibrated WSE-TSR tracking control strategy may not necessarily enhance power capture. Instead, it offers the advantage of reducing torque fluctuations, contributing to improved load minimisation. This distinction holds particular relevance when considering larger wind turbines, such as the NREL 5 MW, where the conclusions drawn from smaller turbine systems might not directly apply. Importantly, we want to clarify that **our paper does not conclude** that higher torque control gains lead to tighter control, with more torque variation and slightly more power output.

The <u>second contribution</u> is to address the increasing complexity of wind turbines and the limitations of conventional control theory. Our research takes a novel approach by adopting a multi-objective optimisation framework. This approach recognises that achieving optimal turbine performance extends beyond mere power capture, encompassing a broader range of objectives, including load minimisation and enhanced stability. By employing a multi-objective optimisation framework, we aim to achieve optimal calibration of the WSE-TSR tracking controller for realistic wind turbine sizes.

Despite our shift towards a multi-objective optimisation framework, frequency domain analysis remains an essential tool for evaluating the stability and performance of the closed-loop system in terms of controller bandwidth. The optimal solutions we observe are still assessed within this frequency domain framework. This enables us to relate control parameter insights (such as stability and controller bandwidth) to meaningful performance metrics (including power maximisation and load minimisation).

In summary, our paper provides a comprehensive and analytically driven exploration of the benefits and drawbacks of the WSE-TSR tracking controller, contradicting earlier assumptions and establishing a new perspective on performance optimisation for realistic wind turbine sizes. We hope this clarifies the objectives and conclusions of our research. However, to further clarify the aim of our paper and to take into account your feedback, we have revised the manuscript as follows:

Lines 90 to 93:" Therefore, applying a frequency-domain framework to evaluate the optimal solutions found by solving the multi-objective optimisation problem enables linking the conflicting control objectives with the stability and performance of the closed-loop system in terms of controller bandwidth."

- When is the torque control bandwidth "high enough", and there are diminishing returns in the power gains? How are drivetrain and tower loads affected? If the reference rotor speed is changed to avoid a natural frequency, how high does the torque control bandwidth need to be? I would have liked to learn the answer to these questions so that users of the WSE-TSR controller have specific guidance for how to tune them.

Author response: We appreciate your insightful questions and feedback. We are glad to address them point by point:

1) When is the torque control bandwidth "high enough", and there are diminishing returns in the power gains?
   You bring up a crucial point regarding the optimal torque control bandwidth and its relationship to power gains. We agree that deriving a direct conclusion from the frequency domain framework can be challenging. To address this, we employed a multi-objective optimisation approach. Through this method, we explored a range of optimal solutions that strike a balance between power extraction and torque variation. This balance, in turn, dictates the selection of a specific control bandwidth. By identifying these optimal points, we aimed to provide more comprehensive insights into the trade-offs between power extraction and control dynamics, enabling users of the WSE-TSR controller to make informed decisions on calibration. Therefore, to provide more clarity on this aspect, we have revised the manuscript as follows:

   Lines 555 to 559:" The resulting Pareto front approximations represent the optimal solutions and controller calibrations, providing a trade-off between the defined objectives and dictating the selection of specific controller bandwidth. A set of Pareto optimal solutions has been evaluated in the frequency and time domains to provide more comprehensive insights into the balance between performance metrics and control dynamics, enabling users of the WSE-TSR tracking control scheme to make informed decisions on its optimal calibration."

2) How are drivetrain and tower loads affected?
   We greatly appreciate your suggestion and have taken steps to enhance the clarity of load impact analysis. As you mentioned in a previous question, Table 2 has been included in the paper to give readers a concise overview of how torque control tuning influences tower and blade load components. This addition aims to give users a clear understanding of the intricate relationship between control parameters and load dynamics, facilitating their decision-making process.
   We acknowledge your question about drivetrain loads' importance in analysing the controller tuning. After thorough consideration and extensive simulations, we have determined that the behaviour of drivetrain loads closely aligns with the trends observed in torque variance. Given this congruence and to maintain a streamlined presentation, we have decided not to include separate drivetrain load analyses in the current paper.

3) If the reference rotor speed is changed to avoid a natural frequency, how high does the torque control bandwidth need to be?
   Your question about changing the reference rotor speed to avoid natural frequencies and its implications on the required torque control bandwidth is insightful. However, we want to clarify that our paper's scope does not encompass tower resonance avoidance through rotor speed adjustments. While this aspect is essential, it falls beyond the aims of this particular study. We appreciate your suggestion and acknowledge its value.

- This paper is on torque control, not general model-based wind turbine controllers as the title suggests.

Author response: Thank you for your suggestion, and we agree. Accordingly, we have changed the title of our paper to: ”Analysis and multi-objective optimisation of wind turbine torque control strategies”.

- For a wind energy science audience, there is a lot of mathematical jargon and nomenclature. To appeal to a wider audience, frame your problem in wind energy terms. Some examples are:

Author response: We appreciate your feedback and understand the importance of making our paper accessible to a broader audience. Regarding your specific points:

  - around L340, where MOO doesn't need to be explained in such mathematical detail
    One of the key objectives of our paper is to present a calibration framework for the WSE-TSR tracking controller using a multi-objective optimisation approach. A formalisation of the multi-objective optimisation problem is presented in Section 4.1 to facilitate understanding. This section provides a compact description of the MOO framework and thereby defines variables and terminology, as this will be used later in the results section.
    We have opted to retain the section as it is needed for the present paper. However, we have revised the section by removing the text and corresponding

equation that elaborates on the resultant minimisation cost function, thereby eliminating mathematical details associated with the multi-objective optimisation. We believe this revision addresses your concerns and ensures the paper remains comprehensible to a broader audience.

In L344, "operational conditions" are described. What exactly are they?

In response to your comment about the term “operational conditions” in Line 344, we clarified the meaning of operational conditions in Section 4.2 (lines 394 to 399), where the implementation of the multi-objective optimisation framework is described. Specifically, the NREL 5MW reference wind turbine is subject to a realistic turbulent wind profile with a mean wind speed of 9 m/s at hub height and a turbulence intensity of 15%. Each simulation lasts 3600 seconds, of which the first 100 seconds are discarded to exclude the transient start-up effects from the results. The resulting time series are then used to compute the considered objective functions: $f1(\Gamma d)$ and $f2(\Gamma d)$.

- I would have liked to see power variance as the cost or perhaps structural loads; this data is available in the simulations.

Author response: Thank you for your comment.

We appreciate your suggestion to include power variance as a cost measure in our analysis. We have considered your feedback and made further analysis in response to your suggestion. Below, you can find a plot where power variance is presented as the first optimisation objective (replacing torque variance).
The results obtained from using torque variance have proven to yield comparable insights to those derived from power variance. Given that both variables provide similar outcomes, we believe that maintaining the consistency of using torque variance throughout the paper enhances the clarity and cohesiveness of our analysis. Therefore, we have decided to continue using torque variance as the primary variable for our analysis.

Regarding your inquiry about incorporating turbine load as an additional objective, we want to clarify that while the presented multi-objective optimisation framework does have the capability to include load objectives, this aspect is devoted to future work. In this paper, we focused on the current power and torque variance objectives.

We hope these explanations address your concerns and highlight our rationale for the choices made. Your feedback has been invaluable in refining the quality of our work, and we are grateful for the opportunity to improve upon it based on your suggestions.

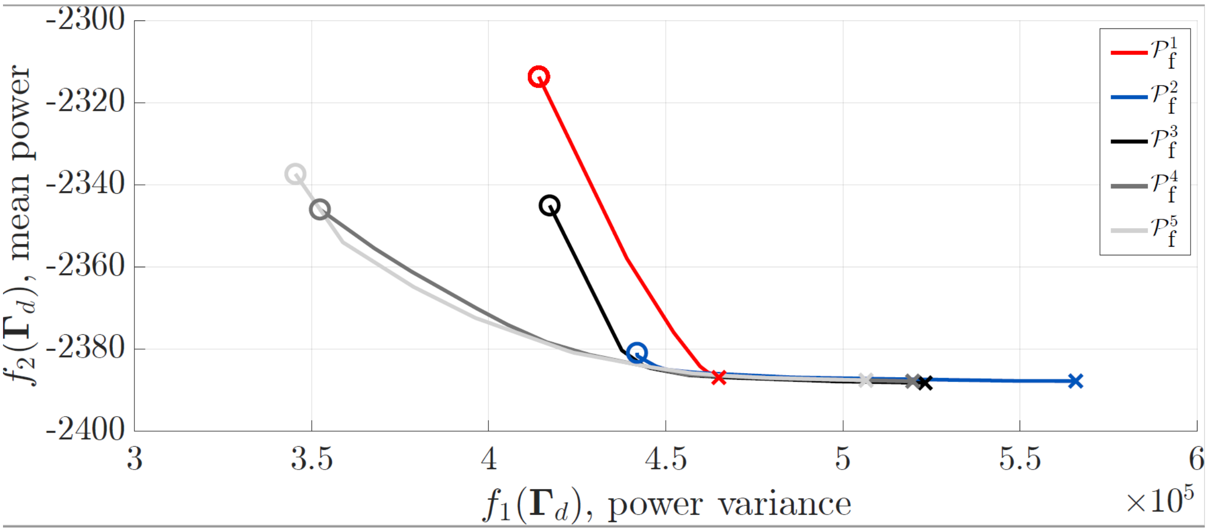

Figure - Pareto fronts obtained for the WSE-TSR tracking control scheme for different sets of estimator-controller design variables: Γ1, Γ2,Γ3, Γ4 and Γ5. Simulations are performed with the NREL's mid-fidelity wind turbine simulation software OpenFAST (NREL,2021) under realistic turbulent wind conditions. The objective functions f1(Γd), i.e. power fluctuations minimisation, and f2(Γd), i.e. power maximisation, define the performance space for the controller. The optimal solutions for f1(Γd) and f2(Γd) are indicated using circles (◦) and crosses (×), respectively.

Minor comments:

- Abstract: Quantitative results would be appropriate here.

Author response: Thank you for your suggestion. We have modified the abstract accordingly. The revised text reads as follows:

Lines 12 to 13: "By lowering the controller bandwidth at the expense of generated power of 2%, the torque actuation effort reduces by 80% with respect to the optimal calibration corresponding to the highest control bandwidth."

- Abstract: does OpenFAST count as high-fidelity simulations? Probably not, so perhaps be specific.

Author response: We agree with the reviewer. We have revised the manuscript by changing "high-fidelity simulations" to "mid-fidelity simulations".

- Intro: L29: k\omega^2 may be common in research articles, but probably not on actual turbines

Author response: Thank you for your comment. We completely understand your point and acknowledge that the komega2 controller may be more commonly discussed in research articles rather than directly implemented in real wind turbines. Therefore, we have changed the manuscript accordingly. The revised text reads as follows:

Lines 28 to 30: "Nowadays, the Kw2 controller is still a commonly considered partial load region wind turbine torque control strategy due to its satisfactory performance, ease of derivation, and simple implementation by only requiring a measurement of the rotor or generator speed (Johnson et al., 2006; Odzemir et al., 2013)."

- Intro: L50: Is this always the case or just for that rotor in that study?

Author response: Thank you for your comment. The 0.5% increase in captured power is observed for the CART rotor by Johnson et al. (2004). However, the authors of that study suggest that the gain reduction strategy can be implemented on any existing wind turbine, providing improved energy capture. No linear correlation is found between the gain reduction factor and the site condition. Therefore, the percentage of increased captured power will depend on the turbulent conditions and the turbine. To resolve any ambiguities, we have revised this part in the introduction.

- Section 2: Theory: L149: constant \lambda_* or single \lambda_*?

Author response: Thank you for pointing this out. We have modified the corresponding line. Accordingly, the revised text reads as follows:

Lines 156 to 157: "It can be observed that a single $\lambda*$ exists, which corresponds to the rotor operating point for maximum power extraction efficiency $C_{p,*}(\lambda*)$."

- Fig 6: it's difficult to tell the black lines apart

Author response: Thank you for your comment. We agree with the reviewer that Figure 6 was unclear. Accordingly, Figure 6 has been modified to increase the contrast of the presented results.

- Are the controllers with a high bandwidth stable? The resonance would lead me to believe they are not. Are there stability margins associated with these controllers? What is the significance of the resonant peak? It seems like it would degrade performance if excited.

Author response: Thank you for your comment. We appreciate your feedback and understand your concerns regarding the potential negative impact of resonance on the system's performance. We have conducted further analysis to investigate this aspect, aiming to grasp the motivation behind the resonance peaks for the controller with higher bandwidth. To this end, we have summarised our findings in the figures below, where we present the Nyquist plot of the controller transfer function $K_{\Omega r \rightarrow Tg}(s)$ multiplied with the wind turbine plant (open loop gain).

Regarding the Nyquist plot, we can see that the system is stable because the Nyquist plots do not encircle point -1, and there are no poles of the open-loop system on the right-half plane. Furthermore, the gain and phase margins for cases A and B, indicated with GmA PmA, GmB and PmB, respectively, are:

- GmA = 3.55 PmA = 66.75;
- GmB = 2.83 PmB = 114.88;

which confirms the stability of the system with these controllers in closed-loop. Upon assessing these results, we can now confidently assert that the resonance peaks observed in cases A and B do not cause instability of the closed-loop system. Rather, they contribute to increasing the bandwidth of the corresponding controllers. This outcome reinforces the synergy between a multi-objective optimisation framework and frequency analysis. This synergy bridges the stability of the controllers with the turbine performance metrics, namely power maximisation and load minimisation.

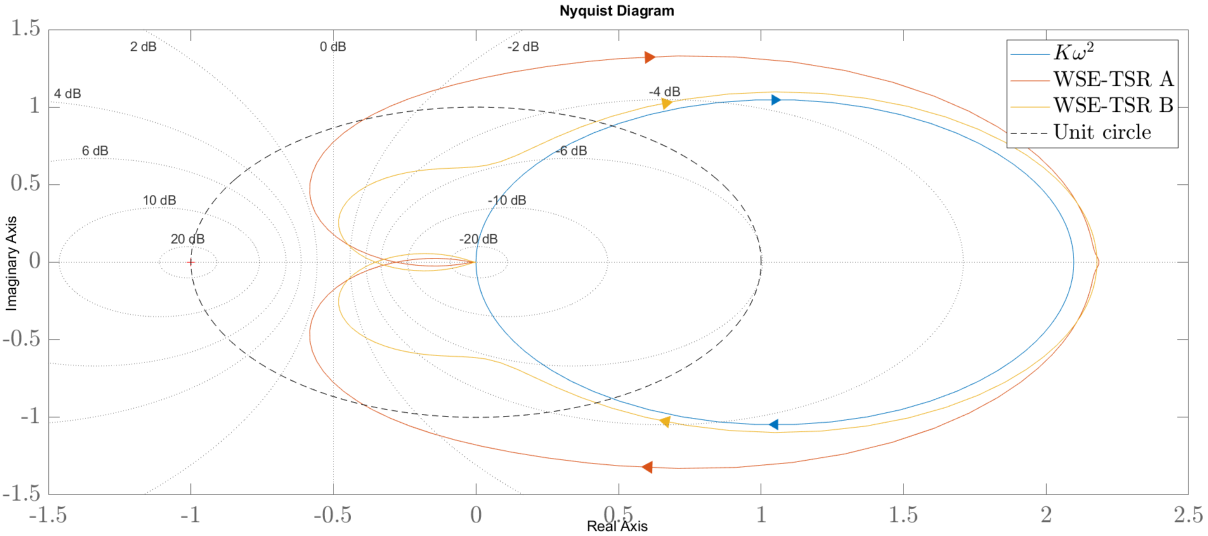

Figure - Nyquist plot