2nd Review of “Scaled testing of maximum-reserve active power control”

Simone Tamaro, Davide Bortolin, Filippo Campagnolo, Franz V. Muehle, and Carlo L. Bottasso

Review date: 02/04/2026

I would like to thank the authors for considering our comments and revising the paper accordingly. All comments have been sufficiently addressed, and in my opinion, the paper can be accepted.

However, I would recommend to elaborate more on the following points (the numbers refer to my questions from the first review):

1. The authors now clearly distinguish between two different aspects of the proposed controller: 1) increasing the available power reserve, and 2) better tracking by maximising the turbines' distance from the greedy operating point, i.e. from the saturation. However, the newly added text contains a small inaccuracy. In line 49 of the revised manuscript with highlighted changes, the authors state that "none of the APC methods developed so far directly accounts for the power reserve available at each turbine". There were some works on model predictive active power control, where the available control actions were implemented directly in the controller formulation -- turbines were constrained not to operate beyond the greedy operating point. Since MPC directly takes constraints into account, it means that the available powers at each turbine were considered by the controller. Some examples include e.g.:

Kazda et al., doi: 10.1109/TCST.2019.2923779

Vali et al., doi: 10.1109/TCST.2021.3053776

Hence, I ask the authors to choose a more precise wording.

4. Thanks to the additional explanation of Fig. 12, it is now much easier to understand the figure. However, there seem to be four different line colours used in the figure, which are not explained. I would suggest either explaining the colours or using a single one. I assume that each of the turbines got a different colour. As for the minimum value (the right plots), the CL+MR uses the same colour as WT2, but the OL uses a slightly different shade of blue, indicating that it is mostly, but not entirely based on WT1. An additional explanation or a legend would lead to a better understanding of the figure. If this is not deemed important, it would be fine to abandon the colour coding of the lines.

5. Thanks to the authors for offering plausible explanations for the discrepancies between their results and the ones from Vali et al., 2019. Another difference that has caught my attention when looking at the provided supplementary files is how the load-reducing loop gets deactivated when a turbine gets saturated. In the current implementation, the error signals is set to zero in CLD LOOP 1/2/3, which might lead to sudden jumps in alpha (i.e. in alfa1, alfa2, alfa3 in the Simulink model). When the turbine gets out of saturation, this might also lead to a jump in alpha. Additionally, since only the loads from non-saturated turbines are considered in the controller,

this might lead to a sudden change in the reference load signal for the non-saturated turbines whenever a turbine gets in or out of saturation. It is my understanding that this approach is slightly different than in Vali et al., 2019. Have you detected sudden changes in alpha, especially for the turbines going in and out of saturation, and could this also contribute to the increased loads?

Just to be clear: I do not expect the authors to perform another wind tunnel campaign or to run a series of simulations getting to the bottom of this. But a quick look into the measurement data could possibly offer another explanation for the increased loads.

8. The original question 8 referred to the transient response of the turbine to changing yaw or pitch angle, i.e. how fast rotational speed and power react to these changes and how that response time compares to upscaled real-world turbines. This is a relevant aspect when validating the capability of a wind farm to provide system services.

It would be interesting to hear about this in the paper.