WES-2021-119

Lidar-assisted model predictive control of wind turbine fatigue via online rainflow-counting considering stress history

Stefan Loew and Carlo L. Bottasso

General Comments

This is a well written paper about the author's contributions to the field of model predictive control for wind turbines. Extensive detail is given to the problem formulation and setup, where fatigue calculations, model predictive control, and lidar-based estimation are all described very thoroughly. This PORFC algorithm which uses the lidar measurements and employs online rainflow counting is shown to produce favorable results in a comprehensive study comparing the proposed controllers against the standard NREL 5MW control algorithm.

Generally, I think this is a well written, novel, and thorough presentation of the completed work.

Detailed Comments

- I am curious as to why a moving horizon estimator was used instead of a Kalman filter. At least some commentary on this would be helpful, especially given the estimator offsets seen in Figure 9.

- I agree with the other reviewer's comments regarding using "profit" as a primary metric of interest. In my opinion, I don't think a major revision is necessary, rather a more explicit acknowledgement of the blanket nature of a metric like this might be more useful. I think it is fine to propose an objective function that subtracts some estimated fatigue cost from revenue, but I think it is important to discuss the limitations of such a metric. In presenting your results and conclusions, it is important to emphasize that the novelty of this work lies in how you are including online fatigue estimations in MPC, and that you are simply using the proposed cost function to demonstrate it. The novelty is not in the cost function in and of itself, and thus, the function could be changed accordingly. You talk a little bit in the last section about how the cost function could be modified for "different business cases", but this is only suggested as a possibility late in the paper.

A Minor Comment/Gripe

- I am not particularly keen on comparisons of controllers that have some inherent optimality/objective (like MPC) to highly generalized (and old) reference controllers. The results of this paper suggest a ~20% profit increase over the NREL5MW controller, but the NREL 5MW controller has not been tuned or optimized with any explicit consideration of fatigue, let alone "profitability". In an ideal world, it would be nice to

see the novel MPC algorithm compared to a reference controller that *has* been explicitly re-tuned to address the same, or similar, objectives.

In an even more ideal world, the reference controller would be more “modern” and have features such as thrust limiting, as this would provide a more realistic basis of what sort of advantages the MPC controller provides over more modern standards (my, obviously biased, opinion is that NREL’s ROSCO controller is a good candidate for this sort of study, though a re-tuned in-house or modified NREL 5MW controller could be sufficient).

That said, I understand that it may be unrealistic to re-tune and re-optimize a reference controller for each project. I don’t think this is necessary for the work presented in this article, but I do think the rudimentary nature of the NREL 5MW controller deserves some mention in this paper. And of course, I hope you keep these sentiments in mind in your work going forward.