Review - WES-2021-119

General assessment

The paper provides a comprehensive study of lidar-based Model Predictive Control using online fatigue damage estimation and is deemed of high scientific value. Lidar-based wind speed estimation, MPC of wind turbines and fatigue are introduced in a pedagogical manner. Then, the novel lidar- and fatigue damage-based algorithm PORFC is presented in detail, including a thorough description of the tuning procedure. A comprehensive cross-comparison of control algorithms is then performed showing the potential of the approach.

Specific comments

- A request for revision lies in how the overall uncertainty in the fatigue cost model is considered. Although it is understood that using fatigue damage directly brings us one step closer to monetary costs compared with indirect fatigue damage metrics, the concept of "profit" directly subtracting fatigue cost from revenue as main quantitative key performance indicator is inappropriate. The uncertainty lying in this last step impairs the excellent scientific value of the paper. In general, the paper should be kept neutral and objective, not trying to subjectively sell the method.

  More into detail, reducing the cost to CAPEX of one single component is overly simplifying the problem. Fatigue costs involve holistic system-level analysis over the entire lifetime of the turbine. For instance, it is expected that O&M costs on drivetrain and pitch actuation system components would change the conclusions dramatically (the term "worth investigating" used in conclusion is a mild statement). As the overall gain in profit is not so evident in terms of order of magnitude, suggesting a reduction in power production in order to increase profit by decreasing fatigue should be taken with caution.

  However, this can be fixed without major revision through targeted changes in writing style keeping objectiveness in mind and emphasizing on revenue and fatigue costs separately, rather than on profit. The latter may still be kept for illustration, but the rough underlying assumptions should be clearly mentioned. The promising potential of the method should rather be promoted through control algorithms that do not (or not much) reduce revenue (in other words, revenue and fatigue costs should be weighted unevenly, ideally with varying weight to encompass multiple scenarios and handle the aforementioned uncertainty).
- A relation between the Moving Horizon Estimator and the much more standard Kalman filter would be welcome, both being optimal state estimators

- Even though OpenFAST has a much higher fidelity than most control-oriented models, it should not be presented as high-fidelity model. It is an engineering model, commonly defined as mid-fidelity, high-fidelity being reserved for CFD.

- Paragraph 5.2.2/Filtering: the IIR filter used behind the filtfilt function has not been specified (a Butterworth filter of order 2 would be typical for this purpose). The moving average filter may perform better than the particular IIR filter used here, but it would be surprising if this can be generalized, so the comparison should probably not be presented here. This said, zero-phase IIR filters are indeed inappropriate for online use on very short records, and the horizon length is an easier tuning parameter to work with than a cutoff frequency in this case, so the moving average filter is a good choice.

Typographical comments

Paragraph 5.1.6: missing right parenthesis after "based on a realistic piecewise S-N-curve"

Paragraph 5.2.2: should it be "if the MPC internal and the plant model are NOT matching"?