Response to reviewer 1

This paper is an excellent piece of research, full of innovative ideas and techniques and very well written. It should definitely be published in WES. My suggestions below are not even minor revisions, but addressing them will clarify a few issues and hopefully improve readability.

I have one somewhat sad comment. To be honest, I was expecting a better performance of the entire modeling system and perhaps more wake analyses (only two turbine wakes were analyzed, qualitatively). I guess I should not be surprised because the data are never what we hope them to be. There were so many mismatches between the real wind farm and the simulated one, from pitch angles to yaw misalignment etc. But unfortunately, because of such data issues, I am not convinced of the goodness of either the SP-Wind solver with the AASM or the initialization technique using TotalControl Flow data. This is really sad because both are innovative and appear to be robust. Maybe a more in-depth data cleaning procedure could be applied to the dataset?

We thank the reviewer for the kind words and appreciation of the manuscript. The main goal of the current work was to highlight the challenges associated with detailed aeroelastic simulations of wind farms and comparison against a rich field data set which includes SCADA, LiDAR and Loads comparisons. Through the detailed comparison, we were able to highlight the areas of numerical modelling which still require improvement for accurate representation of wind farm performance. While some data cleaning was conducted to remove channels with high noise or missing data, it was not clearly evident that the remaining errors are due to errors in measurements or a real feature due to the differences between our modelling approach and the real conditions at the wind farm.

Please find below responses to the reviewers point comments, which we believe have improved the quality and clarity of the manuscript.

1. Either “dataset” or “data set” should be used, but not “data-set”. Please replace it throughout.

This suggestion has be incorporated in the manuscript and data set is now used consistently throughout.

2. Line 20: you should expand your literature review because citing only two studies about LES applications to wind farms is just not acceptable. There are so many more such studies, including several about Lillgrund. Here is an incomplete list:

Archer et al. (2013), Bhaganagar and Debnath (2015), Calaf et al. (2010, 2011), Chaudhari et al. (2017), Churchfield et al. (2012), Fleming et al. (2014), Ghaisas and Archer (2016), Ghaisas et al. (2017), Han et al. (2016), Lu and Porte-Agel (2011), Martínez-Tossas et al. (2015), Meyers and Meneveau (2011), Xie and Archer (2017)

We thank the author for a comprehensive list of references, and have now included additional references in the text for LES and aerodynamics of wind farms.

3. Lines 22-23: you should not list a series of issues that LES has apparently helped address without citing the literature studies that you think did it. You need to add at least one citation for each of these terms: gusts, atmospheric stratification, and turbine-wake interactions. I assume that Mehta at al. (2014) was about local wind climate only.

The paper by Mehta et al (2014) provides a literature review of aerodynamics of wind farm LES, which further expands upon the mentioned effects. As the issues are not explored in our work and for the sake of brevity, we chose to include a reference for the review for completeness. We have now also included an additional review paper (Porte-agel 2020) for wind farm flows.

4. Around line 190: What exactly is a “Pressure-driven” boundary layer? Is it neutral? It seems that there is no thermal inversion in those. What is a “conventionally neutral” BL? Please add some more description about the two datasets used (PDBL and CNBL).

A PDBL is a neutral boundary layer which ignores the Coriolis effects and free atmosphere stratification, whereas a CNBL includes these effects and is characterized by a neutral boundary layer capped with stable free atmosphere. Both the PDBL and CNBL are neutral in the surface layer. This has been clarified in the text in section 4.1 in the following lines, and a reference has been included which details the generation of the precursor datasets.

“ The TotalControl flow fields are initialized using mean velocity profiles upon which random divergence-free perturbations are added. These initial conditions are then advanced in time for 20 physical hours, so that the influence of the unphysical perturbations has disappeared, and the flow has reached a fully turbulent and statistically stationary state. The PDBL is a simple representation of a neutral atmospheric boundary layer which ignores the effects of rotation and thermal stratification, while CNBL provides a more realistic representation by including these effects. Both these boundary layer types have been extensively used in wind-farm LES.“

5. Related to question 2 above, it would really help is the same information could be provided for all 5 cases (PDBL and CNBL). In Figure 6, it would be great if you could add the same profiles but for CNBL. In Figure 7, it would be even better if you could add the profiles for the PDBL, especially the potential temperature profiles.

Both the boundary layers in the dataset are fundamentally different, hence it was not possible to have similar figures for all the 5 cases. This is reason why the 2 PDBL were grouped together in Figure 6, and the 3 CNBL in Figure 7. The PDBL are unidirectional and do not include temperature effects, therefore it is not possible to plot veer and potential temperature effects as was done for the CNBL profiles.

6. Figure 6a: is the value of z0 reported for PDkhi correct? It is inconsistent with the value in Table 1.

This was an inconsistency and the value in Table 1 has now been corrected.

7. Lines 200-205: It is brilliant to re-scale the precursor runs from Total Control Flow via appropriate z0 and u*. In the text, it sounds like u* is “imposed” in the Total Control Flow fields (together with z0). Is that true? Can u* be imposed during initialization? I am just curious. If so, please add the values of imposed u* in Table 1 in another column.

The following lines have been added in section 4.1 to clarify these points

“ All the boundary layers have been intialized using a friction velocity of u_*=0.28 m/s , which is a typical value for marine boundary layers. For the PDBL, this requires a driving pressure gradient of $\nabla p_\infty/\rho =$ -5.2267×10^{-5} m/s^2 for a boundary layer height of 1500 m according to the equation, $u_\star = \left(-\frac{H}{\rho} * \nabla \mathrm{p}_\infty\right)^{\frac{1}{2}}$ ”

The pressure gradient in the Lillgrund simulations is adapted for the new values of friction velocities required for the transformation. As all the flow fields have the same friction velocity, we include this information as a line in the text instead of another column in the table.

8. (5): what is vector e1?

Vector e1 represents a unit vector in the x direction, symbolizing that the off-set to the mean flow through a new surface roughness is applied only to the u velocity. This has now been clarified.

9. Table 1: Is z0 for case PDkhi really 2 x 10^-5 m? In Figure 6a, it appears to be 2 x 10^-3 m. There is an inconsistency.

This was an inconsistency. z0 for the case PDki is 2 x 10^{-3} m, and has now been fixed in the table.

10. Page 13: This idea is really cool, but the notation is tough. For example, u and v are usually the two horizontal components of the wind vector. Here, they are used to indicate wind speed, but u for LES and v for Lidar. A bit confusing. X and Y are both wind speeds, but usually x and y are Cartesian coordinates. The overbar indicates what exactly? A 75-minute time average at each point? Or is there some sort of horizontal average first to obtain "time average profiles at range gate locations", which are then weighted with w? Why is it so important to minimize the covariance distance? Given the relatively large errors in wind direction (Figure 8b, especially PDk3), I was thinking that maybe a wind direction distance could be minimized instead? And finally why is the simple sum of the two distances chosen? Isn't wind speed more important? These are all requests to add some clarifications in the text.

We thank the reviewer for the suggestions and have reworded this section to improve clarity. X and Y are now instead replaced with P and Q to avoid confusion. The overbar indicates a 75 minute time average. Before averaging, data is extracted from the LES domain at points equivalent to the LiDAR measurement range gate locations. Along with the mean velocity profiles, the covariance distance is also minimized to aid in the selection of transformation parameters which gave similar turbulence intensity as the LiDAR measurements. Without including the covariance distance, we observed that while the mean profile could match well, the turbulence intensity can have large differences. Figure 8b does not show differences in wind direction, but in wind veer across the rotor area. All the PDBL simulations have no spanwise velocities, hence zero veer. This emphasizes that the PDBL is not a realistic representation of the atmospheric boundary layer, as can be seen in figure 8b where the field measurements have significant veer across the rotor area. Since the optimization framework minimizes the distances between both the u and v velocity components, the wind direction is implicitly minimized as well and does not need to be including separately. We agree that choosing a direct sum might be too simplistic, and in future work a parametric study could be conducted to determine which metric has more significance when determining similarity. These points have now been clarified in the text in more detail.

11. Table 3 is not cited anywhere in the manuscript. Perhaps at line 253?

A reference to Table 3 is now included in section 4.2

12. Line 2054: This domain is huge! Why was such a large domain chosen? I thought that perhaps that was the domain size in the PDBL and CNBL datasets (if so, please mention it). If not, then how do you "extract" a portion of the data from those datasets to match your domain?

The domain size in the current simulations is indeed restricted by the domain size of the precursor datasets, which was originally designed for a larger wind farm and this has now been clarified in the text in the following lines in section 5.1

"The choice of domain size is restricted by the one used for the precursor data sets, which was initially designed for simulations with the much larger TotalControl reference wind farm to avoid blockage effects"

13. Line 277: this phrase does not make sense … What does "made to advance" mean? Who advances what? What do you mean by "the inflows"? I thought you took initial and boundary conditions from TotalControl Flow, after correcting them with the scaling parameters. Please rewrite to clarify.

These lines aim to explain how the concurrent precursor methodology of windfarm LES has been modified to include flow transformation for representing the inflow conditions at the Lillgrund wind farm. The entire three-dimensional instantaneous flow field from the TotalControl precursor database is used as the initial starting conditions for a precursor simulation, which are then advanced in time using the Runge Kutta scheme. A slab is taken from the precursor domain (without turbines), transformed by the scaling parameters to match the inflow conditions at the Lillgrund wind farm, and added to a second domain which contains wind turbines through a fringe region. This has been now clarified in the text in section 5.1.

"First, the inflows from the previously generated TotalControl precursor database are made to advance in time in a domain without wind turbines, called the precursor domain. Concurrently, the flow is transformed using the identified transformation parameters through the optimization framework to obtain the five cases identified in Table 2 and fed into a second domain which contains wind turbines as body forces F_R through a fringe region"

14. Line 301: It sounds like you should have cleaned up the SCADA data and removed high-pitch cases. Why did you not do it?

It was not directly clear whether the high pitch cases were a result of corrupted data, or a feature of the controller operational on the Lillgrund turbines. Hence it was not straightforward to remove these cases. As the pitch values were not unrealistically high, it could be that the Lillgrund turbines exhibit some degree of pitching in region 2 as well, which was not replicated in our controller implementation. In hindsight, it would've been better to screen the data to learn about such issues in advance and this can be incorporated into future workflow.

15. Figures 16 and 17: I assume that the orange lines correspond to the location of the vertical profiles. Did you also average along those lines or did you pick the wind speed at the exact points downstream at 0.5D, 1D etc?

Figures 16 and 17 represent wake deficit in the horizontal plane behind the turbine at different downstream locations, and not the vertical profiles. The horizontal profiles were compared at a cross section spanning 2 rotor diameters at the same points behind the rotor in the LES results and the LiDAR wake measurements.