Response to Reviewer #2

We appreciate the constructive and insightful comments from the reviewer. Detailed replies to each of the reviewer's points (in blue italics) are provided below. Given the opportunity, we feel that incorporating the reviewer's comments will produce a substantially clearer and stronger paper as a result.

General Comments

The paper deals with an important subject, e.g. the maximal power output of a wind plant by optimal placing of the turbines. The combination of RANS and adjoint is nowadays mainly used in shape optimization, so this application shows an approach that is not so widely used yet.

The chosen test cases and final case are well suited to build up the understanding of the optimization process, but I found some sketchiness of the analysis. So I would mainly suggest a more careful analysis of the results.

But overall, the paper is well written and for the most part well structured.

Specific Comments

section 2.2

Generally you talk often about something like realistic or fully turbulent flows. But in RANS it's still "just" a snapshot of the averaged flow, perhaps use steady-state sometimes instead of turbulent flow field. Reading this evoked a picture of a fluctuating flow field, which is not true.

> We agree with the reviewer that it is important to convey that our simulations are steady-state, and we will provide additional references to this fact in the abstract, introduction, and conclusions.

second sentence: the information that cheap computational cost allows gradient-free methods implies that gradient methods are suitable for high computational cost. But that's only the case for low amount of design variables or very few methods. A little note on the limits would be nice.

> We agree that gradient-based methods are suitable for optimizing high computation cost functionals because minima can be found with orders of magnitude fewer functional evaluations than a gradient-free method. It is also true that certain gradient methods may be limited to a relatively small number of control variables if the cost of calculating the gradient scales with the number of control variables, such as in a finite difference approach. However, a core point of our paper is that the adjoint approach finds gradients at a cost that is independent of the number of control variables, allowing for gradient-based optimization of high-dimensional optimization problems with many controls variables. In order to avoid confusion, we will more clearly make a distinction in this section between other gradient-based methods and the present adjoint approach to finding gradients.

last sentence of first paragraph: Do you expect improvements regarding the found optimum or regarding the computational effort of the optimization process?

> We expect that using a higher fidelity flow model would result in a different optimal layout than optimizing with a linear wake model because the optimization algorithm would see different representations of the wake and underlying flow physics. We are preparing a follow-up paper

comparing optimized layouts obtained from our model and an engineering wake model when both layouts are simulated in an LES model. We will make mention of this in our future work plans at the end of the paper.

section 2.3

generally: how do constraints and the amount of constraints influence the effort of the adjoint approach? As a well-posed optimization problem is crucial, a sentence on this would be nice. This could also be mentioned when introducing the optimization problem in section 3.1, as you could directly refer to your case.

The most challenging constraints are related to the inter-turbine minimum spacing as the number of these constraints grows exponentially and requires the use of the SLSQP optimization algorithm instead of a bounded BFGS algorithm. Following the reviewer's recommendation, we will expand our discussion of the constraints and their impact on selecting an optimization algorithm in Sections 2.3 and 3.1

Eq. 5: Usually this equation is given with other derivation (du/dm instead of partial operators, see: Giles and Pierce 2000) operators. Depending on the definition of u and m, this could be used differently. Perhaps add one explanation here.

Based on this comment and a comment from another reviewer we will reformulate the adjoint derivation in terms of reduced functionals that should avoid any confusion.

after eq. 9: Can you give a reason for your statement regarding the accuracy of gradients?

Given that the accuracy of the gradients depends heavily on the specific calculation method (e.g., 2nd vs. 4th order finite differences), we have chosen to remove this sentence in order to avoid ambiguity.

section 3.3

Beginning: From my knowledge of wind farm simulations, the grid dimensions seem to be quite small. Was there a grid study or former set-up validation of the simulations you could refer to?

We agree that the grid dimensions in this study could be increased, but our primary interest here is in introducing the adjoint optimization and flow solver tool. For this we chose a more 'manageable' grid size in order to explore a wider range of cases. In the future, we will implement and test various improvements to the tool, including more sophisticated RANS models, complex terrain, and more grid points. We will expand our discussion of future work in the conclusions to point out that larger problem sizes are an important area of development.

Is " 'do nothing' boundary condition" perhaps better defined by saying "Neumann boundary condition"?

The 'do nothing' boundary condition is the typical outflow boundary condition in finite element simulations of incompressible flows where the boundary integral terms arising from integration by parts of the variational form are set to 0, i.e., $\nu\partial_n u - pn = 0$ on the outflow boundary. For readers seeking more detail on this boundary condition, we will add a reference to Heywood, Rannacher, and Turek (Int. J. of Num. Meth. in Fluids, 1996).

last paragraph: I am not quite sure when looking at your flowchart in the next section when you are calculating your adjoints. Only on base of the averaged flow field or for each of the K wind states? Because if you could compute your complete gradients based on the very last, averaged step, this would mean an extraordinary reduction of complexity and computational effort of this approach.

The adjoint is calculated after solving for each of the flow states and calculating a total power. This total summation involves the flow field solutions from each state, and as a result a corresponding adjoint simulation is required for each one. In order to make this more clear we will add a sentence further explaining this point.

section 3.4

second sentence after the numbering: you state that less optimization iterations are necessary for multiple inflow states. This is not intuitive on the first look, so perhaps move the sentence about the more convex problem upwards in order to give an explanation.

We will reorder the sentences accordingly.

You mention line search without naming the optimization algorithm you will use. Perhaps move the description (or a shorter version of this part) of SLSQP from the end of section 3.5 to this section.

We appreciate this suggestion and we will move the discussion of optimization algorithms from Section 3.5 up to Section 3.4.

section 4.1

Plot: The speed-up is very difficult to see on the lower two plots. Perhaps adding the mean velocity in thin lines could be helpful to show where the current velocity is higher or below that value. Another idea would be to mark the areas of increased velocity, for example with shaded backgrounds.

In order to clarify this figure, we will reduce the number of profiles shown in the bottom two panels and will also add vertical dashed lines to each profile showing the baseline value of velocity for comparison. We also wish to point out that this figure was not intended to demonstrate the speedup so much as to demonstrate correct wake and boundary layer profiles. As the reviewer correctly notes in a later comment, the speedup effects are more apparent between two adjacent turbines than around just a single turbine.

This section is the part where the flow field, that is the base of the optimization, is investigated. Therefore, the statements are a bit weak. Can you either give a quantifying number of the error or cite some publication where this flow solver set-up is validated?

In this section our intent was to show that the flow fields produced are qualitatively consistent with prior results. Moreover, the primary purpose of this paper is to introduce the combined adjoint and flow solver optimization tool. In reality, a wide range of different RANS models could be used with this tool, and exploring different models is a direction for future work. In order to make the purpose of this section more clear, we will note at the beginning and end of the section that the purpose of these flow solver tests is simply to demonstrate qualitative agreement with prior results and that a wide range of other flow solvers and RANS models could be implemented in the tool.

section 4.2

You justify the staggered layout with the speed-ups. But the speed-up only appears very close to the inducing turbine. Or do you talk about something like jets between two neighbouring turbines? Also in figure 5 this becomes not very clear to me, as the maximal speed-up is less than 5% and only visible right behind the rotor disc. If you talk about two different speed-ups a clear wording has to be found.

> The reviewer is correct that the speedup effect we refer to might be better described as a jet between two turbines. The speedup effect is enhanced by the presence of two turbines and is less noticeable for a single turbine in isolation. We will clarify our description of the speedup effect accordingly. The power produced by the turbines scales with the cube of the windspeed, so even slight windspeed increases can produce substantial power increases.

Figure 6: The first plot for two wind directions shows a complex layout. I would have guessed either a fully occupied front row of turbines at the north and south boundary of the farm and then some turbines in between with maximal spacing in between in order to reduce the impact of wakes. Or if the speed-up effect is high enough, a pair-wise placing of the turbines along the diagonal from lower left to upper right corner of the farm boundaries, because with this layout the turbines either see the full wind or the speed-up wind from the one in front. Did you check these layouts? If so, how do they perform? How do the turbines move? Because the shown displacement of some turbines seems to be quite big, what is a bit surprising for a local optimum search. Perhaps a movement trajectory of some representative turbines helps to understand the final layouts. Discussing this most easy case more in detail also helps to understand and analyse the other results.

> In response to this comment and one from another reviewer, we will add a reference case with two parallel turbine rows and discuss improvements with the optimal layout.

Technical Comments

Figure 2: Axis labels and numbers are quite small

> We will increase the font size.

In the figures where you show velocity fields it can be helpful to reduce the amount of colour steps (eg. between red and blue). That way differences (and speed-ups) can become more obvious.

> We appreciate this suggestions and we will try reducing the number of discrete colors to see if the effects become more apparent.

Sincerely, the authors.