2018 Annual American Control Conference (ACC)
June 27–29, 2018. Wisconsin Center, Milwaukee, USA

Online model updating by a wake detector for wind farm control

C.L. Bottasso, J. Schreiber

***Abstract*— An engineering wake model is updated online based on measurements provided by shaded turbines. Departing from other approaches, the measurements include information on the impinging wakes, obtained by a wake detector based on measured rotor loads. The updated model exhibits improved prediction capabilities, and it can be used for implementing a model-based wind farm controller.**

I. Introduction

Each turbine in a wind farm emits a wake characterized by reduced velocity and increased turbulence, leading to losses in power production and increased loads on downwind turbines. The negative effects of wake interactions may be mitigated by wake management strategies [1]. One possible implementation of such strategies is based on a wind farm flow model: the predictions of the model are used by a controller, whose aim is to energize and/or redirect wakes for improved energy yield and/or reduced loading.

The performance of any such model-based control method is *inherently limited by the accuracy of its underlying model*. Unfortunately, any model has limitations –at least in some situations–, and especially the simple reduced-order or engineering models used for control synthesis. However, the fidelity of a model can be corrected and improved at run-time based on measurements made on the plant. Figure 1 illustrates this concept.

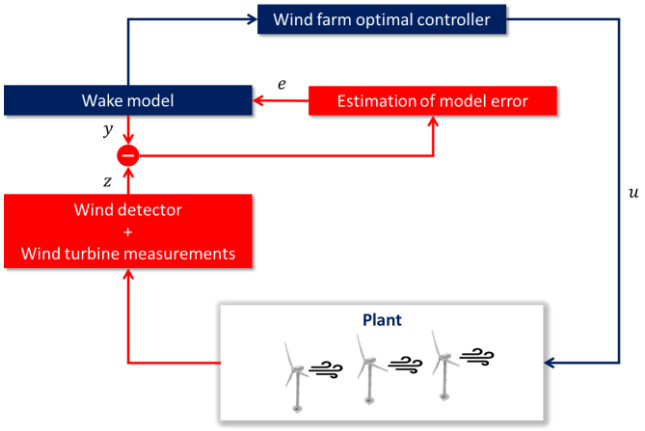

Figure 1. Wind farm control with model updating.

To correct model predictions, one might think of using standard and already available measurements of power and hub-height wind speed, for example using a Kalman filter. Unfortunately, this might not work in general because power and rotor speed might not carry enough informational content to correct for some model errors, as shown later on in this work. In fact, in the case of a wrong power prediction at a downstream wind turbine, one cannot distinguish whether the error is caused by a wrong wind speed in the wake (for example, due to an inaccurate modeling of wake recovery) or by a wrong location of the wake with respect to the impinged rotor disk.

This impasse is solved by using our newly-developed *wake detector* [2]: by using rotor loads, the detector reveals the presence of a wake by mapping blade loads into local estimates of the wind speed over sectors of the rotor disk. This way, a wake model can be improved online during operation of the wind farm (according to the scheme of Fig. 1), generating high quality predictions of the wake *speed* and *position* within the farm. In turn, this improves the control action computed on the basis of these predictions. This idea is developed in the present work with reference to a static wind farm flow model, although nothing in this approach prevents its extension to the dynamic case. Similar concepts of state estimation have been explored in the context of dynamic wake models in [3, 4]. However, it is unclear whether such formulations are able to cope with simultaneous errors in wake recovery and trajectory, as the method presented herein.

This paper is organized as follows. Section II formulates the model update approach, the wind farm model and the load-based wind detector. Section III describes different possible implementations of the model update method. The various options are then tested with reference to experimental measurements obtained on a scaled wind farm facility operated in a large boundary layer wind tunnel. Finally, Section IV summarizes results and conclusions, and gives an outlook towards future work.

II. Methods

A. State update

The model update method is formulated here based on a generic non-linear static wind farm model. A similar formulation could also be derived for a dynamic model, leading in that case to a standard Kalman filtering problem. The static model is written as

$$x = f(u, m, p), \tag{1}$$

$$y = g(x), \tag{2}$$

where f is a non-linear static function, which depends on the model formulation. The control inputs are noted u, and include the yaw and induction of each wind turbine in the farm. Measurements of ambient conditions are noted m, and include density and free stream wind speed and direction (typically estimated by the upstream wind turbines). Physical tunable coefficients of the model and the wind farm layout are represented by the vector of parameters p. The model states are indicated as x, and in the present study they include the velocity and lateral position of the wake of each turbine.

C. L. Bottasso and J. Schreiber are with the Wind Energy Institute, Technical University of Munich, Boltzmannstraße 15, 85748 Garching bei München, Germany (carlo.bottasso@tum.de, johannes.schreiber@tum.de).

A set of outputs y is defined by function g. As shown later on, the outputs may be represented by the turbine power at the downstream turbines, but they may also include estimated flow velocities at the downstream rotors.

In general, the predictions of the model states will be in error, due to a lack of model fidelity, mistuning of the parameters or inaccuracies in ambient conditions. This can be corrected by introducing a state error e. The corresponding corrected state $\hat{x}$ becomes

$$\hat{x} = x + e. \tag{3}$$

A maximum likelihood estimate of the state error can be readily obtained by solving the following problem

$$\min_e \ (z - \hat{y})^{\mathrm{T}} R^{-1} (z - \hat{y}), \tag{4}$$

where z are measurements and $\hat{y}$ the corresponding updated model outputs ($\hat{y} = g(\hat{x})$). For a given fixed covariance R, this procedure corresponds to the method of least squares.

Note that, as ambient wind conditions are often uncertain, the presented formulation could be extended by including these same conditions within the list of states. However, it is also clearly necessary to ensure the observability of all chosen states. For example, a wrong wind direction might not be distinguishable from a wrong wake location. The development of a general formulation for the estimation of wind farm flow model states is a problem of great interest [4], which is however outside of the scope of the present paper.

B. Wind farm model

The wind farm model includes two components: a wake model and a power model. The wake model is based on the double Gaussian profile proposed by [5], combined with the yaw-induced wake deflection developed in [6]. The combination of the two models gives the evolution of the flow speed within the wake downstream of each rotor disk, together with its spatial location. The power model yields the turbine power output by computing the mean flow speed at the rotor using a disk-attached grid. The turbine power coefficient $C_{\mathrm{P},\gamma=0}$ is assumed to be constant below rated wind speed. To take into account the power reduction in misaligned conditions, the following relationship is used

$$C_{\mathrm{P}}(\gamma) = C_{\mathrm{P},\gamma=0} \cos(\gamma)^{p_{\mathrm{P}}}, \tag{5}$$

where γ is the turbine misalignment angle and p_{P} a tunable parameter.

When implementing the state update for wake speed u, Eq. (3) is modified as $\hat{u} = u + re$, where r is the Keane wake reduction (see Eq. (22) of [5]). Since the Keane wake model uses a Gaussian shape for the speed deficit –and hence does not have a well-defined wake width–, this form of the error avoids changing the ambient wind speed away from the wake.

B. Wind detector

A load-based wind speed detector [2] is used to estimate the flow at the downstream wind turbine. As shown in Fig. 2, the detector works by mapping blade loads into local estimates of the wind speed. These are then averaged over sectors of the rotor disk. The resulting sector-effective (SE) wind speed measurements on the left and right parts of the rotor (noted $V^{\mathrm{SE,left}}$ and $V^{\mathrm{SE,right}}$, respectively) are then used in the state update formulation described earlier on.

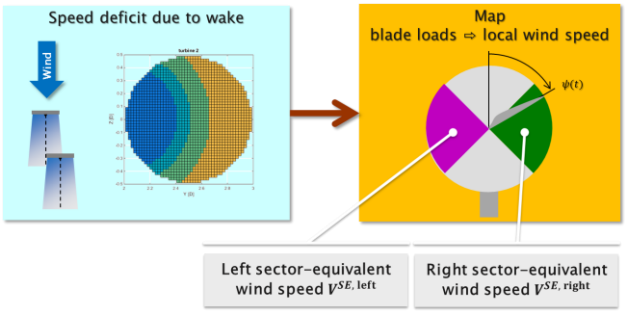

Figure 2. Wind detector estimating the left and right sector-equivalent wind speeds.

III. Implementation & results

A. Implementation

To evaluate the proposed method, three versions of the state update formulation are implemented for a simple farm consisting of two wind turbines. In the notation used below, the upstream wind turbine is indicated as WT1, while the downstream one as WT2.

The *simplistic method* (subscript s) is intended to demonstrate that, by only using power measurements at the downwind turbine ($P_{\mathrm{WT2,exp}}$), it is in general not possible to correct at the same time for errors in lateral wake position (d_{WT1}) and speed (u_{WT1}) of the upstream wind turbine. In contrast to the simplistic method, the *power method* (subscript p) is well-posed, as it only tries to correct the wake speed and not its position based on downstream power measurements. The *wind-sensing method* (subscript ws) includes as measurements also the SE wind speeds $V_{\mathrm{WT2,exp}}^{\mathrm{SE,left/right}}$ obtained by the wind detector on the downwind turbine. This way, the method is able to correct for both speed and position in the wake.

Table 1 gives an overview of the three different approaches. For all cases, the ambient conditions are obtained from the front wind turbine: wind direction is measured by the on-board wind vane, while the ambient wind speed is computed by the rotor effective wind speed corrected for yaw misalignment using Eq. (5).

The diagonal entries of the covariance matrix R are initially set to $1/P_r^2$ for power and to $1/V_r^2$ for the SE wind speed model outputs, where $(\cdot)_r$ indicates a rated quantity. When using maximum likelihood, the covariance is updated after each iteration based on the residuals. Problem (4) is solved using the Nelder-Mead simplex algorithm implemented in the MATLAB function *fminsearch* [7].

TABLE I. STATE UPDATE IMPLEMENTATIONS

Method:	Simplistic ($*= s$)	Power ($*= p$)	Wind-sensing ($*= ws$)
$x_{(*)} =$	$\begin{bmatrix} d_{\mathrm{WT1}} \\ u_{\mathrm{WT1}} \end{bmatrix}$	$[u_{\mathrm{WT1}}]$	$\begin{bmatrix} d_{\mathrm{WT1}} \\ u_{\mathrm{WT1}} \end{bmatrix}$
$\hat{x}_{(*)} =$	$\begin{bmatrix} d_{\mathrm{WT1}} + e_{\mathrm{d}} \\ u_{\mathrm{WT1}} + re_{\mathrm{u}} \end{bmatrix}$	$[u_{\mathrm{WT1}} + re_{\mathrm{u}}]$	$\begin{bmatrix} d_{\mathrm{WT1}} + e_{\mathrm{d}} \\ u_{\mathrm{WT1}} + re_{\mathrm{u}} \end{bmatrix}$
$\hat{y}_{(*)} =$	$[P_{\mathrm{WT2}}]$	$[P_{\mathrm{WT2}}]$	$\begin{bmatrix} P_{\mathrm{WT2}} \\ V_{\mathrm{WT2}}^{\mathrm{SE,right}} \\ V_{\mathrm{WT2}}^{\mathrm{SE,left}} \end{bmatrix}$
$z_{(*)} =$	$[P_{\mathrm{WT2,exp}}]$	$[P_{\mathrm{WT2,exp}}]$	$\begin{bmatrix} P_{\mathrm{WT2,exp}} \\ V_{\mathrm{WT2,exp}}^{\mathrm{SE,right}} \\ V_{\mathrm{WT2,exp}}^{\mathrm{SE,left}} \end{bmatrix}$

B. Experimental setup

Experimental tests with scaled wind turbine models were used to study the performance of the various state update formulations. The scaled turbines, designed for realistic wake behavior, were operated in the boundary layer wind tunnel of the Politecnico di Milano at an ambient hub-height wind speed of 5.8 m/sec and a turbulence intensity of about 5%. A detailed description of the turbines and the wind tunnel can be found in [8, 9]. The wind farm layout is depicted in Fig. 3, where γ_{WT1} is the yaw misalignment of the upstream wind turbine with respect to the wind vector, positive as indicated in the figure. The two turbines are operated at a longitudinal distance of 4 diameters (D) with no lateral displacement.

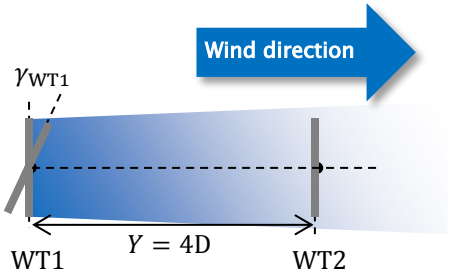

Figure 3. Wind farm layout, top-view.

The wind farm model parameters p include the power loss exponent p_P (see Eq. 5), the coefficients k^*, ϵ, and n that define the wake shape through the expression $\sigma = k^* Y^n + \epsilon$ (where σ is the standard deviation of the double Gaussian wake deficit), and finally the scaling factor c_- [5]. The parameters were first manually tuned with the objective of obtaining a good fit of the model predictions with the experimentally measured wake speed, downstream turbine power and SE speeds at various yaw misalignments of WT1. Figure 4 shows in the upper subplot a comparison between measured (subscript exp) and modeled power at both turbines. The lower subplot shows the SE wind speeds for the left and right sectors of WT2. Each experimental data point represents the mean value of a 60 sec time recording. As the scaled turbine models used in these particular experiments are not equipped with blade load sensors, blade loads were reconstructed from shaft loads using the Coleman Transformation as described in [9]. To account for the fact that the reconstructed experimental blade loads do not contain frequencies above 1P (one per revolution), also the SE wind speed computed from the wind farm model was accordingly filtered. This was obtained by first best-fitting over the turbine rotor disk a linear wind field, and then computing from it the desired quantities $V_{\mathrm{WT2}}^{\mathrm{SE,left/right}}$.

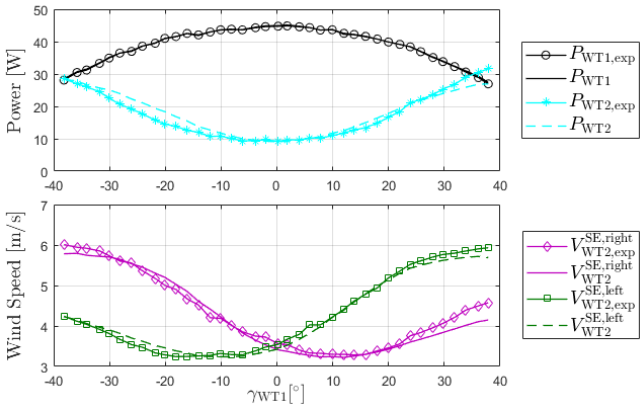

Figure 4. Comparison between experimental and modeled turbine power and SE wind speeds.

C. Results

An experimental time sequence was obtained by stacking one after the other a number of recordings, each one corresponding to a different constant yaw setting of the front machine. Since the flow is turbulent, wake dynamics induced by turbulent fluctuations, including meandering, are included in the recordings. However, the effects of transient changes from one yaw set point to the next are not, including the corresponding travel-time wake delays, which can be estimated to be approximatively equal to 1 sec. Since delays are not included in the static model used here, all signals were filtered with a moving average of 4 sec. The filter window size was chosen to reduce effects of short-term fluctuations, which are believed to be of limited interested for plant-level control.

Figure 5 shows the performance of the simplistic state update method. The upper subplot shows the time history of the upwind turbine yaw position γ_{WT1}, which changes in three steps from 0 deg to 30 deg. Previous experiments indicated that the last yaw position in the plot is the approximate point of maximum power production for the present wind farm configuration. The second subplot shows the experimentally measured power produced by the downwind turbine ($P_{\mathrm{WT2,exp}}$), together with the state updated model prediction (noted $P_{\mathrm{WT2,s}}$, where the second subscript indicates the simplistic formulation). The two lines are essentially identical, indicating an almost perfect prediction of power output by the model. The plot also shows that power increases after each yaw step, which is indeed caused by the wake deflecting laterally and thereby reducing its effects on the downstream rotor.

The third subplot shows the SE wind speeds in the left and right turbine sectors. The experimental measurements from the wind speed detector (solid lines with marker) show the direction of wake deflection: with increasing time and yaw,

the flow velocity in the left sector increases, implying that the wake center is moving to the right. The SE wind speeds of the updated simplistic method are also shown on the same plot. These curves reveal that the model-predicted flow velocities, which were not explicitly taken into account by the method, behave in a radically different way from the measured ones. In fact, the simplistic state update method corrects the wake center position by moving it to the left of the downwind turbine, instead of to the right as it should be. The last subplot of Fig. 5 shows the corresponding state errors. The large error in wake speed significantly alters the wake deficit, while the error in wake position implies that the wake center is located to the left of the rotor.

The simplistic method is clearly ill-posed, as two independent states are corrected using only one measurement. Therefore, multiple combinations of wake speed and displacement can be obtained that, although completely wrong, still apparently lead to a very good power estimate. A controller using the predictions of such a model is invariably bound to fail.

Notice that the ill-posedness of the present formulation is rather obvious, by considering that one single global rotor measurement as power cannot distinguish between changes due to a different wake recovery or position. Indeed, in the context of the present formulation, a well-posedness check can be formulated by considering the linearized version of Eq. (2), i.e. $y = Cx$, where $C = \partial g / \partial x$. The problem can be considered to be well posed if state x can be deduced from measurements z of y, a condition that is satisfied only if $\text{null}(C) = \{0\}$, i.e. if C is of full column rank. This is akin to the observability condition for dynamical systems, specialized to the present static case. For the simplistic approach, C is a 1×2 matrix that cannot satisfy this observability condition.

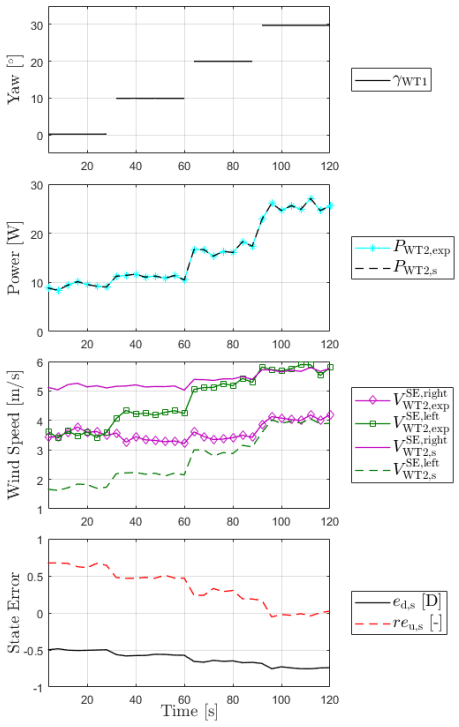

Figure 5. Predictions of the on-line corrected simplistic model (s) compared to experimental measurements (exp).

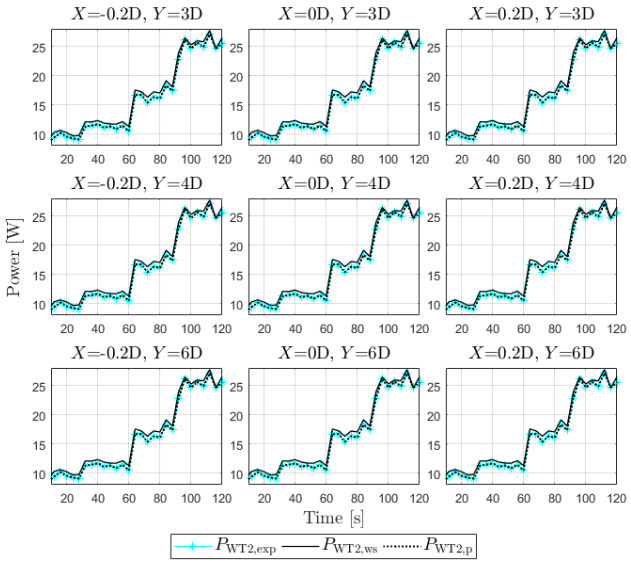

Figure 6. Downwind turbine power predicted by the wind-sensing method ($P_{WT2,ws}$) and the power method ($P_{WT2,p}$), compared with the experimental one ($P_{WT2,exp}$), for various modeling errors.

After having illustrated the ill-posedness of the simplistic method, the power and wind-sensing approaches are compared. In both cases, the problem is now well-posed: for the power method, only wake speed is corrected based on measured power, while for the wind-sensing method the presence of the wake detector allows for the separation of the effects caused by wake speed from those caused by position. To better understand the characteristic of the methods, artificial errors were imposed on the wind farm model. An error in wake recovery and expansion is simulated by changing the modeled longitudinal distance Y between the turbines with respect to the one of the experiments. In addition, to simulate an error in the modeled wake position, the lateral distance X is also varied.

For nine combinations of modeling errors, Fig. 6 reports the model-predicted power together with the experimentally measured one. Independently of the modeling error, it appears that power is always well predicted. The SE wind speeds at the downwind turbine are shown in Fig. 7. Solid lines with markers represent experimental measurements, dash-dotted lines the power method and solid lines without markers the wind-sensing method flow speeds. The wind-sensing method provides predictions that are very close to the experimental measurements, independently of the modeling error. In fact, both wake speed and wake position can be corrected independently by this approach. On the other hand, the power method only corrects wake speed. Therefore, it provides good results only in the case of model errors in the longitudinal displacement (middle column of the subplots). However, as soon as there is also an error in the wake position, flow velocities do not match anymore. These discrepancies may translate into significant deficiencies when it comes to utilizing the wind farm model for control purposes.

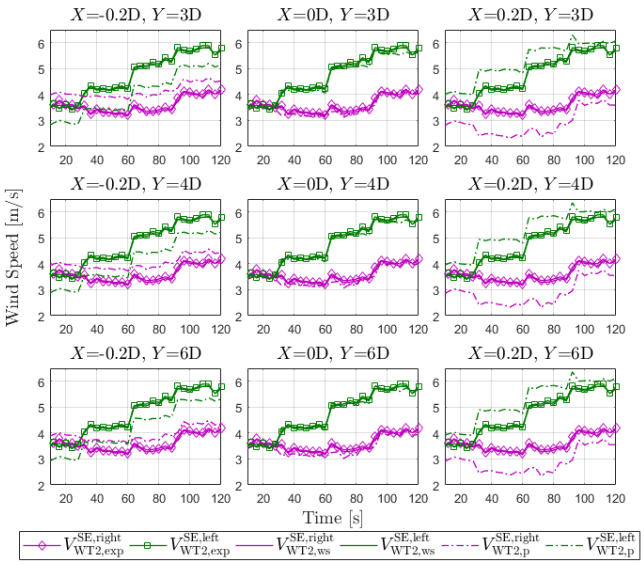

Figure 7. Downwind turbine SE wind speeds predicted by the wind-sensing method ($V_{WT2,ws}^{left/right}$) and the power method ($V_{WT2,p}^{left/right}$), compared to the experimental ones ($V_{WT2,exp}^{left/right}$), for various modeling errors.

To illustrate this point, Fig. 8 shows, for one of the nine cases considered above, the maximum possible wind farm power predicted by the model by yawing the upwind turbine to its optimal position. In the experiment, the optimal position is approximatively equal to 30 deg, which are reached after 90 sec. Even though the power method is apparently able to match the downwind turbine power during the experiment, this is in reality based on a wrong prediction of the flow within the farm. Hence, the maximum predicted power is highly overestimated. On the other hand, the wind-sensing method, being capable of a more faithful prediction of the actual flow, provides for a realistic estimate of the maximum achievable power throughout the whole test case. This highlights the importance of correctly modeling the flow within the wind farm for control purposes.

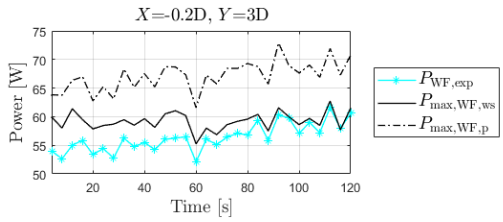

Figure 8. Measured wind farm power ($P_{WF,exp}$) and model-predicted maximum available wind farm power $P_{max,WF}$. For $t > 90$s, the experiment reaches the optimal solution.

IV. Conclusions

A model-based wind farm control algorithm can only be as good as its underlying model. In a realistic scenario, various sources of uncertainties and model defects limit the predictive capabilities of any wind farm flow model. After having calibrated the model offline, the only remaining way to improve this situation is to correct the predictions of the model online, by using measurements obtained on the plant.

The present paper has considered the problem of model-updating, in the context of a well-known static engineering wake model. The model can predict the flow speed within the wake, as well as its geometry and spatial location depending on environmental and wind turbine operational parameters.

Three possible implementations of the method have been considered. The first, and possibly the most natural, tries to correct model predictions by using power measurements on the downstream turbine. Unfortunately, but quite obviously, the method was shown to fail because of its inability to distinguish between effects caused by wake speed or position.

The second approach uses power to correct only for wake speed. This avoids the problem being ill-posed, but clearly cannot correct the predictions of the model whenever the wake position is in error. It was shown that, even in the very simple two-wind turbine case considered here, this formulation may lead to significant errors in maximum power predictions.

Finally, a novel method based on a wake detector was proposed. The wake detector is capable of estimating the local wind speeds on the left and right sectors of the rotor disk. Clearly, the two velocities carry information on the actual location of the wake with respect to the affected rotor. This allows one to distinguish between wake speed and location, and results in the correct update of both states of the engineering model.

The present work is to be considered only as a preliminary study, and further investigations are planned. These include studies of observability in the case of only partially impinging wakes, as well as the investigation of more complex wake interference scenarios. The model update formulation will also be exploited for designing wind farm control laws using optimal model-based approaches.

ACKNOWLEDGMENT

This work has been supported by the CL-WINDCON project, which receives funding from the European Union Horizon 2020 research and innovation program under grant agreement No. 727477.

REFERENCES

[1] Fleming, Paul A.; Gebraad, Pieter M. O.; Lee, Sang; van Wingerden, Jan-Willem; Johnson, Kathryn; Churchfield, Matt et al. (2014): Evaluating techniques for redirecting turbine wakes using SOWFA. In: Renewable Energy 70, S. 211-218. DOI: 10.1016/j.renene.2014.02.015.

[2] Bottasso, C. L.; Cacciola, S.; Schreiber, J. (2018): Local wind speed estimation, with application to wake impingement detection. In: Renewable Energy 116, S. 155–168. DOI: 10.1016/j.renene.2017.09.044.

[3] Shapiro, Carl R.; Meyers, Johan; Meneveau, Charles; Gayme, Dennice F. (2017): Dynamic wake modeling and state estimation for improved model-based receding horizon control of wind farms. In: 2017 American Control Conference (ACC). 2017 American Control Conference (ACC). Seattle, WA, USA, 24.05.2017 - 26.05.2017: IEEE, S. 709–716.

[4] Doekemeijer, B. M.; Boersma, S.; Pao, L. Y.; Van Wingerden, J. W. (2017): Ensemble Kalman filtering for wind field estimation in wind farms. In: 2017 American Control Conference (ACC). 24-26 May 2017. Seattle, WA, USA. American Control Conference; American Automatic Control Council; ACC. Piscataway, NJ: IEEE, S. 19–24.

[5] Keane, Aidan; Aguirre, Pablo E. Olmos; Ferchland, Hannah; Clive, Peter; Gallacher, Daniel (2016): An analytical model for a full wind turbine wake. In: J. Phys.: Conf. Ser. 753, S. 32039. DOI: 10.1088/1742-6596/753/3/032039.

[6] Jiménez, Ángel; Crespo, Antonio; Migoya, Emilio (2010): Application of a LES technique to characterize the wake deflection of a wind turbine in yaw. In: Wind Energ. 13 (6), S. 559–572. DOI: 10.1002/we.380.

[7] Lagarias, Jeffrey C.; Reeds, James A.; Wright, Margaret H.; Wright, Paul E. (1998): Convergence Properties of the Nelder--Mead Simplex Method in Low Dimensions. In: SIAM J. Optim. 9 (1), S. 112–147. DOI: 10.1137/S1052623496303470.

[8] Campagnolo, Filippo; Petrović, Vlaho; Schreiber, Johannes; Nanos, Emmanouil M.; Croce, Alessandro; Bottasso, Carlo L. (2016): Wind tunnel testing of a closed-loop wake deflection controller for wind farm power maximization. In: Journal of Physics: Conference Series 753 (3), S. 32006. DOI: 10.1088/1742-6596/753/3/032006.

[9] Campagnolo, Filippo; Schreiber, Johannes; Garcia, Andrea M.; Bottasso, Carlo L. (2017): Wind Tunnel Validation of a Wind Observer for Wind Farm Control. In: The proceedings of the Twenty-seventh International Ocean and Polar Engineering Conference, ISOPE.

[10] Schreiber, J.; Nanos, E. M.; Campagnolo, F.; Bottasso, C. L. (2017): Verification and Calibration of a Reduced Order Wind Farm Model by Wind Tunnel Experiments. In: Journal of Physics: Conference Series 854, S. 12041. DOI: 10.1088/1742-6596/854/1/012041.