Wind farms providing secondary frequency regulation: Evaluating the performance of model-based receding horizon control[*]

Carl R. Shapiro[1], Johan Meyers[2], Charles Meneveau[1], and Dennice F. Gayme[1]

[1]Department of Mechanical Engineering, Johns Hopkins University, 3400 N Charles St, Baltimore, Maryland 21218, USA
[2]Department of Mechanical Engineering, KU Leuven, Celestijnenlaan 300A, 3001 Leuven, Belgium

Correspondence to: Dennice Gayme (dennice@jhu.edu)

Abstract. We investigate the use of wind farms to provide secondary frequency regulation for a power grid using a model-based receding horizon control framework. In order to enable real-time implementation, the control actions are computed based on a time-varying one-dimensional wake model. This model describes wake advection and wake interactions, both of which play an important role in wind farm power production. In order to test the control strategy, it is implemented in a large eddy simulation (LES) model of an 84-turbine wind farm using the actuator disk turbine representation. Rotor-averaged velocity measurements at each turbine are used to provide feedback for error correction. The importance of including the dynamics of wake advection in the underlying wake model is tested by comparing the performance of this dynamic-model control approach to a comparable static-model control approach that relies on a modified Jensen model. We compare the performance of both control approaches using two types of regulation signals, "RegA" and "RegD", which are used by PJM, an independent system operator in the Eastern United States. The poor performance of the static-model control relative to the dynamic-model control demonstrates that modeling the dynamics of wake advection is key to providing the proposed type of model-based coordinated control of large wind farms. We further explore the performance of the dynamic-model control via composite performance scores used by PJM to qualify plants for regulation. Our results demonstrate that the dynamic-model controlled wind farm consistently performs well, passing the qualification threshold for all fast-acting RegD signals. For the RegA signal, which changes over slower time scales, the dynamic-model control leads to average performance that surpasses the qualification threshold, but further work is needed to enable this controlled wind farm to achieve qualifying performance for all regulation signals.

1 Introduction

Recent market trends are rapidly changing the composition of power grid energy sources, replacing conventional dispatchable power sources with non-dispatchable, variable resources, such as wind energy. These changes are putting pressure on the power system by reducing the number of resources available to provide a wide range of grid services traditionally provided by conventional power plants (Aho et al., 2012). A particularly important example is grid frequency regulation, which is closely tied to short-term imbalances in active power generation and load over time scales ranging from milliseconds to tens

---

[*]This paper is an extended version of our paper presented at the 2016 TORQUE conference: Shapiro, C. R., Meyers, J., Meneveau, C., and Gayme, D. F.: Wind farms providing secondary frequency regulation: Evaluating the performance of model-based receding horizon control, Journal of Physics: Conference Series, 753, 052012, 2016.

of minutes (Rebours et al., 2007). In order to deal with this challenge, a number of independent system operators (ISOs) are beginning to consider requiring wind plants to provide frequency regulation services and expanding frequency regulation markets to include wind plants (Aho et al., 2012; Díaz-González et al., 2014).

Secondary frequency regulation, where participating generators track a power signal sent by an ISO over tens of minutes, is an area of growing interest. Recent work (Aho et al., 2013; Jeong et al., 2014) has shown that stand-alone wind turbines can effectively provide secondary frequency regulation, but recent fluid dynamics simulations (Fleming et al., 2016) have shown that interactions between wakes can lead to poor tracking performance when these single turbine control strategies are applied to an array of turbines (Aho et al., 2013; Jeong et al., 2014). This poor performance is not unexpected because aerodynamic interactions between turbines occur at timescales commensurate with those of the secondary frequency regulation signals. Such considerations have led to recent emphasis on approaches that consider the collective behavior of the farm (Annoni et al., 2016; Gebraad et al., 2015), but few of these studies provide real-time implementable algorithms that can respond to changing wind farm power output levels.

Our recent work (Shapiro et al., 2017) sought to overcome these challenges by developing a time-varying extension of the classic Jensen wake model (Katić et al., 1986) that accounts for the dynamics of wake advection through the farm. This new wake model was incorporated into a predictive model-based receding horizon control framework to coordinate an array of wind turbines to provide secondary frequency regulation by modulating the thrust coefficients of individual turbines. This approach used predictions from the underlying model to iteratively solve an online optimization problem representing the power tracking goal. Feedback from measurements of the velocity at each turbine was used to correct modeling errors. This approach showed promising results when tested in a large eddy simulation (LES) model of a wind farm where turbines were represented using actuator disk models (Shapiro et al., 2017). In these simulations, we used setpoint reductions of only 50% of the maximum regulation provided, but were able track a sample regulation signal with the wind farm test system used. In previous studies (Aho et al., 2013; Jeong et al., 2014), where control design was done at the individual turbine level, successful power tracking required setpoint reductions exactly equal to the maximum change in power production requested by the ISO. The ability to lower the setpoint reduction represents an important advantage over single turbine approaches as the amount of setpoint reduction corresponds directly to the amount of power that wind farms are sacrificing in the bulk energy market to provide regulation. In fact, previous studies have shown that setpoint reductions equal to the one-sided regulation signal variation may not be economically prudent (Rose and Apt, 2014).

The feasibility of providing secondary frequency regulation with wind farms was demonstrated by our initial results (Shapiro et al., 2017). In this work we further evaluate the performance of this approach and consider the effect of reducing the control design and wake model complexity. In particular, we evaluate the importance of explicitly modeling the dynamics of wake advection by comparing the performance of the dynamic-model approach to a similar static-model approach; i.e. we replace the dynamic-wake model with a wake model that does not include wake advection (Katić et al., 1986). In order to make appropriate comparisons, the static-model controller solves an online optimization problem with feedback similar to that solved in the dynamic-model controller.

Both controllers are tested using regulation test signals from PJM, an ISO in the United States Eastern Interconnection (PJM, 2012, 2015). In lieu of a real wind farm, a LES with actuator disk turbine models is used as a "virtual wind farm" to test the controllers. Our results show that the static wake model leads to poor tracking performance, which indicates that the complexity of this particular control design cannot be reduced by ignoring the dynamics of wake advection.

The remainder of the paper further evaluates the performance of the dynamic-model controller using PJM regulation test signals and PJM's performance evaluation criteria (PJM, 2012, 2015). PJM implements secondary frequency regulation through two regulation signals that are based on the Area Control Error (ACE) signal, a combined measure of the power imbalance and deviation of the frequency from its nominal operating value. The "RegA" signal is a low pass filter of the ACE that is generally followed using traditional regulating resources, such as fossil fuel plants. The "RegD" signal is a high pass filter of the ACE that can be followed by more quickly responding resources, such as energy storage devices.

In order to participate in PJM's regulation market, power plants must pass the Regulation Qualification Test for the particular regulation signal being supplied. This test is carried out over a 40-minute period, and the tracking capability is quantified using a composite performance score, which is the weighted sum of accuracy, delay, and precision scores (PJM, 2012, 2015). The accuracy score measures the ability of the signal to respond to a change in the ISO regulation signal. The delay score measures the delay in the plant's response to the regulation signal. The precision score measures the difference between the requested power and the plant's power output. A minimum score of 75% is needed to qualify to participate in each of the two regulation services. Once qualified for a particular service, a plant is continuously evaluated; if its average score over the last 100 hours drops below 40%, then the plant is disqualified from providing the service and must retake the initial performance test to requalify. In this paper, we calculate the performance scores for the dynamic-model controller to determine whether the controlled wind farm system can meet PJM's 75% threshold for qualification. These simulations also allow us to evaluate whether wind farms with this control strategy are better suited to provide traditional (RegA) or fast-acting (RegD) regulation.

The remainder of this paper is organized as follows. The static and dynamic wake models are described in Section 2, and the respective model-based controller designs are outlined in Section 3. Sections 4 and 5 describe the virtual wind farm test system and the simulation cases. The two controllers are compared in Section 6. The performance of dynamic-model control is further explored in Section 7 using PJM's performance criteria. Finally, we present conclusions and discuss directions for future work in Section 8.

2   Wake models

The two wake models employed here are respectively based on static and dynamic adaptations of the classic Jensen wake model (Katić et al., 1986). In this presentation of the Jensen model, we consider regularly arranged wind farms with N rows and M columns of turbines, where each column is aligned with the prevailing wind direction, as shown in Figure 1. The streamwise coordinate is denoted as x, and the n-th turbine row is located at $x = s_n$. Every turbine is assumed to have the same rotor diameter D.
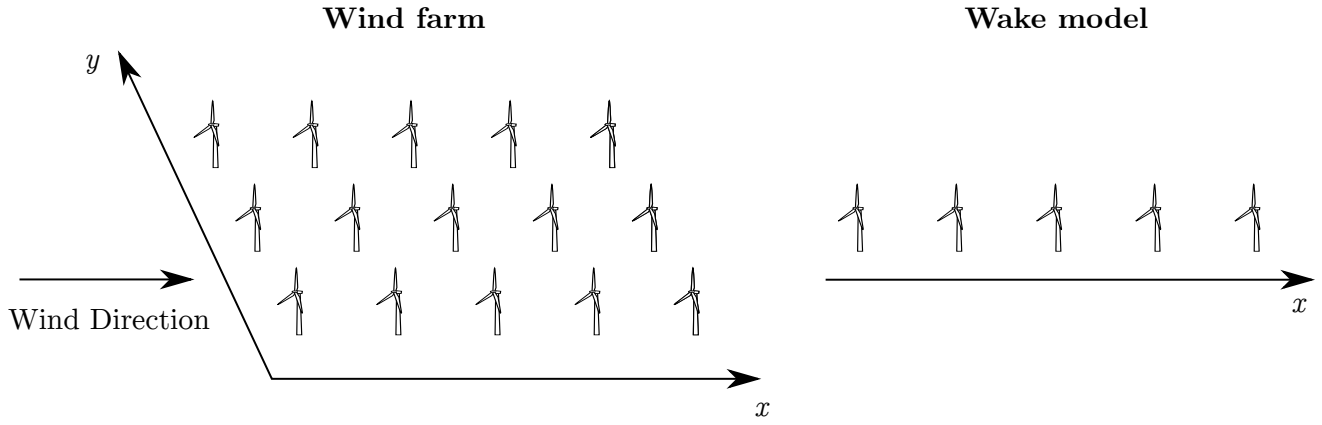
Figure 1. Diagram of a regular wind farm (left) with $N = 5$ rows and $M = 3$ columns of turbines and the corresponding row-averaged wake model representation (right). Each column is aligned with the streamwise coordinate x, and each row is aligned with the spanwise coordinate y.

The standard Jensen model assumes each turbine generates a wake region that expands radially at a linear rate k with increasing downstream distance from the turbine. This leads to following definition of the wake diameter

$$D_w(x) = D + 2kx, \tag{1}$$

where x is the streamwise distance from the turbine rotor plane. Conservation of mass leads to the following velocity deficit in
5 the wake of the m-th turbine in the n-th row

$$\delta u_{nm}(x) = \frac{2U_\infty a_{nm}}{\left[D_w(x - s_n)/D \right]^2}, \tag{2}$$

where a_{nm} is the induction factor and U_∞ is the velocity upstream of the wind farm. This representation yields top-hat profiles of velocity deficits in each cross-stream plane. The velocity field experienced by the each turbine is found by superimposing the squared velocity deficits

10 $$u_{\infty nm} = U_\infty - \sqrt{\sum_{(j,k) \in \mathcal{S}_{nm}} \delta u_{jk}^2(s_n - s_j)}, \tag{3}$$

where $\mathcal{S}_{nm}$ is defined as the set of turbines whose wakes lie within the swept area of the turbine rotor of the m-th turbine in row n. The definition of these sets means that Eq. (3) reduces to $u_{\infty 1m} = U_\infty$ for the first row of turbines. The power production of each turbine is subsequently found using

$$\hat{P}_{nm} = \frac{1}{2}\rho \frac{\pi D^2}{4} C_{Pnm} u_{\infty nm}^3, \tag{4}$$

15 where C_{Pnm} is the power coefficient of the turbine in row n and column m.

For ease of implementation, each wake model used in this paper makes the following modifications to the standard Jensen model. First, we consider each row of turbines collectively (as shown in Figure 1; cf. Shapiro et al., 2017), which means that each modeled value is homogeneous in the spanwise direction and we neglect the spanwise merging of wakes. To reflect this modification, the column index m used in the Eqs. (2)–(3) are dropped in subsequent equations. Second, to account for entrance

5 effects in the farm and compensate for the neglected spanwise wake interactions, we allow each wind turbine row to have a separate wake expansion rate k_n.

Furthermore, we express the turbine power production using the local thrust coefficient C'_{Tn} and modeled velocity at the turbine rotor $\hat{u}_n$. Simple momentum theory can be used to show that (Meyers and Meneveau, 2010; Goit and Meyers, 2015; Shapiro et al., 2017)

10 $$a_n = \frac{C'_{Tn}}{4 + C'_{Tn}}, \qquad \hat{u}_n = (1 - a_n)u_{\infty n}, \quad \text{and} \qquad C'_{Tn} = \frac{C_{Tn}}{(1 - a_n)^2}. \tag{5}$$

Similarly one can show that $C'_{Pn} = C_P/(1 - a_n)^3$, from which we conclude $C'_{Tn} = C'_{Pn}$. Replacing the induction factor a_n in Eq. (2), the modeled upstream velocity $u_{\infty n}$ in Eq. (3), and the power power coefficient C_{Pn} in Eq. (4) with these equations, the power production can be rewritten as

$$\hat{P}_n = M \frac{1}{2} \rho \frac{\pi D^2}{4} C'_{Tn} \hat{u}_n^3. \tag{6}$$

15 These idealized conditions assume that the electrical power generated by the turbine is proportional to the power extracted from the flow and the control actions do not significantly affect the aerodynamic efficiency of the blades (Goit and Meyers, 2015, Appendix A). The following subsections describe the static and dynamic wake models in more detail.

2.1 Static wake model

The static wake model used in this work is the Jensen model with the modifications described above. In order to use this
20 static wake model in a model-based controller for the farm power production, the model must be augmented to account for time-varying changes in the local thrust coefficient $C'_{Tn}(t)$. Including time dependency in the thrust coefficient and replacing the induction factor in Eq. (2) with the expression in Eq. (5) gives the following expression for the velocity deficit for the n-th turbine

$$\delta u_n(x,t) = \frac{C_{Tn'}(t)}{4 + C'_{Tn}(t)} \frac{2U_\infty}{[1 + 2k_n(x - s_n)D]^2}. \tag{7}$$

25 With this approach, thrust coefficient changes instantaneously affect the velocity deficit everywhere; i.e., the wakes implicitly have an infinitely fast advection speed. Finally, the velocity at the turbines of the n-th row can be found by explicitly writing out the set of upstream turbines in Eq. (3) affecting the velocity at the n-th turbine and using the equation for the rotor-averaged velocity in Eq. (5)

$$\hat{u}_n(t) = \left(1 - \frac{C_{Tn'}(t)}{4 + C'_{Tn}(t)}\right) \left(U_\infty - \sqrt{\sum_{m=1}^{n-1} \delta u_n^2(s_n - s_m)}\right). \tag{8}$$

Eqs. (7) and (8) are therefore the static wake model equations $\mathbf{W}_s(\mathbf{C}'_T, \mathbf{q}_s) = \mathbf{0}$, where $\mathbf{q}_s = [\boldsymbol{\delta}\mathbf{u}, \hat{\mathbf{u}}]$ denote the model states and boldface indicates vectors.

2.2 Dynamic wake model

The dynamic wake model also borrows from the classic Jensen model, but instead allows the wake velocity deficit to move
5   downstream at a finite velocity. The resulting one-dimensional time-varying wake model, which assumed that the wake travels with the inlet velocity U_∞, was previously proposed and validated against LES of wind farms at startup (Shapiro et al., 2017). This led to the velocity deficit being governed by

$$\frac{\partial \delta u_n}{\partial t} + U_\infty \frac{\partial \delta u_n}{\partial x} = -w_n(x)\delta u_n(x,t) + f_n(x,t), \tag{9}$$

where $w_n(x)$ is the wake decay function and $f_n(x,t)$ is a forcing function used to account for the effect of the turbine on the
10   flow field. The wake decay function

$$w_n(x) = 2\frac{U_\infty}{d_n(x)}\frac{d}{dx}d_n(x) \tag{10}$$

is determined by assuming that the wake diameter normalized by the rotor diameter $d_n(x) = D_{wn}(x)/D$ at a fixed location x is constant in time. Momentum theory shows that as the air flows through the turbine rotor, the velocity reduces to $U_\infty - 2U_\infty C'_{Tn}/(4+C'_{Tn})$ (Shapiro et al., 2017). In order to retrieve this expected velocity reduction, the forcing function is specified
15   as

$$f_n(x,t) = \frac{2U_\infty^2}{d_n^2(x)}\frac{C'_{Tn}(t)}{4+C'_{Tn}(t)}G(x-s_n), \tag{11}$$

where $G(x-s_n)$ is a smoothing function that integrates to unity, centered at the streamwise location of the turbine $x = s_n$. A Gaussian function with characteristic width Δ

$$G(x-s_n) = \frac{1}{\Delta\sqrt{2\pi}}e^{-\frac{(x-s_n)^2}{2\Delta^2}} \tag{12}$$

20   maintains smoothness in the velocity deficit fields.

In the Jensen model (Katić et al., 1986), the dimensionless diameter of the wake generated by turbine row n is $d_n(x) = 1 + 2k_n(x - s_n)/D$, where k_n is an empirical wake expansion coefficient. We make two modifications to this equation. First, the linear expansion is assumed to begin at $x = s_n + 2\Delta$ to prevent the wake expansion from occurring within the induction zone imposed by the Gaussian forcing. The second modification addresses the fact that the equation for the standard Jensen
25   dimensionless wake diameter is ill-posed upstream of the turbine, where it can vanish or become negative. Therefore, we use the following modified function that smoothly approximates the linear expansion in the far wake while avoiding becoming less than unity close to the turbine

$$d_n(x) = 1 + k_n \ln\left[1 + \exp\left(\frac{x - s_n - 2\Delta}{D/2}\right)\right]. \tag{13}$$

As in the static model, squared deficits (Katić et al., 1986) are superposed to calculate the estimated streamwise velocity $\hat{u}_n$ at the turbine

$$\hat{u}_n(t) = U_\infty - \int_0^L \left(\sum_{m=1}^N \delta u_m^2(x,t) \right)^{1/2} G(x - s_n)\, dx. \tag{14}$$

Finally, the total estimated power $\hat{P}_n$ of the M turbines in row n is found using Eq. (6). The dynamic wake model equations Eqs. (9)–(14) are written as $\mathbf{W}_d(\mathbf{C}_T', \mathbf{q}_d) = \mathbf{0}$, where $\mathbf{q}_d = [\boldsymbol{\delta u}, \hat{\mathbf{u}}]$ denote the model states.

3 Controlled wind farm designs

The model-based controllers implementing the static and dynamic wake models discussed above are designed to track the power reference signals $P_{\text{ref}}(t)$ sent by an ISO by modulating the thrust coefficients of each turbine row $C_{Tn}'(t)$. In both cases, feedback is included by measuring the row-averaged, rotor-averaged wind speed $u_n(t)$. The resulting feedback term $\epsilon_n(t)$ is fed into the controller and used to correct the predicted power output of the wake model. A block diagram of the controlled wind farm system is shown in Figure 2.
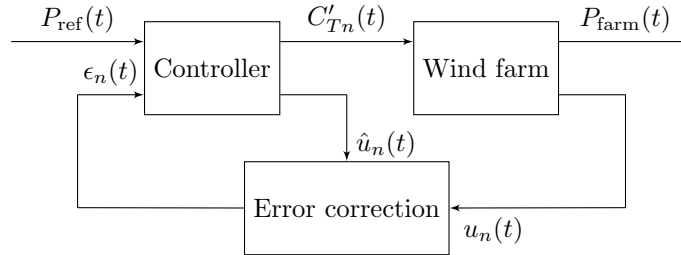

Figure 2. Block diagram of the controlled wind farm system for both controllers. Each controller computes a thrust coefficient command signal $C_{Tn}'(t)$ using the reference signal $P_{\text{ref}}(t)$ and an error correction term $\epsilon_n(t)$. The error correction is computed using the measured velocity $u_n(t)$ and the predicted velocity $\hat{u}_n(t)$ from the underlying wake model.

3.1 Controller designs

Each controller calculates the local thrust coefficient trajectories by repeatedly solving a minimization problem of the form

$$\underset{\mathbf{C}_T', \mathbf{q}}{\text{minimize}} \quad \mathcal{J}(\mathbf{C}_T', \mathbf{q}) + \mathcal{R}(\mathbf{C}_T') \tag{15}$$

$$\text{subject to} \quad \mathbf{W}(\mathbf{C}_T', \mathbf{q}) = 0, \tag{16}$$

where $\mathbf{W}(\mathbf{C}_T', \mathbf{q}) = \mathbf{0}$ and $\mathbf{q}$ are placeholders for the static and dynamic wake model and states, which were previously indicated by the subscripts s and d, respectively. The cost functional $\mathcal{J}$ represents the reference tracking goal, and the functional $\mathcal{R}$

are regularizations to maintain well-behaved control trajectories. These regularizations include a penalization for fast changes in the thrust coefficients to avoid excessive oscillations in the control and a penalization for deviations away from the pre-control reference power to prevent thrust coefficients from moving outside of physical bounds.

Although both controllers solve an online optimization problem, the mechanics of the implementation are quite different. Since the equations for the static model have no dynamics, every instance in time is uncoupled. Therefore, the static-model controller can consider each point in time as a separate minimization problem, and the cost functionals can be written solely in terms of the current state. With this approach, the power tracking cost functional at time t is written as

$$\mathcal{J}_s(\mathbf{C}_T', \mathbf{q}_s) = \frac{1}{\mathcal{P}^2}\left(\sum_{n=1}^{N} \hat{P}_n(t) - P_{\text{ref}}(t)\right)^2,$$ (17)

and the regularization terms are

$$\mathcal{R}_s(\mathbf{C_T'}) = \eta \sum_{n=1}^{N} \left(C_{Tn}'(t) - C_{T,\text{ref}}\right)^2 + \gamma T^2 \sum_{n=1}^{N}\left(\frac{dC_{Tn}'}{dt}\right)^2.$$ (18)

The constants $\mathcal{P}$ and T in Eq. (17)–(18) are used to make each term in the power tracking cost functional and regularization functional dimensionless and of comparable magnitude. Here we choose $\mathcal{P} = M\frac{1}{2}\rho\frac{\pi D^2}{4}U_\infty^3$ and time T as the time horizon of the reference signal considered. The constants η and γ are the respective weights of each regularization term.

The dynamic-model controller (Shapiro et al., 2017), on the other hand, accounts for the time-dependent advection of turbine wakes. We therefore employ a model-based receding horizon framework, which is a predictive approach that uses the model to plan future control actions. The receding horizon method works by iteratively solving a finite-time minimization problem over a time horizon T. The solution is implemented for a shorter period T_A before re-solving the minimization problem. More details about this procedure can be found in (Bewley et al., 2001; Goit and Meyers, 2015). With this predictive framework, the reference tracking goal is represented by the cost functional

$$\mathcal{J}_d(\mathbf{C}_T', \mathbf{q}_d) = \frac{1}{\mathcal{P}^2 T}\int_0^T \left(\sum_{n=1}^{N} \hat{P}_n(t) - P_{\text{ref}}(t)\right)^2 dt,$$ (19)

and the regularization functional is defined as

$$\mathcal{R}_d(\mathbf{C}_T') = \frac{\eta}{T}\sum_{n=1}^{N}\int_0^T \left(C_{Tn}'(t) - C_{T\text{ref}}'\right)^2 dt + \gamma T \sum_{n=1}^{N}\int_0^T \left(\frac{dC_{Tn}'}{dt}\right)^2 dt.$$ (20)

Consistent with the distinction between the non-predictive and predictive nature of the static-model and dynamic-model controllers, the functionals for the static wake model are not integrated forward in time. In other words, the static wake model is not a receding horizon method because the modeled system does not include dynamics.

All minimizations are solved using the modified unconstrained reduced cost functional $\tilde{\mathcal{J}}(\mathbf{C}_T') = \mathcal{J}(\mathbf{q}, \mathbf{C}_T')$ (Bewley et al., 2001; Goit and Meyers, 2015), instead of the cost functional defined in Eq. (15). Minimizations are performed using the gradient-based nonlinear Polak-Ribière conjugate gradient method (Press, 2007) combined with the Moré-Thuente line search

method (Moré and Thuente, 1994) and terminated after 100 iterations. For the static wake model, gradients are obtained using finite differencing, which can be implemented efficiently because there are only N control variables per minimization. For the dynamic wake model, gradients are obtained using one backward simulation of the adjoint equations of the wake model using the formal Lagrangian method (Goit and Meyers, 2015; Borzì and Schulz, 2011). The full procedure is detailed in detail in Shapiro et al. (2017). This approach was chosen because it is computationally efficient for systems with large state spaces, such as the discretized PDE system described by the dynamic wake model.

3.2 Measurement feedback

As shown in Figure 2, controllers employ closed-loop feedback for velocity measurements at each turbine to correct modeling error and assumptions. The row-averaged power and row-averaged, rotor-averaged are defined from velocities u_{nm} measured at every turbine in the wind farm

$$P_n = \sum_{m=1}^{M} P_{nm}, \quad \text{and} \quad u_n = \frac{1}{M}\left(\sum_{m=1}^{M} u_{nm}^3\right)^{1/3}. \tag{21}$$

The definition of the row-average velocity at the turbine disk is necessary to ensure that $P_n = M \frac{1}{2}\rho\frac{\pi D^2}{4} C'_{Tn} u_n^3$. These measurements are used to calculate an error term ϵ_n and provide feedback by replacing Eq. (6) with

$$\hat{P}_n = M \frac{1}{2}\rho\frac{\pi D^2}{4} C'_{Tn} (\hat{u}_n + \epsilon_n)^3. \tag{22}$$

For the static wake model, the error term for turbine row n at time step k is calculated using the difference between the measured and estimated velocity from the previous time step $k-1$

$$\epsilon_n^k = u_n^{k-1} - \hat{u}_n^{k-1}. \tag{23}$$

For the dynamic wake model, the error correction at the receding horizon iteration starting at time t_c is

$$\epsilon_n(t) = (u_n(t_c) - \hat{u}_n(t_c)) e^{-(t-t_c)/\tau}. \tag{24}$$

The exponential decay accounts for the reduced future accuracy of the error term in the receding horizon prediction.

4 Virtual wind farm test system

A LES model of a wind farm with wind turbines represented using actuator disk models is used to test the two control approaches. The wind farm is composed of $N = 7$ rows of $M = 12$ aligned columns of turbines. Each turbine has a 100 m rotor diameter D and a 100 m hub height. The spacing between turbines is $7D$ in the streamwise direction and $5D$ in the spanwise direction. Prior to the initiation of the control actions, all of the turbines are operated at a constant reference local thrust coefficient of $C'_{T,\text{ref}} = 1.33$, which is assumed to be representative of wind turbines operating in region 2 (Calaf et al., 2010; Stevens et al., 2014a).

These simulations are performed using JHU's LES code (LESGO) (Calaf et al., 2010; Stevens et al., 2014b; VerHulst and Meneveau, 2015), which uses pseudo-spectral discretization in the horizontal directions with periodic boundary conditions. The code also employs second-order Adams-Bashforth time integration, second-order finite differencing in the vertical direction, and the dynamic scale-dependent Lagrangian Smagorinsky subgrid stress model (Bou-Zeid et al., 2005). Inlet conditions for the wind farm are generated using the concurrent-precursor method (Stevens et al., 2014b). The force exerted by and the power extraction of the m-th turbine of the n-th row are both a function of the filtered rotor-averaged velocity u_{nm} (Calaf et al., 2010) and the thrust coefficient C'_{Tnm}. The force is modeled as a drag force $F_{nm} = -\frac{1}{2}\rho\frac{\pi D^2}{4}C'_{Tnm}u_{nm}^2$, and the power extraction is $P_{nm} = -F_{nm}u_{nm}$. An instantaneous color contour plot of the streamwise velocity field from one of these simulations is shown in Figure 3.
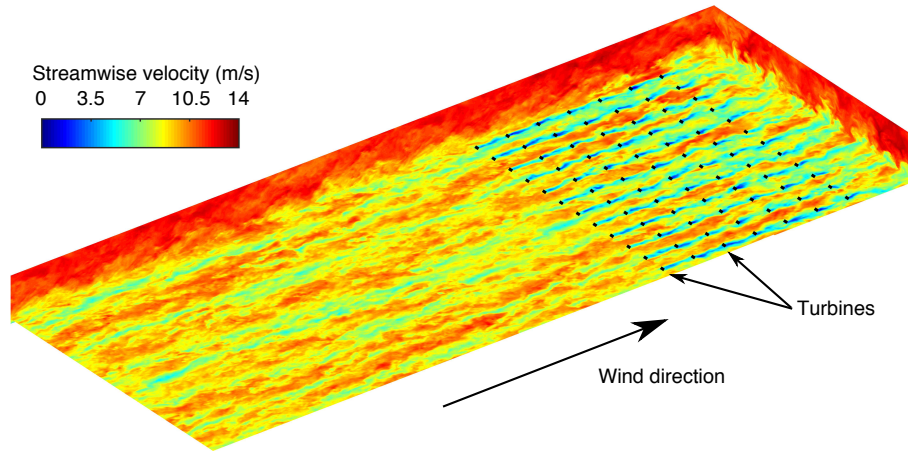

Figure 3. Instantaneous streamwise velocity contours for a large eddy simulation with actuator disk turbine models, which are indicated by black lines. The inlet conditions are generated using a concurrent precursor simulation (Stevens et al., 2014b) shown at the beginning of the plotted domain.

5 Test cases

The performance of the controlled wind farms are tested using PJM's published RegA and RegD test signals as well as historical RegA and RegD signals from three hours in 2015 (PJM, 2016). For each regulation signal, we examined three initial conditions for the wind farm and two levels of power setpoint reduction, which we also refer to as derates. For each controller test, the reference signal is defined as $P_{\mathrm{ref}}(t) = [1 - \alpha + 0.08r(t)]P_{\mathrm{base}}$, where P_{base} is the 5-minute average power prior to initiation of the control, α is the derate amount, and $r(t) \in [-1, 1]$ is the regulation signal from the ISO. As a result, the reference signal varies by $\pm 8\%$ of the baseline power P_{base}. In all cases, the baseline power is determined using an uncontrolled simulation

Table 1. Test case identifiers

Identifier	Type	Description
RegA	Signal	Traditional RegA regulation signal
RegD	Signal	Fast-responding RegD regulation signal
D4	Derate	Power setpoint is reduced by 4% of P_{base}
D6	Derate	Power setpoint is reduced by 6% of P_{base}
IC1	Initial condition	Initial condition 1
IC2	Initial condition	Initial condition 2
IC3	Initial condition	Initial condition 3
TS	Period	PJM test signals
H1	Period	PJM historical hour 1
H2	Period	PJM historical hour 2
H3	Period	PJM historical hour 3

with a local thrust coefficient of $C'_T = 1.33$, which corresponds to region 2 operation (Stevens et al., 2014a) with a total farm power of approximately 100 MW.

A total of 48 test cases are used through a combination of the four variables discussed above. Each test case is given a unique identifier that is a combination of identifiers for each of the variable types shown in Table 1. "Signal" refers to the regulation signal type (RegA or RegD), "Derate" refers the to derate amount (4 or 6%), "Initial condition" refers to the initial condition of the controlled plant simulation, and "Period" refers to the regulation signal period, which is either the PJM test signals or one of the selected hours in 2015. For example, the test case "RegA.D4.IC1.TS" refers to the case with the RegA test signal, 4% derate, and the first initial condition.
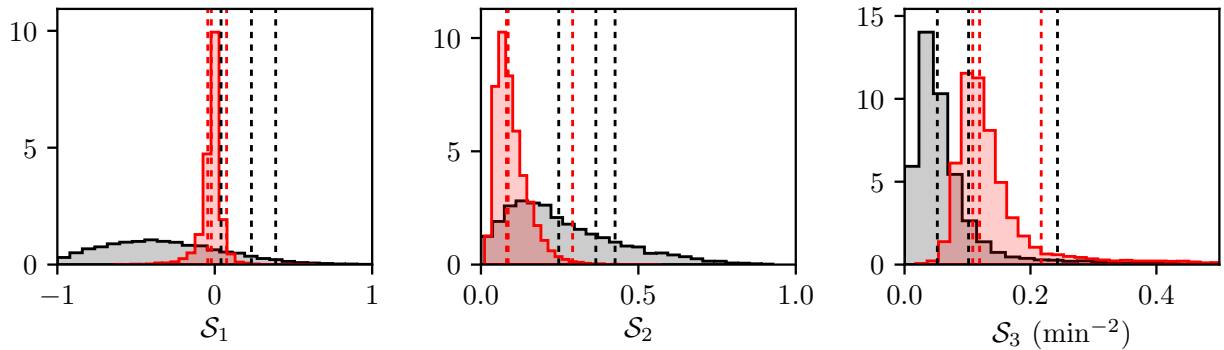

Figure 4. Probability density functions of $\mathcal{S}_1$–$\mathcal{S}_3$ defined in Eq. (25) for RegA (black) and RegD (red) during 2015. The three selected historical hours are shown in the PDFs as vertical dashed lines.

The number of historical hours used to test the controlled wind farm is constrained by the computational cost of running the model wind farm LES. As a result, it is impractical to select enough hours to sample the entire range of possible regulation signals provided by PJM. To prevent systematic bias, the three hours were selected without considering the characteristics of the regulation signals during those periods.

After selecting the hours, we compare the statistics of the selected signals to the range of all possible regulation signals provided by PJM in 2015. We measure the characteristics of the hourly PJM signals using three statistics

$$\mathcal{S}_1 = \frac{1}{T} \int_0^T r(t)\,dt \qquad \mathcal{S}_2 = \frac{1}{T} \int_0^T r^2(t)\,dt - \mathcal{S}_1^2 \qquad \mathcal{S}_3 = \frac{1}{T} \int_0^T \left(\frac{dr}{dt}\right)^2\,dt, \tag{25}$$

where $T = 60$ min, $\mathcal{S}_1$ is the mean of $r(t)$, $\mathcal{S}_2$ is the variance of $r(t)$, and $\mathcal{S}_3$ is the variance of $\frac{dr}{dt}$. For each of these statistics, the probability density function (PDF) is calculated using all possible hourly signals provided by PJM in 2015, and is shown in Figure 4. These PDFs demonstrate the differences between the RegA and RegD signals. The RegA signals have a larger mean and variance than the RegD signals, but the variance of $\frac{dr}{dt}$ is smaller. The values of these statistics for the three selected hours is compared to the PDFs over the entire year in Figure 4. Using these statistics, the selected historical signals represent a reasonable cross section of the possible PJM regulation signals. The only exception, the high percentile ranking in $\mathcal{S}_1$ of the RegA signals, represents a more difficult test for the controlled wind farm because more energy is requested than the average.

We set the regularization parameters for both methods to $\eta = 0.005$, $\gamma = 2.083 \times 10^{-5}$, and $\tau = 120$ s. The required parameters of the static and dynamic wake models, inlet velocity U_∞ and wake expansion coefficients k_n, are calculated for each initial condition using measurements form the five minutes prior to initialization of the control. The inlet velocity is set using the relationship $U_\infty = (4 + C'_{T,\mathrm{ref}})C'^{-1}_{T,\mathrm{ref}} T^{-1} \int_{-5\mathrm{min}}^0 u_1(t)\,dt$, and the inlet velocities are found using a least squares fit between the measured power and the modeled power using the static model. The resulting parameters are shown in Table 2.

Table 2. Wake model inlet velocity U_∞ and wake expansion coefficients k_n by initial condition.

Initial condition	U_∞ (m/s)	k_1	k_2	k_3	k_4	k_5	k_6	k_7
1	9.65	0.028	0.049	0.041	0.047	0.053	0.054	0.054
2	9.32	0.026	0.046	0.043	0.047	0.054	0.052	0.052
3	9.64	0.026	0.040	0.040	0.037	0.044	0.041	0.041

6 Comparison of control methods

The power tracking performance and control trajectories for the two control methods are compared for both the RegA and RegD test signals with 4% derates in Figures 5 and 6, respectively. The dynamic-model control demonstrates good overall tracking performance, although it has some trouble meeting the reference signal during the last 5–10 minutes of the RegA.D4.IC1.TS

and RegA.D4.IC4.TS cases. On the other hand, the static-model control demonstrates poor overall tracking performance, although it is able to track the signal for several periods, such as around minute 20 of all cases in Figure 5 and the first several minutes of the cases in Figure 6.

5   The performance of the static-model control provides an interesting demonstration of the importance of including time dependency. Consider the change in operating state that occurs around minute 5 of the RegA.D4.IC2.TS case in Figure 5. The controller quickly moves to a state of low upstream thrust coefficients and the Betz-optimal thrust coefficient $C_T' = 2$ (Goit and Meyers, 2015) for the last row. In the context of the static Jensen model, this change is modeled as an instantaneous reduction of the wake velocity deficit throughout the farm. Since the reduced wake deficit must travel through the farm before this change in state is reached, the power production of the farm plunges, slowly recovering to the desired setpoint around minute 15.

10   Similarly, all of the static-model control cases in Figures 5 and 6 demonstrate significant overshoot in the power production during the first 30 seconds of the simulations as the thrust coefficients quickly move away from the pre-control level. In the static-model control case, the Jensen model assumes that changes in rotor-averaged velocity occur instantaneously, producing large transient variations when moving between thrust coefficient states. In reality, however, the air around the turbine will slowly respond to a sudden change in the thrust coefficient, producing a delayed response in the rotor-averaged velocity.

15   Neglecting the wake advection time in the Jensen model thus explains the poor performance of the static model controller. Unlike the static-model based control, the dynamic-model control includes the time-dependent effects of wake advection and does not produce large overshoots when changing between states. In the next section we further study the performance of this dynamic-model control approach.

7   Performance evaluation of dynamic-model control

20   The time evolution of the total farm power in LES are compared to the reference signals for initial condition 3 and a 4% derate in Figure 7, which shows all regulation signals (RegA or RegD) and regulation period combinations. These results demonstrate the good overall tracking performance of the controlled wind farm, except for a few specific periods of under-performance. Furthermore, the results demonstrate that the dynamic-model based receding horizon control method is also able to reduce the natural turbulent fluctuations in the power production of the wind farm. Quantitative measures of the performance of

25   each regulation signal type (RegA or RegD) for derate values of 4% and 6% are shown in terms of PJM's performance scores in Figure 8. The controlled wind farm performs better for the RegD signals, meeting the composite score threshold for qualification of 75% in all cases. The performance of the RegA signals is also satisfactory for PJM participation, but the controller would not have qualified in all tests. These lower composite scores may be explained by the large values in $\mathcal{S}_1$, which represents total energy requested in the signals, compared to other PJM signal. However, in cases where the controlled wind

30   farm had poor performance for the RegA signal with 4% derate, increasing the derate to 6% markedly improved the overall performance.
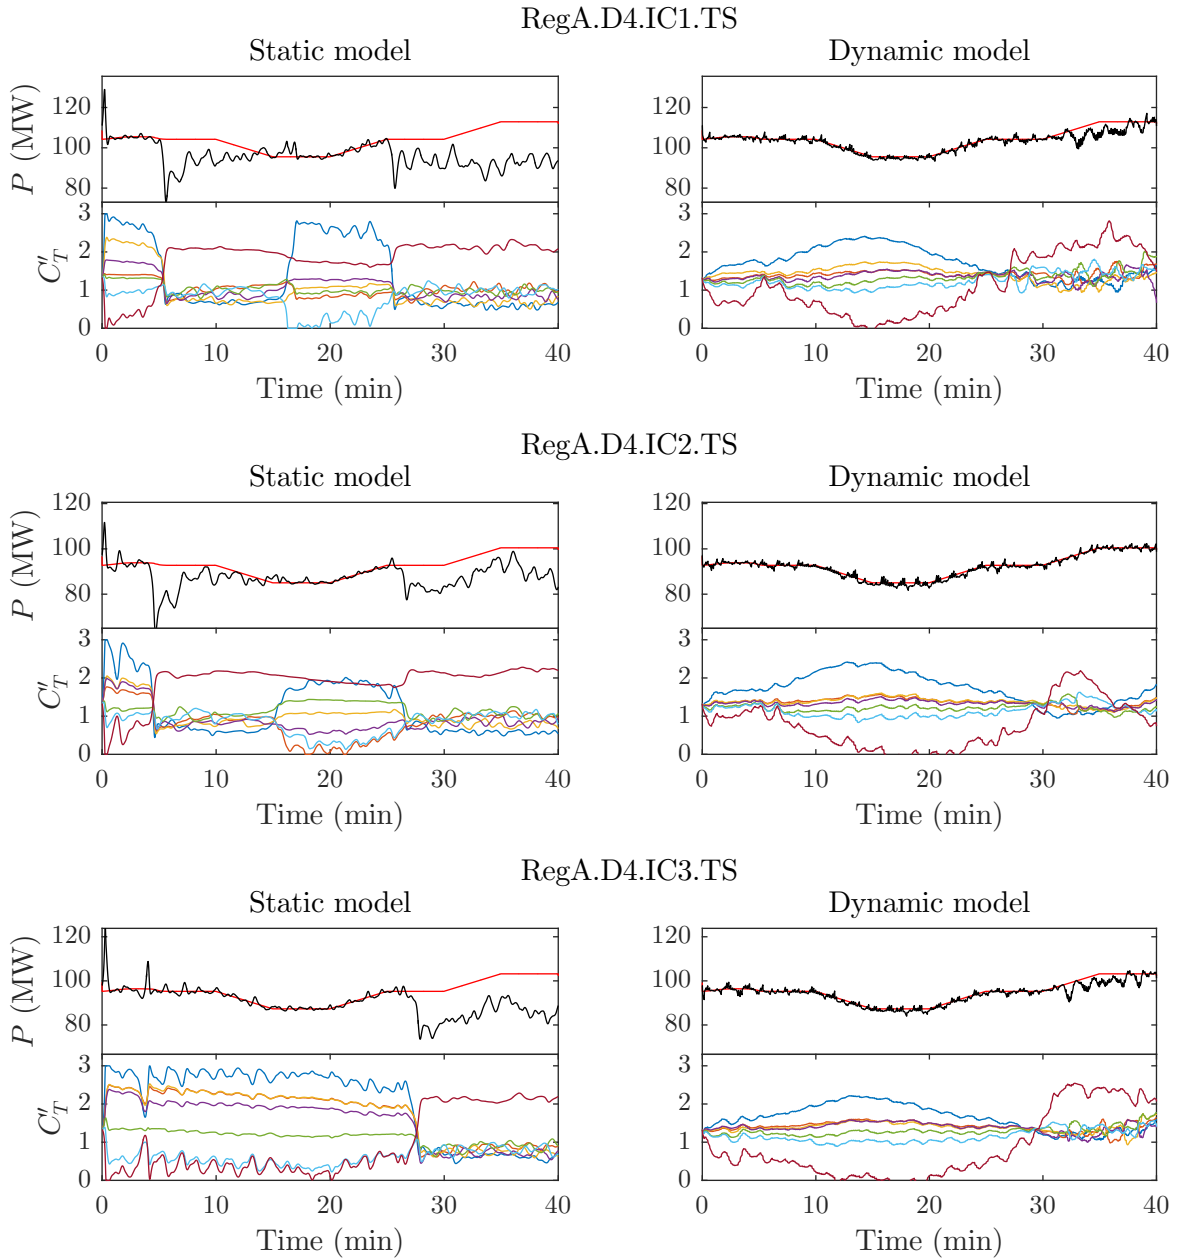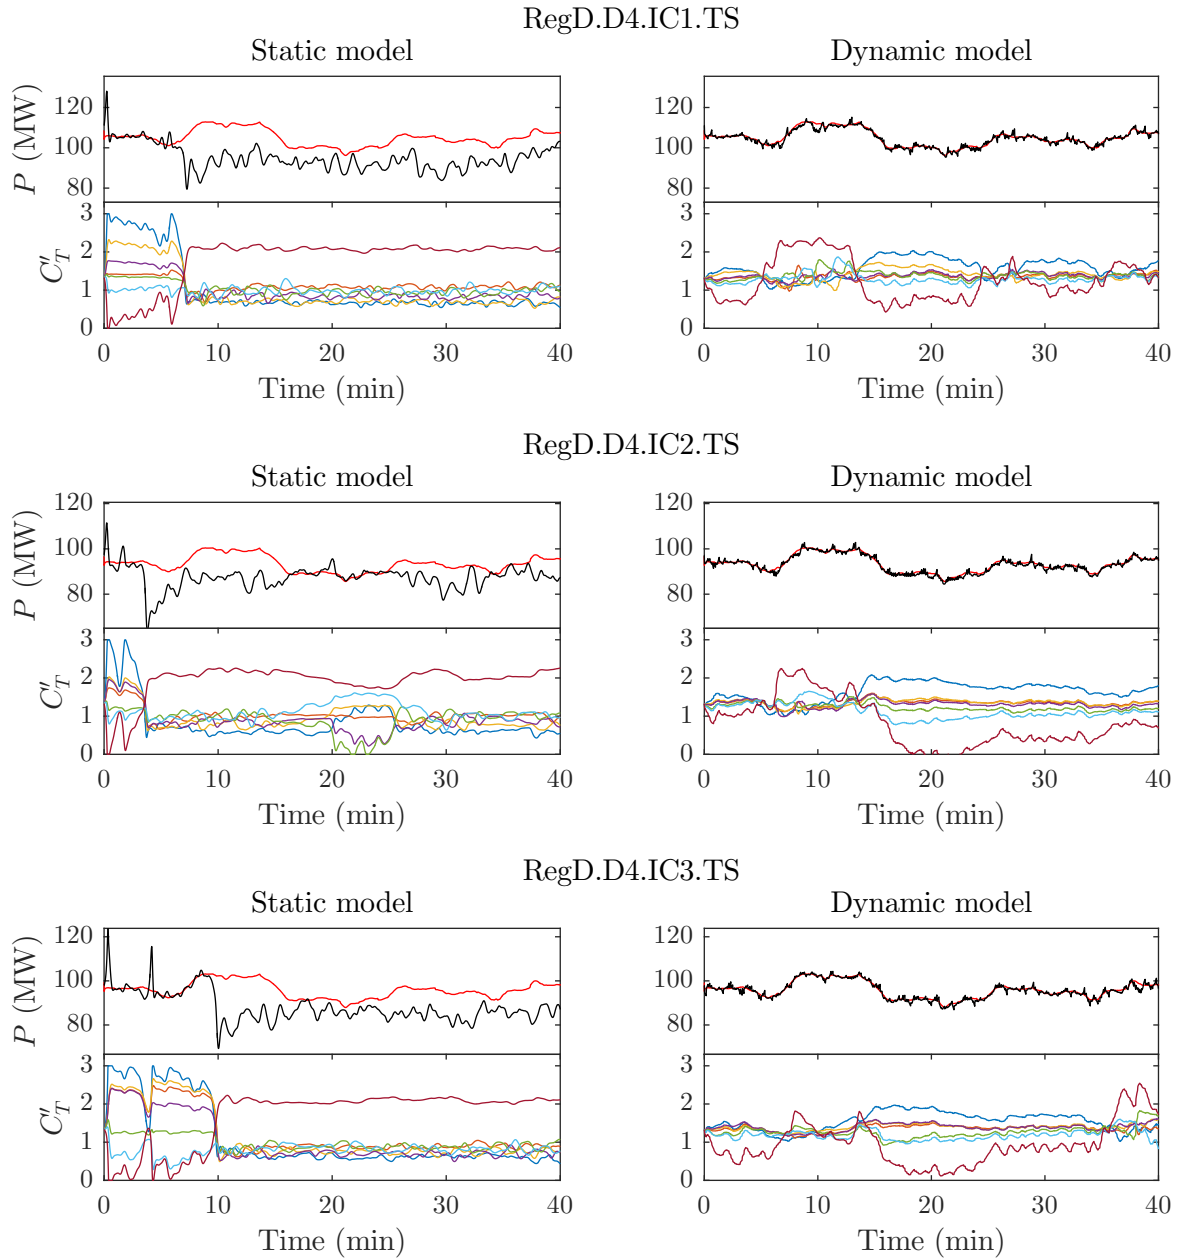
Figure 5. Comparison of static-model (left) and dynamic-model (right) control methods for RegA test signals with 4% derates. All three initial conditions 1–3 are shown from top to bottom. The top panel in each row shows the controlled wind farm power production (——) compared to the reference signal (——). The bottom panel in each row shows the local thrust coefficients calculated by control methods by row: row 1 (——), row 2 (——), row 3 (——), row 4 (——), row 5 (——), row 6 (——), row 7 (——).

Figure 6. Comparison of static-model (left) and dynamic-model (right) control methods for RegD test signals with 4% derates. All three initial conditions 1–3 are shown from top to bottom. The top panel in each row shows the controlled wind farm power production (——) compared to the reference signal (——). The bottom panel in each row shows the local thrust coefficients calculated by control methods by row: row 1 (——), row 2 (——), row 3 (——), row 4 (——), row 5 (——), row 6 (——), row 7 (——).
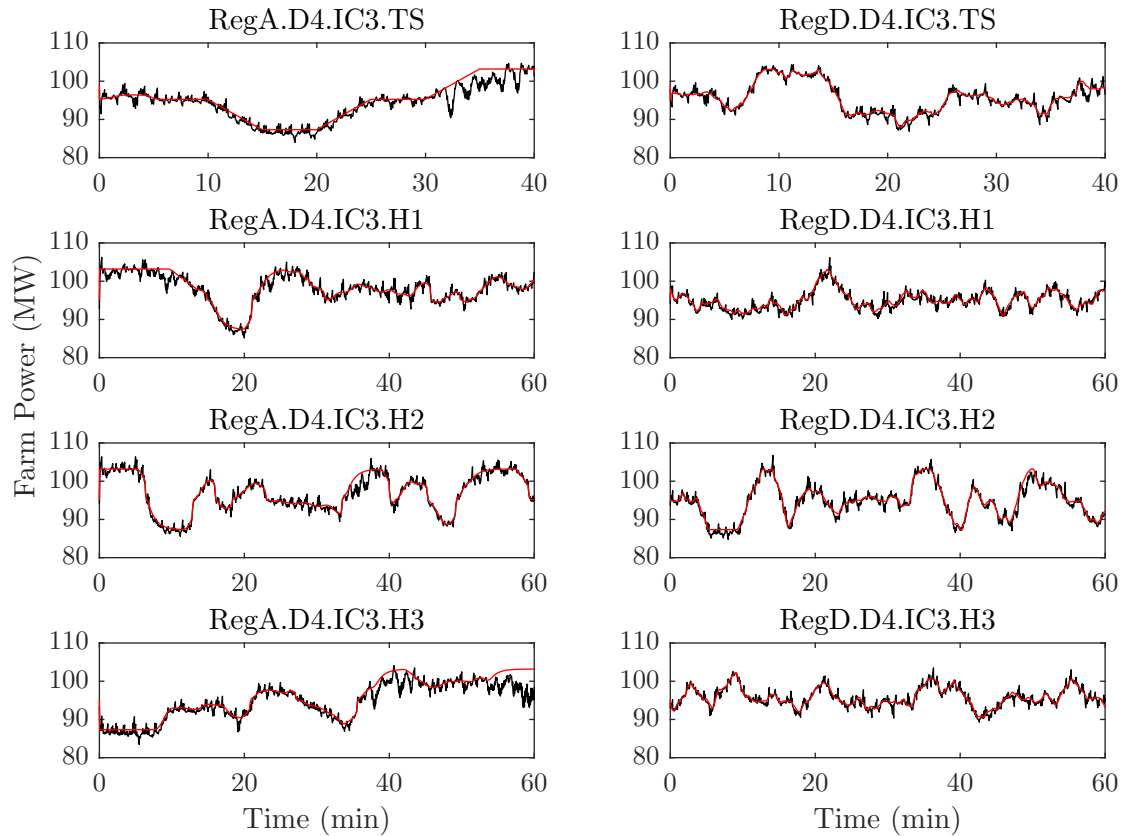
Figure 7. Power tracking performance of dynamic-model controlled wind farms comparing simulated farm power from a controlled LES wind farm model (——) and power reference signals (——) for 4% derates and initial condition 3.

The results shown in Figures 5–8 provide important insights into the possible strengths and limitations of the proposed approach to wind farm control for frequency regulation. These results suggest that wind farms may be well suited to act as a quickly responding resource for grid regulation services. The consistent passing of the composite performance score for the RegD signals indicates that dynamic-model controlled wind farms are able to provide this service reliably.

5   The power tracking results in Figure 7 demonstrate that the controller is able to track the up-regulation portions of the RegA signals at the beginning of the control period, such as during the first 5–10 minutes of the signals on 2/14 and 8/17. However, when up-regulation is requested for prolonged periods or towards the end of the control interval, such as the last 10 minutes of the "RegA.D4.IC3.TS" and "RegA.D4.IC3.H3" cases, the controller does not perform as well. A possible explanation is that the available energy in the wind is slowly changing as the atmospheric boundary layer evolves. Since estimates of available

10   energy are readily available over short time horizons, more frequent market clearing may allow wind farms to more effectively provide regulation. Ultimately, future work is needed to determine whether this is a fundamental limitation of the wind farm dynamics or the control strategy.
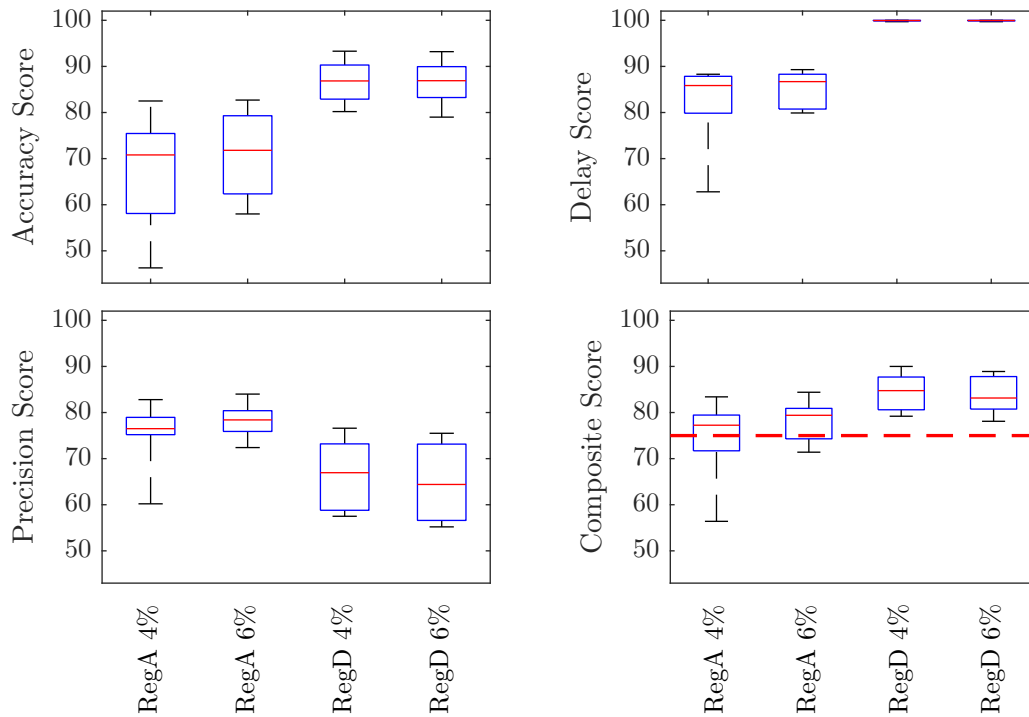
Figure 8. Boxplots of PJM performance scores for dynamic-model controlled wind farm for all regulation signal types (RegA or RegD) with derate values of 4% and 6%. The qualification threshold of 75% for the composite score (— — —) is shown in the lower right panel. The average controlled system performance exceeds the 75% for both signal types, but only the RegD cases pass all of the time.

8 Conclusions

In this study we further characterize the performance of wind farms providing secondary frequency regulation using the dynamic-model control framework proposed in Shapiro et al. (2017). This model-based receding horizon approach relies on a simple one-dimensional time-varying wake model to provide thrust coefficient trajectories for individual turbines within a wind farm. As in previous work, the control approach is tested using a "virtual wind farm" represented by LES of an 84-turbine wind farm with turbines modeled as actuator disks.

First, we evaluate the relative importance of including the dynamics of wake advection in the control scheme by comparing the performance of the dynamic-model controller to a comparable static-model controller. Tests using regulation signals from PJM indicate that the dynamic-model control demonstrates good overall tracking performance, whereas static-model control failed to match the reference signal for all simulated cases. These results indicate that the complexity of including the dynamics of wake advection is indeed required in model-based coordinated wind farm controls.

The tracking performance of the dynamic-model control method is then further quantified using PJM's performance metrics. Tests for both regulation signal types, RegA and RegD, exceed the PJM threshold for regulation participation on average, but the only the RegD signal exceeds the threshold in all cases. These results indicate that this model-based receding horizon controller design could allow wind farms to meet industry design standards and allow wind farms to fully participate in regulation markets, particularly in fast-acting regulation markets.

The potential for reducing the derate required to participate in these regulation markets was also explored. Participating in frequency regulation markets currently requires a trade-off between revenue losses in the bulk power market and revenues generated in the frequency market. Previous approaches (Aho et al., 2013; Jeong et al., 2014) required power setpoint reductions of an amount equal to the regulation amount, which directly reduces bulk power revenue by this amount. For this study we took a more aggressive approach by reducing the power setpoint by only 75%, and even only 50%, of the maximum regulation provided. For both of these derates, the controller is able to track fast-acting RegD signals. The potential for reducing the required derate has important economic implications for wind farms participating in both energy and regulation markets, a situation that will become increasingly common as more ISOs require wind farms to contribute to this grid service.

Although the dynamic-model controller design showed promising results, more work is needed to push this approach towards the implementation phase. We used the local thrust coefficient as a surrogate for real turbine control variables, such as generator torque and blade pitch angle. Improvements to our representation of these variables through actuator line methods and the inclusion of drivetrain dynamics in the control method are needed. Including rotational inertia may allow for further reductions in the amount of derate because rotational kinetic energy can compensate for short term power shortages (De Rijcke et al., 2015). Furthermore, we assumed that the ISO provided the regulation signal at the beginning of the control period; however, PJM provides this reference at a 2 second scan rate. This shortcoming could be addressed by adding estimated reference trajectories to the control design.

9 Data availability

Data sets containing the output files from the simulations will be provided with the final paper.

References

Aho, J., Buckspan, A., Laks, J., Fleming, P., Jeong, Y., Dunne, F., Churchfield, M., Pao, L., and Johnson, K.: A tutorial of wind turbine control for supporting grid frequency through active power control, in: American Control Conference, pp. 3120–3131, Montreal, Canada, doi:10.1109/ACC.2012.6315180, 2012.

5 Aho, J., Buckspan, A., Pao, L., Fleming, P., et al.: An active power control system for wind turbines capable of primary and secondary frequency control for supporting grid reliability, in: AIAA Aerospace Sciences Meeting, Grapevine, TX, 2013.

Annoni, J., Howard, K., Seiler, P., and Guala, M.: An experimental investigation on the effect of individual turbine control on wind farm dynamics, Wind Energy, 19, 1453–1467, doi:10.1002/we.1930, 2016.

Bewley, T. R., Moin, P., and Temam, R.: DNS-based predictive control of turbulence: An optimal benchmark for feedback algorithms, Journal

10 of Fluid Mechanics, 447, 179–225, 2001.

Borzì, A. and Schulz, V.: Computational Optimization of Systems Governed by Partial Differential Equations, SIAM, Philadelphia, PA, 2011.

Bou-Zeid, E., Meneveau, C., and Parlange, M.: A scale-dependent Lagrangian dynamic model for large eddy simulation of complex turbulent flows, Physics of Fluids, 17, doi:10.1063/1.1839152, 2005.

Calaf, M., Meneveau, C., and Meyers, J.: Large eddy simulation study of fully developed wind-turbine array boundary layers, Physics of

15 Fluids, 22, 2010.

De Rijcke, S., Driesen, J., and Meyers, J.: Power smoothing in large wind farms using optimal control of rotating kinetic energy reserves, Wind Energy, 18, 1777–1791, doi:10.1002/we.1790, 2015.

Díaz-González, F., Hau, M., Sumper, A., and Gomis-Bellmunt, O.: Participation of wind power plants in system frequency control: Review of grid code requirements and control methods, Renewable and Sustainable Energy Reviews, 34, 551–564, doi:10.1016/j.rser.2014.03.040,

20 2014.

Fleming, P., Aho, J., Gebraad, P., Pao, L., and Zhang, Y.: Computational fluid dynamics simulation study of active power control in wind plants, in: American Control Conference, pp. 1413–1420, Boston, MA, doi:10.1109/ACC.2016.7525115, 2016.

Gebraad, P. M. O., Fleming, P. A., and van Wingerden, J. W.: Wind turbine wake estimation and control using FLORIDyn, a control-oriented dynamic wind plant model, in: American Control Conference, pp. 1702–1708, Chicago, IL, doi:10.1109/ACC.2015.7170978, 2015.

25 Goit, J. P. and Meyers, J.: Optimal control of energy extraction in wind-farm boundary layers, Journal of Fluid Mechanics, 768, 5–50, 2015.

Jeong, Y., Johnson, K., and Fleming, P.: Comparison and testing of power reserve control strategies for grid-connected wind turbines, Wind Energy, 17, 343–358, doi:10.1002/we.1578, 2014.

Katić, I., Højstrup, J., and Jensen, N.: A simple model for cluster efficiency, in: European Wind Energy Association Conference and Exhibition, pp. 407–410, Rome, Italy, 1986.

30 Meyers, J. and Meneveau, C.: Large eddy simulations of large wind-turbine arrays in the atmospheric boundary layer, in: 50th AIAA Aerospace Sciences Meeting, Orlando, FL, 2010.

Moré, J. J. and Thuente, D. J.: Line search algorithms with guaranteed sufficient decrease, ACM Transactions on Mathematical Software, 20, 286–307, 1994.

PJM: PJM Manual 11: Energy & Ancillary Services Market Operations, 2012.

35 PJM: PJM Manual 12: Balancing Operations, 2015.

PJM: Ancillary Services, http://www.pjm.com/markets-and-operations/ancillary-services.aspx, 2016.

Press, W. H.: Numerical recipes 3rd edition: The art of scientific computing, Cambridge University Press, Cambridge, UK, 2007.

Rebours, Y., Kirschen, D., Trotignon, M., and Rossignol, S.: A Survey of Frequency and Voltage Control Ancillary Services—Part I: Technical Features, IEEE Transactions on Power Systems, 22, 350–357, doi:10.1109/TPWRS.2006.888963, 2007.

Rose, S. and Apt, J.: The cost of curtailing wind turbines for secondary frequency regulation capacity, Energy Systems, 5, 407–422, doi:10.1007/s12667-013-0093-1, 2014.

5   Shapiro, C. R., Bauweraerts, P., Meyers, J., Meneveau, C., and Gayme, D. F.: Model-based receding horizon control of wind farms for secondary frequency regulation, Wind Energy, 2017.

Stevens, R. J. A. M., Gayme, D. F., and Meneveau, C.: Large eddy simulation studies of the effects of alignment and wind farm length, Journal of Renewable and Sustainable Energy, 6, 023105, doi:http://dx.doi.org/10.1063/1.4869568, http://scitation.aip.org/content/aip/journal/jrse/6/2/10.1063/1.4869568, 2014a.

10   Stevens, R. J. A. M., Graham, J., and Meneveau, C.: A concurrent precursor inflow method for Large Eddy Simulations and applications to finite length wind farms, Renewable Energy, 68, 46–50, doi:10.1016/j.renene.2014.01.024, 2014b.

VerHulst, C. and Meneveau, C.: Altering Kinetic Energy Entrainment in Large Eddy Simulations of Large Wind Farms Using Unconventional Wind Turbine Actuator Forcing, Energies, 8, 370–386, 2015.