Scalable Minimum Fatigue Control of Dispatchable Wind Farms*

Daria Madjidian†

December 28, 2015

Abstract

As wind energy penetration increases, wind power plants may be required to regulate their power production according to the load-balancing needs of the power system. This presents an opportunity: when a wind farm tracks a power set-point, its wind turbines are free to continuously vary their power production as long as the sum of their productions meets the power demand. Here, we present an intuitive wind turbine coordination policy that uses this flexibility to minimize the aggregate fatigue load on the turbines. An important property is that the policy is scalable enough to be applied to any wind farm size. Specifically, the computational effort required to compute and reconfigure the optimal coordination policy is the same as that of a single stand-alone turbine, and the only centralized information processing needed to implement it is a single averaging operation. The efficiency of the coordination policy is illustrated in a simulation study based on real wind farm data.

1 Introduction

In order to reduce emissions, many regions have set goals on renewable energy production. In particular, Europe aims to produce 20% of its electricity from renewable sources by 2020 [1], the United States is looking into 20% wind power penetration by 2030 [2], and Denmark aims to have all of its electricity supplied by renewables in 2050 [3]. However, unlike conventional generation, the intermittent nature of renewable sources, such as wind and solar, poses a significant challenge to the load-balancing mechanism of the power system. Among the several renewable-integration measures that are being actively explored, one option is to let wind farms behave as dispatchable power plants that regulate their power production according to the balancing needs of the electrical power system [4, 5, 6]. This policy has already been translated into practice. For instance, several regions have grid codes where large wind farms are required to respond to power demands from the system operator [7, 8, 9]. Moreover, in the United Kingdom and Spain wind power plants participate in electricity markets where they bear full responsibility for forecasting and balancing their own production [10].

While operating a wind farm as controllable power plant leads to reduced energy production, it also offers an opportunity to mitigate this cost. When a wind farm tracks a power set-point, there is flexibility in distributing the total power production among its wind turbines. This can be used to improve additional aspects of wind power plant operation. For instance, in [11], the flexibility in distributing the power production was used to reduce active power losses in the transformers and lines inside the wind farm. Another possibility, which is the topic of this paper, is to use the flexibility to reduce the structural loads on the wind turbines [12, 13]. Load reduction is an important objective in wind turbine control because it can increase wind turbine life span and reduce maintenance costs [14]. Alternatively, reduced loads might relax the structural rigidity requirements of the wind turbine components and reduce material costs [15].

The underlying idea in this paper is that, instead of each turbine following a fixed portion of the power demand, it can be allowed to continuously adjust its power production in response to local wind speed fluctuations. Since wind conditions are not uniform across the wind farm, wind turbines can coordinate their power production so that changes that benefit one turbine are compensated for by units with opposite needs. Previous studies suggest that, by allowing wind turbines to coordinate their production in this way, it is possible to significantly reduce loads to both the tower and the low-speed shaft of the turbines [12, 13]. The mechanisms that convert the ability to adjust

*This work was supported by the Swedish Research Council through the LCCC Linnaeus Center and by the European commission through the project AEOLUS.

†Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA. Email: dariam@mit.edu

power production into load reduction were investigated in [16]. For work on a related idea, where the objective is to find wind turbine operating points that minimize fatigue loading, see [17].

The main challenge in coordinating the production of a large number of wind turbines is that the associated control problem becomes complex. The problem can be formulated as a cooperative control problem among a group of autonomous subsystems that are coupled by a linear coordination constraint on their inputs. The need to continuously reconfigure the controller to account for changes in wind conditions, external power demands, and the number of wind turbines in operation, sets severe limitations in terms of computational complexity of the control. In [13], computational complexity was reduced by using a distributed control architecture, where the wind turbines update and execute their own local control laws based on information from a limited number of neighboring units. In [18], a model predictive control approach was used to coordinate the turbines. The method exploits the structure of the problem in order to allocate the bulk of computations among the individual turbines, resulting in a speed up of the computation time by several orders of magnitude.

In this paper, we cast the wind turbine coordination problem in classical linear-quadratic-Gaussian (LQG) control framework. The main contribution is to show that the optimal solution is inherently scalable, both in terms of implementation and the computational effort required to obtain it. Our approach is based on the observation that wind turbines in wind farms are often identical in their design and operate around similar wind conditions. The control design problem can then be formulated as a coordination problem among homogeneous agents, which was studied in [19]. The optimal controller for this problem has several useful properties in terms of wind farm control. First, the optimal coordination policy is simple and intuitive. Specifically, the optimal power adjustment at each turbine is comprised of two terms: a local adjustment that disregards the wind farm power-set point and tries to minimize the fatigue load of the wind turbine, and coordination term that corrects for deviations from the power-set point. The coordination term is the same for all turbines and simply equals the average of the local adjustments. As a result of its special form, the only centralized information processing required to implement the coordination policy is a single averaging operation. This can easily be performed for very large wind farms. Moreover, the optimal control law can be obtained by solving a single LQG-problem for a stand-alone turbine. This means that the computational effort required to design and reconfigure the controller is very low, regardless of the number of turbines. Finally, for large wind-farms, an interesting observation is that under the optimal coordination policy, there is practically no trade-off between meeting the power demand to the wind farm and minimizing fatigue loads on the wind turbines.

The rest of the paper is organized as follows. In Section 2 we describe the wind turbine and wind farm models. In Section 3 we formulate the control design problem. The optimal controller and its properties are discussed in Section 4. To illustrate the behavior of the controller, Section 5 provides a simulation study based on real wind farm data. Conclusions are provided in Section 6.

2 Modeling

2.1 Wind turbine model

We adopt a non-linear model of the NREL 5-MW wind turbine from [20]. A schematic overview of the model is depicted in Figure 1. The main non-linearities in the model are in the aerodynamics block, which is implemented as a static model based on the power and thrust coefficients, and in the generator model. The drive-train and tower models are linear and contain poorly damped resonant modes. The pitch actuator is modeled as a first-order linear servo system with an internal loop delay. The wind turbine controller manipulates the generator torque and the blade pitch angle in order to meet a prescribed power demand. It has three main modes of operation, usually called “operating regions”. See Figure 2. The first two modes are identical to those of the standard NREL base-line controller described in [21], whereas the third mode has been modified to allow the turbine to track a power set-point. The controller operates in the third (derated) mode if the wind turbine is capable of producing the power demand, that is, the demand does not exceed the power that can be captured by the turbine. In this case excess wind power is curtailed in order to satisfy demand. This is achieved by modifying the pitch angle to keep the rotor speed close to its rated value, and adjusting the generator torque to maintain the power set-point. In this paper, we assume that the controller operates in Region 3.

The full non-linear model will be used for simulation in Section 5. For control design, we follow standard control design procedure and use a linear model to approximate the wind turbine dynamics around an operating point. Here, we use the linear model in [12]. A block-diagram of the model is illustrated in Figure 3, where P

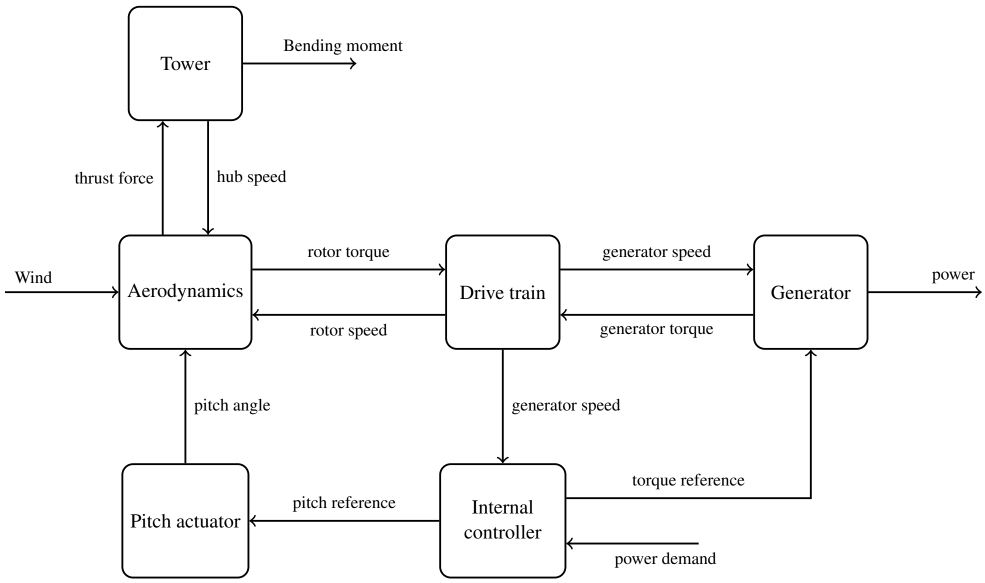

Figure 1: Schematic overview of the NREL wind turbine model

represents the wind turbine and P_v accounts for the spectral density of the wind speed fluctuations. The inputs to the wind turbine, v and p_{ref}, are the deviations in the wind speed and the power reference from their mean values, respectively. To illustrate that p_{ref} is the control input to the turbine, it is also denoted by u. It is important to note that the linear model neglects electrical generator dynamics, which means that u equals the actual power production of the turbine. The outputs F, ω and β are the deviations in thrust force, rotor speed and pitch angle from their nominal values, respectively. The wind speed is modeled as $v = P_v n$, where n is Gaussian white noise with unit intensity. The filter P_v was identified from real wind turbine measurements in [22] and is described by the transfer function

$$P_v(s) = \frac{7.4476}{(1/0.0143)s + 1},$$

2.2 Wind farm model

We consider a wind farm with N wind turbines, and use the subscript i to refer to ith turbine in the wind farm. For instance, β_i denotes the pitch angle deviations at the ith wind turbine. We make two simplifying assumptions, which are discussed in Remark 4.3:

- The operating point is the same for all wind turbines. The operating point is determined by the mean wind speed and the nominal power production, which are set to 10 m/s and 2 MW, respectively. Since all wind turbines are identical in their design, the assumption implies that they have identical dynamics.
- The wind speed variations experienced by different turbines are uncorrelated. This assumption is only made to simplify the control design and is not used in the simulations in Section 5.

Based on the assumptions above, the wind turbines in the wind farm can be described by

$$\frac{d}{dt}x_i(t) = Ax_i(t) + B_u u_i(t) + B_n n_i(t), \quad i = 1, \ldots, N, \tag{1}$$

where n_i and n_j are independent for $i \neq j$. The state vector is given by $x_i = \begin{bmatrix}\beta_i & \omega_i & q_i & v_i\end{bmatrix}^T$, where q_i is a filtered rotor speed measurement that is used by the internal wind turbine controller. At the operating point defined above, the state space matrices take the following values

$$A = \begin{bmatrix} 0 & 1.2 \cdot 10^{2} & -9.2 \cdot 10^{-1} & 0 \\ -8.4 \cdot 10^{-3} & -3.2 \cdot 10^{-2} & 0 & 1.6 \cdot 10^{-2} \\ 0 & 1.5 \cdot 10^{2} & -1.6 & 0 \\ 0 & 0 & 0 & 1.43 \cdot 10^{-2} \end{bmatrix} \qquad B = \begin{bmatrix} 0 \\ -2.1 \cdot 10^{-8} \\ 0 \\ 0 \end{bmatrix} \qquad B_n = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0.11 \end{bmatrix}.$$

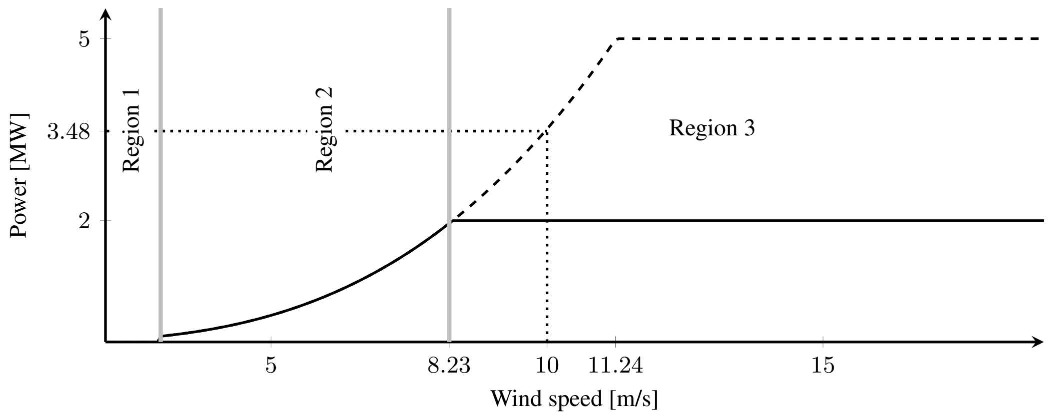

Figure 2: Power curve of the NREL 5-MW wind turbine. The turbine operates in Region 3 if the power set-point is below the power that the turbine is capable of producing (dahsed). The solid curve shows the power production at different wind speeds for a power set-point of 2 MW. At 10 m/s, the turbine spills 1.48 MW of the available power to satisfy the demand. The solid gray lines indicate transitions between operating regions. With a power reference of 2 MW, the wind turbine operates in Region 3 if the wind speed exceeds 8.23 m/s.

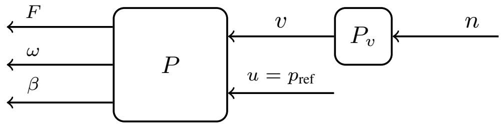

Figure 3: Block-diagram of the linear wind turbine model. The plant P describes the wind turbine and P_v accounts for the spectral density of the wind speed fluctuations. The signal v represents the wind speed fluctuations, u is the deviation of the power reference from its nominal value, and n is white noise that generates v. The outputs F, ω and β are the deviations in thrust force, rotor speed and pitch angle from their nominal values, respectively. The model neglects generator dynamics, which makes u the actual power production of the turbine.

3 Formulation of the optimal control problem

As explained above, we consider a wind farm with N wind turbines, where each turbine operates at a mean wind speed of 10 m/s and a nominal power production of 2 MW. The power set-point to the wind farm is the sum of the nominal wind turbine production, i.e. $2N$ MW.

The aim of the wind farm controller is to use the flexibility in distributing the power among the wind turbines to reduce structural loads. Here, we limit our attention to fatigue loads on the tower of the turbines[1]. The standard method for evaluating fatigue, which is used during the simulations in Section 5, is based on counting stress cycles [15]. However, this fatigue measure cannot be directly addressed in a control design framework based on a linear model. Therefore, we will design a control law that reduces the variance of the thrust force. The idea is that, since the tower motion is mainly driven by the thrust force, a reduction in thrust should lead to less excitation and hence, lower fatigue loading.

We define the cost function for the ith turbine as

$$\mathcal{J}_i = \mathbf{E}\left(z_i^T(t) z_i(t)\right).$$

where $\mathbf{E}$ denotes expectation and

$$\begin{aligned} z_i(t) &= \begin{bmatrix} F_i(t) \\ k_u u_i(t) \end{bmatrix} = \begin{bmatrix} \text{-}5.8\cdot 10^4 & \text{-}1.5\cdot 10^5 & 0 & 7.4\cdot 10^4 \\ 0 & 0 & 0 & 0 \end{bmatrix} x_i(t) + \begin{bmatrix} 0 \\ k_u \end{bmatrix} u_i(t) \\ &= C_z x_i(t) + D_z u_i(t). \end{aligned}$$

Note that we have included u_i in the cost function in order to keep the power variations at each turbine at a reasonable level. Here, the parameter $k_u > 0$ is a weight that can be used to tune the level of power fluctuations that can be accepted at a turbine.

For each wind turbine, we assume that we have access to noisy measurements of the pitch angle and the rotor speed

$$y_i = \begin{bmatrix} \beta_i + k_\beta e_i^\beta \\ \omega_i + k_\omega e_i^\omega \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} x_i + \begin{bmatrix} k_\beta & 0 \\ 0 & k_e \end{bmatrix} \begin{bmatrix} e_i^\beta \\ e_i^\omega \end{bmatrix} = C_y x_i + D_y \begin{bmatrix} e_i^\beta \\ e_i^\omega \end{bmatrix} \tag{2}$$

where e_i^β and e_i^ω are mutually independent Gaussian white noise processes with unit intensities, and $k_\beta > 0$ and $k_\omega > 0$ are design parameters that can be used to account for different signal-to-noise ratios. It is also possible to include the internal controller state, q_i, in the measurement, but since this is only a filtered rotor speed measurement, it does not contain any additional information.

If the wind turbines were able to freely vary their power production, each of their control laws could be obtained by solving the following decentralized optimal control problem

$$\begin{aligned} &\text{minimize} \quad \mathcal{J}_i && \text{(3a)} \\ &\text{subject to} \quad \text{(1) and (2).} && \text{(3b)} \end{aligned}$$

This is a standard linear quadratic Gaussian (LQG) control problem and its solution is well-known [24]. To form the solution, it is necessary to find the satbilizing solutions to the following algebraic Riccati equations

$$A^T X + XA + C_z^T C_z - (XB_u + C_z^T D_z) R_u^{-1} (B_u^T X + D_z^T C_z) = 0 \tag{4}$$

$$AY + YA^T + B_w B_w^T - (YC_y^T + D_y B_w^T) R_w^{-1} (C_y Y + B_w D_y^T) = 0, \tag{5}$$

where $R_u = D_z^T D_z$ and $R_w = D_y D_y^T$. The solutions to these equations are said to be stabilizing if

$$F = -R_u^{-1}(B^T X + D_z C_z) \tag{6}$$

$$L = -(YC_y^T + B_w D_y^T) R_w^{-1} \tag{7}$$

[1]To keep the presentation clean, we only include fatigue loads on one component. We have chosen the tower because it constitutes the most expensive component of a wind turbine [23]. However, the framework in this paper does not depend on this choice, and fatigue loads on other components, such as the low-speed shaft, can be included in a straightforward manner.

are such that the matrices $A + BF$ and $A + LC$ are stable[2]. The solution to (3), denoted u_i^{loc}, is then given by the observer-based control law

$$\frac{d}{dt}\hat{x}_i(t) = (A + BF)\hat{x}_i + L(C\hat{x}_i - y_i) \tag{8a}$$

$$u_i^{\text{loc}}(t) = F\hat{x}_i(t). \tag{8b}$$

Since the wind farm is required to meet a certain power demand, the wind turbines cannot apply (8) directly. Instead, they must coordinate their power production to ensure that the set-point is met. Since u_i is the power deviation from the nominal production of each turbine, the power demand is satisfied if $\sum_{i=1}^{N} u_i = 0$. The overall control problem is defined as

$$\text{minimize} \quad \mathcal{J} = \frac{1}{N}\sum_{i=1}^{N} \mathcal{J}_i \tag{9a}$$

$$\text{subject to} \quad (1) \text{ and } (2) \tag{9b}$$

$$\text{and} \quad \sum_{i=1}^{N} u_i = 0. \tag{9c}$$

Remark 3.1. *The LQG control law* (8) *has an interesting and intuitive interpretation. First, it can be shown that if the full state x_i could be measured perfectly, the optimal control law would be $u_i = Fx_i$. However, since x_i cannot be measured, the optimal control law in* (8) *replaces it with its estimate, $\hat{x}_i$. The term $(A + BF)\hat{x}_i$ in the state estimator equation* (8a) *takes into account the dynamics of the wind turbine, and last term incorporates new information, $y_i - C_y\hat{x}_i$, obtained from measurements (i.e. the innovations). It can be shown that when the innovation weight, L, is chosen as in* (7)*, then $\hat{x}_i$ is the best least squares estimate of x_i. The property that the LQG control law can be obtained by solving separate state-feedback control and estimation problems is known as the* separation principle of estimation and control *[24, Section 14.9].*

4 Optimal control law

It can be shown[3] that the optimal solution to (9) is given by

$$u_i(t) = u_i^{\text{loc}}(t) - \frac{1}{N}\sum_{j=1}^{N} u_j^{\text{loc}}(t), \quad i = 1, \ldots, N, \tag{10}$$

where u_i^{loc}, defined by (8), is the optimal control law for the local problem (3). This means that the optimal control policy is for each wind turbine to behave without concern for the power set-point of the wind farm, and then compensate for a portion, $\frac{1}{N}$, of the resulting deviation. The structure of the control law is illustrated in Figure 4 for a wind farm with two turbines. The turbines have their own local controllers, which compute $u_1^{\text{loc}}(t)$ and $u_2^{\text{loc}}(t)$ based on measurements y_1 and y_2, respectively. If these local control signals were applied to the wind turbines directly, they would result in a deviation, $u_1^{\text{loc}}(t) + u_2^{\text{loc}}(t)$, from the power-set point of the wind farm. This is remedied by a coordinator that computes correction terms,

$$-\frac{1}{2}(u_1^{\text{loc}}(t) + u_2^{\text{loc}}(t)),$$

which are added to the decentralized control signals.

An important property of the optimal control law is that it scales well as the the number of turbines grows large. In order to obtain (10), we only need to compute the local feedback and observer gains F and L in (8), which

[2]Stabilizing solutions to (4) and (5) can be computed using standard computer programs, e.g. by using the command `are` in Matlab

[3]A state-feedback version of (9) was solved in [19]. The solution can be extended to the output-feedback case by using the separation principle [24, Section 14.9].

amounts to solving problem (3) for a single turbine. This means that the controller can be easily reconfigured to account for changes in the mean wind speed, power set point, and the number of units in operation. Moreover, the only information processing required in addition to the local control laws, is a single averaging operation to form the correction terms, which can easily be performed for a very large number of turbines.

If the controller is implemented as suggested in Figure 4, the local control laws will be independent of the number of turbines in the farm. This provides a degree of robustness against failures. For instance, if communication fails between a wind turbine and the coordinator, meaning that a correction term cannot be transmitted, the wind turbine can revert to its nominal power production, $u_i = 0$. Coordinating the remaining wind turbines in an optimal manner is then simply a matter of the coordinator updating the number of actively participating units $N \to (N-1)$. The architecture in Figure 4 is also useful when the wind farm operates in Δ-mode, where the objective is to track a portion of the available wind power as opposed to a predetermined power set-point [25]. Wind turbines can then easily be added or removed from operation without any redesign of the local control laws. Only N needs to be updated to reflect the number in operation.

A natural question is how much better a wind turbine performs under the optimal control law (10) compared to when it maintains its nominal production, i.e. under $u_i = 0$. This is partly investigated in a simulation study in Section 5. In the remainder of this section, we provide some general insight based on results in [19].

Let $\mathcal{J}_{\text{loc}}$ and $\mathcal{J}_0$ denote the values of the wind turbine cost function $\mathcal{J}_i$ under the local control laws $u_i = u_i^{\text{loc}}$ and $u_i = 0$, respectively. Note that $\mathcal{J}_{\text{loc}}$ is the cost of a turbine when it "greedily" minimizes its fatigue loads without any regard for the wind farm power-set point, and $\mathcal{J}_0$ is the cost when the wind turbine is not allowed to adjust its power production at all. It is immediate that $\mathcal{J}_{\text{loc}} \leq \mathcal{J}_0$. The value of the cost function of a wind turbine under the optimal coordination policy (10) can then be expressed as

$$\mathcal{J}_i = \mathcal{J}_0 - (1 - \frac{1}{N})(\mathcal{J}_0 - \mathcal{J}_{\text{loc}}). \tag{11}$$

This expression states that the performance improvement obtained by an average turbine as a result of coordinating power production is $(1 - \frac{1}{N})(\mathcal{J}_0 - \mathcal{J}_{\text{loc}})$. This quantity increases with the number of units in the wind farm. This is intuitive as a large number of turbines provides better opportunity for matching positive and negative power adjustments needs. In particular, $\mathcal{J}_i \to \mathcal{J}_{\text{opt}}$ as the number of turbines $N \to \infty$, which means that, in large wind farms, there is practically no trade-off between satisfying the power demand and reducing fatigue loads.

Remark 4.1. *The theory in [19] allows for several modifications to the optimization problem* (9)*, which can be used to account for additional requirements. For instance, it is possible to prioritize the wind turbines differently by replacing objective* (9a) *by* $\sum_i \lambda_i \mathcal{J}_i$*, where a high value on* $\lambda_i > 0$ *would reflect a high priority for reducing loads on the ith turbine. Also, instead of aiming for perfect tracking of the wind farm power set-point, it is possible to consider a trade-off between tracking and load reduction. It is also possible to incorporate frequency-weighted tracking requirements. For instance, it is possible to suppress slow deviations from the wind farm power set-point while allowing for faster fluctuations. The solution to these modified problems can be extracted from [19]. Importantly, these modifications do not alter the structure of the optimal control law in* (10)*. That is, local wind turbine control is corrected by a single centralized (possibly weighted) averaging operation.*

Remark 4.2. *The quantities* $\mathcal{J}_{loc}$ *and* $\mathcal{J}_0$ *can be computed as*

$$\mathcal{J}_{loc} = \text{tr}(B_w^T X B_w) + \text{tr}(R_u F Y F^T) \qquad \textit{and} \qquad \mathcal{J}_0 = \text{tr}(B_w^T X_0 B_w),$$

where X, Y*, and* F *are defined in* (4), (5), (6)*, respectively, and* X_0 *is the solution to the Lyapunov equation*

$$A^T X_0 + X_0 A + C_z^T C_z = 0,$$

which can be solved using standard computing programs[4]*. See [24] for details.*

Remark 4.3. *It is important to discuss the validity and implication of the assumptions introduced in Section 2.2. First, practical studies suggest that wind speed fluctuations (variations around mean wind speed) are only correlated at very low frequencies [26, 22]. Since our objective is to reduce fatigue loads, which are excited mainly by faster wind speed fluctuations, we do not expect any significant effect on our results due to correlation. Moreover, we stress that this assumption is only made to simplify the the control design, and is not used in the simulation*

[4]For instance, using `lyap` or `are` in as Matlab®

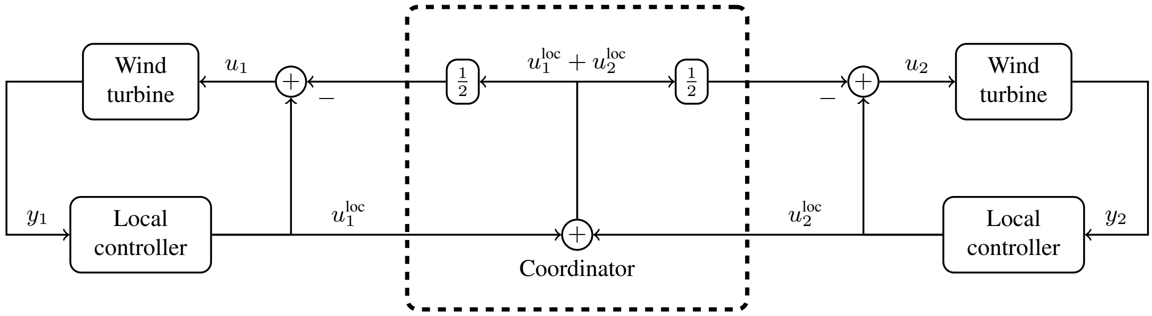

Figure 4: Structure of the optimal control law for a wind farm with two wind turbines. Each turbine has its own local controller, defined by (8), which communicates u_i^{loc} to a wind farm coordinator. The coordinator aggregates all local control signals and returns a correction term, $-\frac{1}{N}\sum u_j^{\text{loc}}$.

study in Section 5, where the wind speed fluctuations are estimated from real wind farm data, which is collected from adjacent wind turbines.

The assumption on identical operating points is more delicate, especially regarding the mean wind speeds. For instance, in wind farms with two dimensional topologies, some turbines will be exposed to wakes, which implies that they will experience lower mean wind speeds. Investigating the robustness of the control laws proposed in this work to non-identical operating points is an important research topic that is not covered in this work.

It is worth mentioning that the design framework proposed in this paper can be easily modified to account for different operating points. One approach would be to split the wind farm into smaller groups, where the wind turbines do have the same operating conditions, and solve (9) *for each group. Note that the power coordination requirement* (9c) *is satisfied for the wind farm if it is satisfied for each subgroup (but not vice versa). This approach is appealing because, according* (11)*, by coordinating groups of say* $N = 10$ *wind turbines we can expect to attain at least* $1 - 1/N = 90\%$ *of the load reduction obtained by coordinating all wind turbines. A second heuristic approach would be to apply* (10)*, and let each* u_i^{loc} *be the optimal control with respect to the ith turbine's specific operating conditions. An evaluation the performance and robustness aspects of this heuristic is outside the scope of this work.*

5 Simulation example

In this section we simulate[5] a wind farm with five NREL 5 MW wind turbines. We let each wind turbine experiences a mean wind speed of 10 m/s and a nominal power production of 2 MW, which is 1.48 MW below its nominal power production ability. See Figure 2. The power demand to the wind farm is 10 MW.

In the simulations, each wind turbine experienced a wind speed of 10 m/s $+\, v_i$. In order to provide realistic simulations, the wind speed variations, v_i, were estimated from real wind farm data. The data was collected from a commercial wind farm consisting of Vestas V90 3 MW wind turbines. During the data collection the turbines operated in Region 2 where they operate at maximum aerodynamic efficiency. In this mode of operation, the power production is approximately a cubic function of the wind speed. The effective wind speed at a wind turbine was estimated as $V_i = (P_i/c)^{\frac{1}{3}}$, were P_i is the power production of the turbine and the coefficient c was obtained by using the Vestas V90 power curve in [28]. Figure 6 shows one of these sequences together with the measured nacelle wind speed. Finally, the wind speed variation was computed from the effective wind speed by subtracting its mean, that is, $v_i = V_i - \text{mean}(V_i)$. The wind direction during the data collection is illustrated in Figure 5. It is parallel to the row of turbines, which implies that the estimated wind speed variations account for possible correlation effects, which were neglected in the control design. A total of nine different ten minute effective wind speed sequences were estimated for each wind turbine.

For each wind speed sequence, we simulated the response of the wind farm under the control law (10). The

[5]The simulations are performed in Simulink®, where we use the wind turbine model implementation in [27].

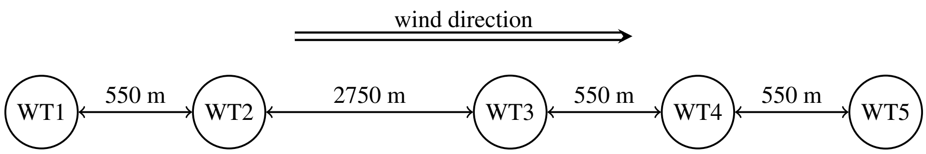

Figure 5: Position of the wind turbines in relation to the wind direction during the data collection. For each turbine, nine wind speed sequences of ten minutes were estimated. Since the estimated wind speeds are based on data collected from neighboring wind turbines, they account for possible correlation effects, which were neglected during the control design.

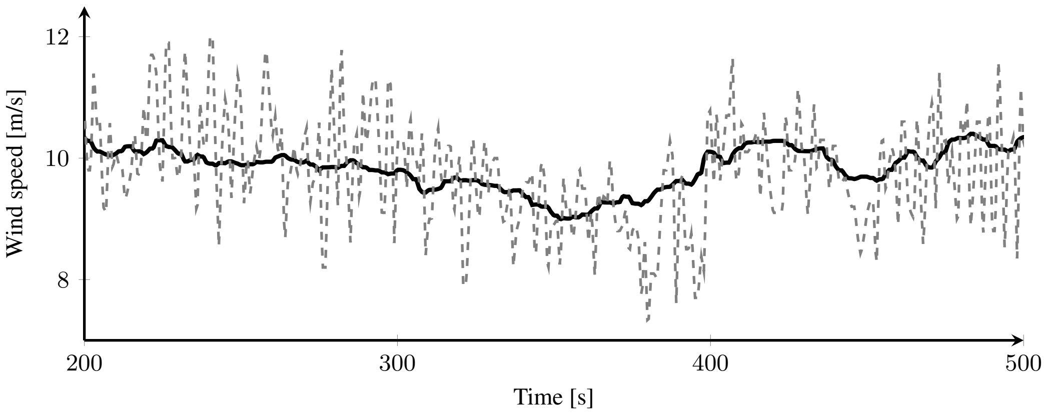

Figure 6: Estimated effective wind speed (solid black) and measured nacelle wind speed (dashed gray) at one of the turbines.

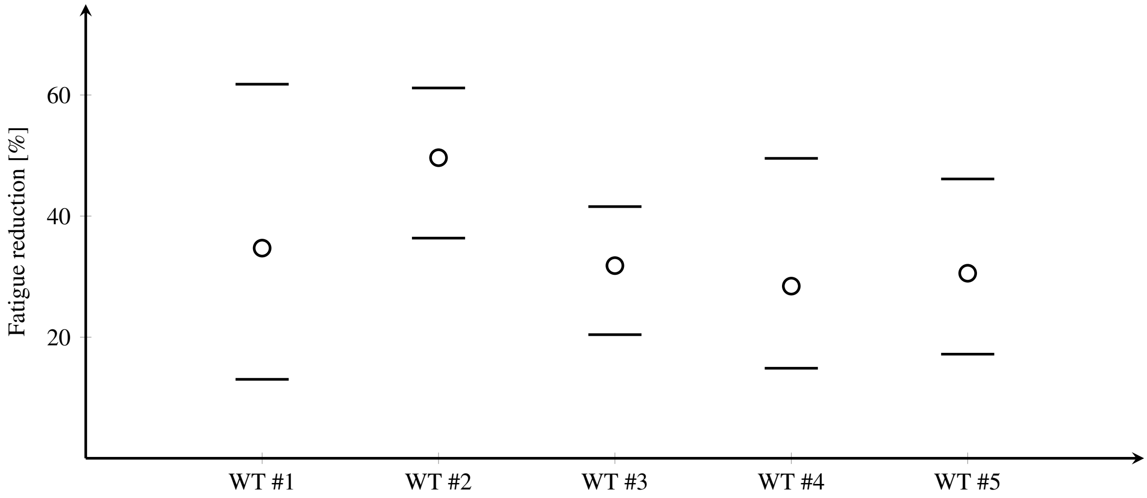

Figure 7: Reduction in accumulated fatigue (%) for each wind turbine when using the control law (10) compared to the case when the turbines maintain a power production of 2 MW at all times. The circles show the average reduction over all 9 simulation runs, while the lines indicate the smallest and largest reduction. The average reduction over all turbines and all simulations was 35%.

control design weights were set to[6]

$$k_u = 10^{-5} \quad k_\beta = 1.5 \quad \text{and} \quad k_e = 10^{-4}.$$

For comparison, we also simulated the case where the turbines do not deviate from their nominal power references, i.e. $u_i = 0,\ i = 1, \ldots, 5$. After each simulation, the accumulated fatigue to each of the wind turbine towers was computed as follows.

1. The rain-flow counting algorithm in [29] was used to extract equivalent stress cycles, $s_1, \ldots, s_r$, from the history of the fore-aft tower bending moment.

2. The fatigue caused by a stress cycle number l was computed based on the S/N-curve as s_l^k, where we use $k = 4$, which is typical for steel structures [15]. The total fatigue was then calculated according to Palmgren-Miner's rule as $\sum_{l=1}^{r} s_l^k$.

Figure 7 summarizes the outcome from the nine different simulations. On average, the accumulated fatigue damage was 35% lower when using the control law (10), compared to the case when the turbines maintained their nominal power production. Figure 8 shows the increase in accumulated fatigue during one of the simulations.

The response of one of the wind turbines to one of the wind speed sequences is shown is Figure 9. An illustrative period starts after approximately 300s, when the wind speed suddenly drops. To avoid exciting the tower bending dynamics, the local wind turbine controller tries to lower the power production, which is indicated by the drop in u_3^{loc}. The actual power production of the turbine, after applying the correction from the coordinator (see Figure 4), is shown by the solid curve. The effect of reducing the power production of the turbine is that a large peak in tower bending moment is avoided.

In Figure 9, the desired power adjustments of the wind turbine u_3^{loc} are in the same direction as the wind speed variations. This is a consistent phenomenon across all simulations and can be explained as follows. A drop in wind speed reduces the aerodynamic torque and has a decelerating effect on the rotor speed. In order to maintain rated speed, the internal wind turbine controller decreases the pitch angle, which in turn excites the tower bending dynamics. The benefit of reducing the power reference is that it helps the speed regulation by forcing the internal controller to reduce the generator torque. This reduces the required amount of pitch activity and leads to less tower bending. For a more extensive discussion on the mechanisms behind the load reduction, see [16].

[6]These values were chosen to obtain a satisfactory response for a single stand-alone wind turbine to a wind gust. Specifically, we solved the uncoordinated problem (3) repeatedly with different weights until a satisfactory response was obtained.

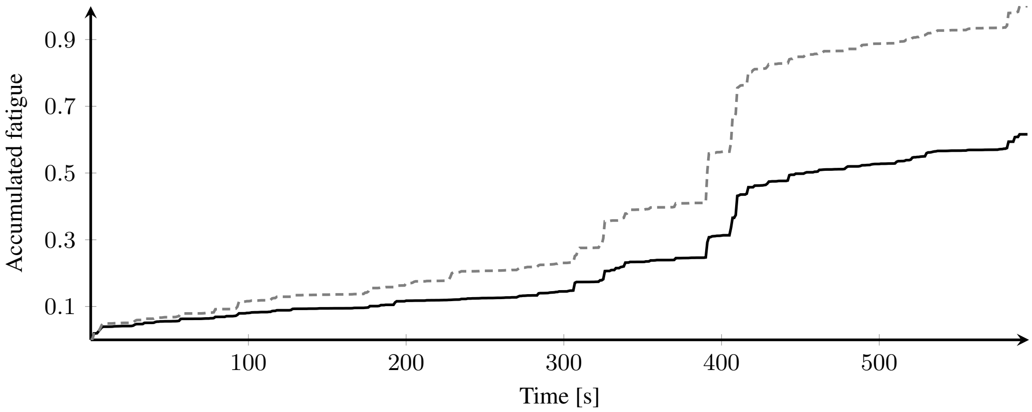

Figure 8: Normalized average accumulated fatigue for the five turbines during one of the simulations. With wind farm controller (solid) and when each wind turbine maintains a production of 2 MW (dashed gray).

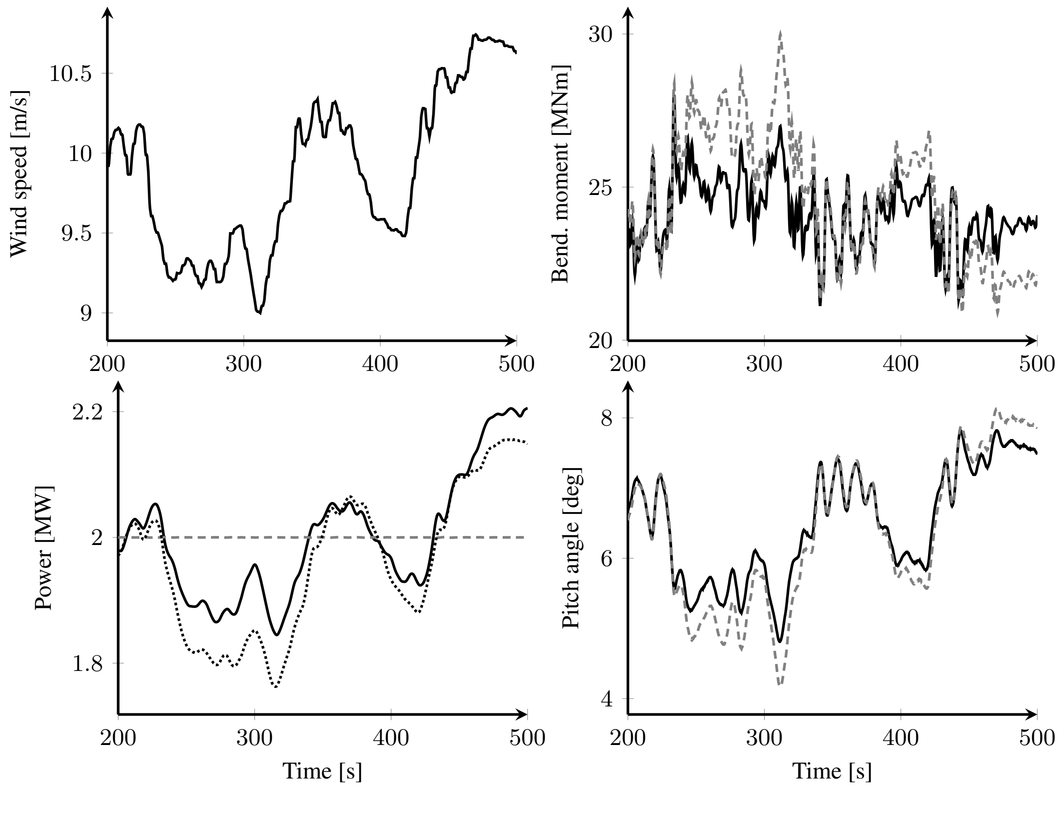

Figure 9: Response of wind turbine 3 to the wind speed in the upper left plot. With wind farm controller (solid) and when each turbine maintains a production of 2 MW (dashed gray). The dotted curve shows the desired power production u_3^{loc}. At time 300 s, the wind speed drops. To maintain the rated rotor speed, the internal wind turbine controller is required to reduce the pitch angle, which in turn excites the tower bending dynamics. In the case when the turbines maintain their nominal power productions, this leads to a large peak in tower bending moment. By allowing the turbine to reduce its power, the wind farm controller (10) manages to reduce the peak.

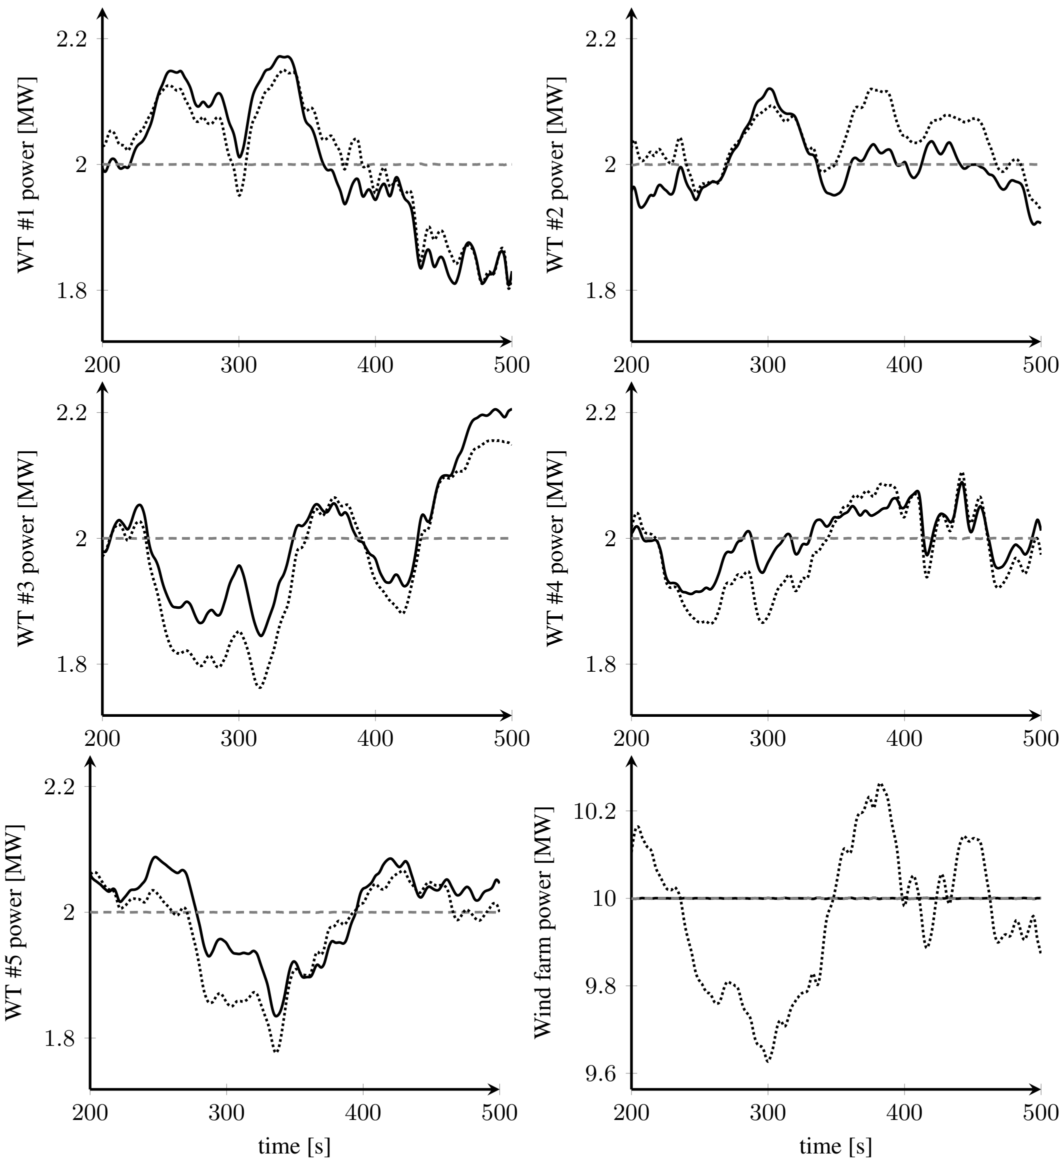

Figure 10: The first five plots show the power production for the wind turbines in one of the simulations. With the wind farm controller (10) (solid) and when each turbine maintains a production of 2 MW (dashed gray). The dotted curve shows the local control signals u_i^{loc}. The lower right plot, shows the total power production of the turbines. Note that the total power production under (10) is indistinguishable from the reference case. The dotted black curve shows the sum of the locally desired power references $\sum_{i=1}^{5} u_i^{\text{loc}}$.

Figure 10 shows the power production for all five wind turbines from one of the simulations. The actual power production follows the local control signals, u_i^{loc}, relatively well, which is a consistent phenomenon across all the simulations. The lower left plot shows the total power production in the wind farm. The production when using (10) is indistinguishable from the case when the wind turbines maintain their nominal power reference.

6 Conclusions

We have proposed a coordination policy for a wind farm that is required to meet a power set-point. The goal is to reduce the fatigue loads to the tower of the wind turbines, while satisfying the power demand. Our approach relied on two assumptions: that wind turbines are identical in their design and operating points, and that the wind speed variations that they experience are uncorrelated. Based on these assumptions, the control design problem was formulated as an optimal coordination problem among homogeneous agents, which was solved in [19].

The proposed wind farm controller has several important properties. First, it has an intuitive and transparent structure. Second, the optimal control law can be applied to any realistically sized wind farm. This is because in order to form the control law, it is only necessary to solve an optimization problem for a stand-alone wind turbine. Also, the only centralized operation required to implement the control law is a single averaging operation to form the correction terms, which can easily be performed for a large number of wind turbines. Moreover, it was possible to provide a simple expression that quantifies the load reduction due to the proposed coordination policy. In particular, for wind turbines operating in large wind farms, there is practically no trade-off between satisfying the power-set point and reducing fatigue loads.

The control law was evaluated in a simulation study based on wind speeds that were estimated from real wind farm data. It should be noted that, since the wind speeds were estimated using data collected from neighboring wind turbines, they account for possible correlation effects which were neglected during the control design. The proposed control law resulted in a 35% average reduction of the fatigue damage to the wind turbine towers compared to if each turbine had maintained its nominal power reference.

References

[1] European Comission. A European strategy for sustainable, competitive and secure energy 2006. URL `http://europa.eu/documents/comm/green\_papers/pdf/com2006\_105\_en.pdf`, COM(2006) 105 final.

[2] US Department of Energy. 20% wind energy by 2030: Increasing wind energy's contribution to U.S. electricity supply 2008. URL `http://www.nrel.gov/docs/fy08osti/41869.pdf`, DOE/GO-102008-2567.

[3] Danish Government. Our future energy 2011. URL `http://www.cphcleantech.com/media/1465822/our_future_energy.pdf`, ISBN 978-87-7844-915-3.

[4] Rodríguez-amenedo JL, Arnalte A, Burgos JC. Automatic generation control of a wind farm with variable speed wind turbines. *IEEE Transactions on Energy Conversion* 2002; **17**(2):279–284.

[5] Bitar E, Rajagopal R, Khargonekar P, Poolla K, Varaiya P. Bringing wind energy to market. *IEEE Transactions on Power Systems* 2011; **27**(3):1225–1235.

[6] Kirby B, Milligan M, Ela E. Providing minute-to-minute regulation from wind plants. *Proceedings of the 9th Annual International Workshop on Large-Scale Integration of Wind Power into Power Systems*, Quebec, Canada, 2010.

[7] Elkraft Systems and Eltra. Wind Turbines Connected to Grids with Voltages above 100 kV – Technical regulation for the properties and the regulation of wind turbines. *Technical Report Regulation TF 3.2.5* 2004.

[8] Hydro-Québec. Technical requirements for the connection of generation facilities to the hydro-québec transmission system – supplementary requirements for wind generation 2005. URL `http://www.hydroquebec.com/transenergie/fr/commerce/pdf/Exig_compl_eolien_transport_mai03_ang.pdf`, retrieved 2014-05-03.

[9] Commission for Energy Regulation. Wind farm transmission grid code provisions – a direction by the commission for energy regulation. *Technical Report CER/04/237* 2004.

[10] Klessmann C, Nabe C, Burges K. Pros and cons of exposing renewables to electricity market risks — a comparison of the market integration approaches in Germany, Spain, and the UK. *Energy Policy* 2008; **36**(10):3646–3661.

[11] de Almeida RG, Castronuovo E, Lopes JAP. Optimum generation control in wind parks when carrying out system operator requests. *IEEE Transactions on Power Systems* 2006; **21**(2).

[12] Spudić V, Jelavić M, Baotić M. Wind turbine power references in coordinated control of wind farms. *Automatika – Journal of Control, Measurements, Electronics, Computing and Communications.* 2011; **52**(2):82–94.

[13] Biegel B, Madjidian D, Spudić V, Rantzer A, Stoustrup J. Distributed low-complexity controller for wind power plant in derated operation. *Proceedings of the 2013 IEEE Multi-Conference on Systems and Control*, Hyderabad, India, 2013; 146–151.

[14] Pao L, Johnson K. Control of wind turbines. *IEEE Control Systems Magazine* 2011; **31**(2):42–62.

[15] Manwell J, McGowan J, Rogers A. *Wind Energy Explained*. 2 edn., John Wiley & Sons: Chichester, UK, 2009.

[16] Madjidian D, Kristalny M, Rantzer A. Dynamic power coordination for load reduction in dispatchable wind power plants. *Proceedings of the 2013 European Control Conference*, Zürich, Switzerland, 2013; 3554–3559.

[17] Juelsgaard M, Schiøler H, Leth J. Wind farm dispatch control for demand tracking and minimized fatigue. *IFAC Symposium on Power Plant and Power System Control*, Toulouse, France, 2012; 381–386.

[18] Spudić V, Baotić M. Fast coordinated model predictive control of large-scale distributed systems with single coupling constraint. *Proceedings of the 2013 European Control Conference*, 2013; 2783–2788.

[19] Madjidian D, Mirkin L. Distributed control with low-rank coordination. *IEEE Transactions on Control of Network Systems* 2014; **1**(1):53–63.

[20] Grunnet J, Soltani M, Knudsen T, Kragelund M, Bak T. Aeolus toolbox for dynamic wind farm modeling, simulation and control. *Proceedings of the 2010 European Wind Energy Conference*, 2010; 3119–3129.

[21] Jonkman J, Butterfield S, Musial W, Scott G. Definition of a 5-mw reference wind turbine for offshore system development. *Technical Report NREL/TP-500-38060*, National Renewable Energy Laboratory, Golden, CO 2009.

[22] Kristalny M, Madjidian D, Knudsen T. On using wind speed preview to reduce wind turbine tower oscillations. *IEEE Transactions on Control Systems Technology* 2013; **21**(4):1191–1198.

[23] Krohn S, Morthorst P, Awerbuch S. The economics of wind energy 2009. URL `http://www.ewea.org/fileadmin/ewea_documents/documents/00_POLICY_document/Economics_of_Wind_Energy_March_2009.pdf`.

[24] Zhou K, Doyle JC, Glover K. *Robust and Optimal Control*. Prentice-Hall: Englewood Cliffs, NJ, 1995.

[25] Aho J, Buckspan A, Laks J, Fleming P, Jeong Y, Dunne F, Churchfield M, Pao L, Johnson K. A tutorial of wind turbine control for supporting grid frequency through active power control. *In Proc. of the 2012 American Control Conference*, Montreal, Canada, 2012; 3120–3131.

[26] Vigueras-Rodríguez A, Sørensen P, Viedma A, Donovan M, Gómez Lázaro E. Spectral coherence model for power fluctuations in a wind farm. *Journal of Wind Engineering and Industrial Aerodynamics* 2012; **102**:14 – 21.

[27] SimWindFarm. Aeolus SimWindFarm Toolbox 2014. URL `http://www.ict-aeolus.eu/SimWindFarm/index.html`, retrieved 2014-04-16.

[28] Vestas V90. Vestas V90 brochure 2014. URL http://vestas.com/en/products_and_services/turbines/v90-3_0_mw, retrieved 2014-04-19.

[29] Hammerum K. A fatigue approach to wind turbine control. M.Sc. thesis, Technical University of Denmark 2006.