

MPC for Wind Power Gradients — Utilizing Forecasts, Rotor Inertia, and Central Energy Storage

Tobias Gybel Hovgaard, Lars F. S. Larsen, John Bagterp Jørgensen and Stephen Boyd

Abstract—We consider the control of a wind power plant, possibly consisting of many individual wind turbines. The goal is to maximize the energy delivered to the power grid under very strict grid requirements to power quality. We define an extremely low power output gradient and demonstrate how decentralized energy storage in the turbines' inertia combined with a central storage unit or deferrable consumers can be utilized to achieve this goal at a minimum cost. We propose a variation on model predictive control to incorporate predictions of wind speed. Due to the aerodynamics of the turbines the model contains nonconvex terms. To handle this nonconvexity, we propose a sequential convex optimization method, which typically converges in fewer than 10 iterations. We demonstrate our method in simulations with various wind scenarios and prices for energy storage. These simulations show substantial improvements in terms of limiting the power ramp rates (disturbance rejection) at the cost of very little power. This capability is critical to help balance and stabilize the future power grid with a large penetration of intermittent renewable energy sources.

I. INTRODUCTION

Today, wind power is the most important renewable energy source. For the years to come, many countries have set goals for further reduced CO₂ emission, increased utilization of renewable energy, and phase out of fossil fuels. In Denmark one of the means to achieve this is to increase the share of wind power to 50% of the electricity consumption by 2020 and to fully cover the energy supply with renewable energy by 2050 [1]. Installing this massive amount of wind turbines introduces several challenges to reliable operation of power systems due to the fluctuating nature of wind power. To mitigate fluctuations, modern wind power plants (WPP) are equipped with variable speed wind turbine (VSWT) technologies, which are interfaced with power electronics converters that are required and designed to fulfil increasingly demanding grid codes (see, e.g. [2], [3]).

The Grid Code (GC) is a technical document setting out the rules, responsibilities and procedures governing the operation, maintenance and development of the power system. It is a public document periodically updated with new requirements and it differs from operator to operator. Countries with large amount of wind power have issued dedicated GCs for its connection to transmission and distribution levels, focused mainly on power controllability and power quality [4], [5]. Particularly, Denmark establishes some of the most

demanding requirements regarding active power control [6]. One of the regulation functions required is a power gradient constraint that limits the maximum rate-of-change of non-commanded variations in the power output from the WPP to the grid. As of today, this constraint is softened if the power production in the WPP drops due to the lack of wind. This is merely out of necessity, and the GCs are expected to tighten further regarding this requirement. Ensuring slow power gradients reduces the risk of instability on the grid, allows the TSO time for counteracting the change, and improves the predictability of power output, enabling the WPP owner to put less conservative bids on the power market.

Energy storage strikes the major problems of wind power and joining energy storage with WPPs to smoothen variations and improve the power quality is not a new idea. In, e.g., [7]–[10] the benefits, economics, and challenges of using different means of storage, *i.e.*, batteries, hydrogen, flywheels etc., in combination with wind power are investigated. [11] uses a Lithium-iron-phosphate battery to achieve power forecast improvement and output power gradient reduction. However, the additional cost of batteries or other energy storages is usually the showstopper, at least as the market is today. In our previous works, we have shown how thermal capacity, *e.g.*, in supermarket refrigeration, can be utilized for flexible power consumption [12], [13]. It is very likely that such techniques (where the capacity is a bi-product of fulfilling another need) can play a major role instead of adding expensive technologies which have storage as their sole purpose. In the rest of this paper, we consider energy storage in general without distinguishing actual storage from flexible power consumption.

Traditionally, the rotor speed of modern wind turbines is controlled for tracking the tip-speed ratio (TSR = angular rotor speed \times rotor radius / wind speed) for maximum power extraction, constrained by the maximum rated speed. However, due to the inertia of the rotating masses in the turbine, there is a potential to improve the quality of the power output by actively letting the rotor speed deviate from the optimal setting. This might of course come at a cost of slightly reduced power output. In, e.g., [14], [15] turbine inertia is used for frequency response and power oscillation damping. In addition, a vast amount of works exist that address power optimization, fatigue load reduction and pitch control for individual turbines in the more traditional sense, *e.g.*, [16]–[18]. Some of these take optimization and model predictive control approaches to solve the problems and many rely on a known operating point (*e.g.*, local wind speed and power set-point) for deriving linearized models.

T. G. Hovgaard and L. F. S. Larsen are with Vestas Technology R&D, DK-8200 Aarhus N, Denmark. {togho, lfsla}@vestas.com

J. B. Jørgensen is with DTU Compute, Technical University of Denmark, DK-2800 Lyngby, Denmark. jbjjo@dtu.dk

S. Boyd is with the Information Systems Laboratory, Department of Electrical Engineering, Stanford University, USA. boyd@stanford.edu

Other works consider the control of large wind farms where the power extracted by upwind turbines reduces the power that is available from the wind and increases the turbulence intensity in the wake reaching other turbines (see, *e.g.*, [19]–[21]).

In [13], we demonstrate the appreciation of a sequential convex programming (SCP) approach [22] for a model predictive control problem, controlling the power consumption for commercial refrigeration with linear dynamics, convex constraints, and a nonconvex objective. Inspired by this, we now turn to the power producers' side of the grid and apply the same technique to a nonlinear wind turbine model. Our method, like sequential quadratic programming (SQP) [23], involves the solution of a sequence of (convex) quadratic programs (QPs), but differs very much in how the QPs are formed. In SQP, an approximation to the Lagrangian of the problem is used; the linearization required in each step can end up dominating the computation [24]. In our SCP method, the convexification step is quite straightforward.

We demonstrate how model predictive control using forecasts of the wind speed can ensure very low power gradients (*e.g.*, less than 5% of the rated power per minute). We do this with a central energy storage added to the WPP and show how we can utilize the inertia in the individual turbines to further improve this and minimize the extra storage capacity needed. Our method gives no guarantee in terms of convergence or optimality but is observed to perform well in practice. [25] uses convex optimization to operate a portfolio of electrical storage devices. In [26], we present a change of variables that renders the problem fully convex and demonstrate efficient closed-loop simulations with real wind data.

II. WIND POWER PLANT

In this section, we describe the dynamic model used for the WPP in the paper. The WPP can have a number of individual wind turbines arranged in a certain geographical topology and one central storage unit. We describe the simplified dynamics of rotational motion, the constraints of the system and the function reflecting the objective of operating the plant.

A. Wind Turbine Model

The WPP in the examples consists of turbines using the NREL 5MW model since this is openly available, but, could easily be substituted with any specific turbine model. The model is described in detail in, *e.g.*, [27], [28]. We simplify the model and derive the system equations as follows. Neglecting the shaft torsion, we describe the turbine itself by two dynamical states, the generator speed, ω_g , in *rad/s* and the generator torque, T_g , in *Nm*.

$$\dot{\omega}_g = \frac{1}{I_g + I_r/N^2} \left(\frac{T_r}{N} - T_g \right), \quad (1)$$

$$\dot{T}_g = \frac{1}{\tau_g} (T_{g,\text{ref}} - T_g), \quad (2)$$

where I_g and I_r are the inertias of generator and rotor respectively, N is the gear ratio, τ_g is the time constant of the generator and $T_{g,\text{ref}}$ is the torque set-point. The torque, T_r , delivered to the rotor by the wind is given by

$$T_r = \frac{1}{2} \rho A C_P(\lambda, \beta) \frac{v^3}{\omega_r},$$

where ρ is air density, A is swept area of the rotor, v is wind speed in *m/s*, ω_r is angular rotor speed in *rad/s*, and the coefficient of power, C_P , is a look-up table (see Fig. 1) derived from the geometry of the blades as a function of TSR and blade pitch angle (β) in degrees. TSR is defined as $\lambda = R\omega_r/v$, where R is the rotor radius in *m*. We use $\omega_r = \omega_g/N$ to eliminate ω_r and describe the power produced in the generator by

$$P_g = \eta_g T_g \omega_g,$$

where η_g is the generator efficiency.

B. Energy Storage Model

We use a simple integrator for illustrating the central energy storage and describe its state-of-charge (Q [J]) in discrete-time by

$$Q(t + T_s) = Q(t) + c(t)T_s, \quad (3)$$

where c is the charge rate in *W* and t denotes time. We assume that the energy storage is lossless. However, in reality batteries have losses just as, *e.g.*, refrigeration systems increase the heat load, and thereby the power loss, as the temperatures are lowered to store extra cooling energy. A loss term could be modeled as $-\eta_{\text{loss}}Q(t)$ which is added to the equation above, but, as our time-scale for storing energy is in the range of seconds to minutes, we neglect this.

We can now find the power supplied from the WPP to the grid

$$P_{\text{grid}} = P_g - c.$$




Fig. 1: Coefficient of power C_P . The peak power coefficient is 0.482.

C. Control

Manipulated variables: Our optimizer manipulates the set-points to the generator torque, $T_{g,\text{ref}}$, and the pitch angle, β_{ref} , for each individual turbine in the WPP. Normally, the pitch is controlled by an inner loop exercising a gain-scheduled PI controller. This controller samples up to 100 times faster than our MPC and we set $\beta_{\text{ref}} = \beta$ as long as the slew rate limit on β_{ref} is observed. Additionally, we manipulate the charge rate, c , to/from the central energy storage.

Measured variables: The controller bases its decisions on measurements of the rotational speed and generator torque in each turbine, on the known current wind speed, and on the filtered, predicted future values of the latter covering the entire prediction horizon, N_p .

D. Constraints

The extracted power, P_g , must be equal to or less than the available power in the wind, P_w , which is a function of the wind speed, v . The turbine is build for a rated power and when the available power in the wind exceeds this level, the blades gradually pitch out of the wind to keep the extracted power at the rated level and reduce loads on the turbine. Likewise, the extracted power can only follow the available power curve down to a certain level, P_{min} , due to the mechanical design. Thus,

$$P_{\text{min}} \leq P_g \leq P_{\text{rated}}, \quad (4)$$

$$P_g \leq P_w(v), \quad (5)$$

For security reasons, the turbine is turned off for wind speeds above 25m/s. Therefore, $P_w = 0$, and the constraint (4) is not relevant for such high wind speeds.

In addition, four physical constraints are given by the system

$$0^\circ \leq \beta \leq 90^\circ, \quad (6)$$

$$-8^\circ/s \leq \frac{\beta(t + T_s) - \beta(t)}{T_s} \leq 8^\circ/s, \quad (7)$$

$$c \leq P_g, \quad (8)$$

and

$$0 \leq Q \leq \mu P_{\text{rated}}, \quad (9)$$

where we introduce the variable μ which is the maximum needed storage capacity in per unit (pu), *i.e.*, normalized by rated power.

The rotational speed is usually limited by a maximum rated speed, mainly due to too high loads on the turbine at higher speeds. However, since we want to put the turbine inertia in play, we allow for higher speeds and introduce the parameter ω_{os} which is the fraction of the rated maximum speed that we accept as over-speed.

$$\omega_{g,\text{rated},\text{min}} \leq \omega_g \leq (1 + \omega_{\text{os}})\omega_{g,\text{rated},\text{max}}. \quad (10)$$

The power supplied to the grid P_{grid} must fulfill the power gradient

$$-\Delta pu \leq \frac{P_{\text{grid}}(t + T_s) - P_{\text{grid}}(t)}{P_{\text{rated}}T_s} \leq \Delta pu, \quad (11)$$

where $\Delta pu \in [0, 1]$ is the grid code for maximum power gradient in per unit with respect to rated power.

In this study, we do not include the wake effects that couple the individual turbines through the downwind wind flow which is affected by the amount of power extracted by upwind turbines. This type of constraint is a focus of our future work.

We define the set Ω as all $(T_{g,\text{ref}}, \beta_{\text{ref}}, c)$ that satisfy the system dynamics (1)–(3) and the constraints (4)–(11).

E. Cost

We assume in this study that the objective of the WPP is to maximize the average power supplied to the grid. Alternative operating modes such as delta production (keeping a reserve by producing less than possible) or frequency response (reacting on frequency deviations on the grid to support stabilization at nominal grid frequency) are thus not considered. The supplied energy, E , over the period $[T_0, T_{\text{final}}]$ is

$$E = \int_{T_0}^{T_{\text{final}}} P_{\text{grid}} dt.$$

Furthermore, we have a cost on the available storage capacity. For a period $[T_0, T_{\text{final}}]$ this is

$$S = \max[\mu]_{T_0}^{T_{\text{final}}} c_{\text{storage}}.$$

We can consider the storage price, c_{storage} , as a tuning parameter or as directly reflecting, *e.g.*, purchase price of batteries divided by their lifetime, a service agreement with a flexible consumer, etc. Thus, S is a cost in the design phase only (or for simulations as we will show here).

F. Nominal Controller

We compare the performance of our proposed method to the solution from the nominal wind turbine control strategy, also defined in the NREL 5MW model. For natural reasons this system is only capable of obeying the power gradient constraint in three cases: 1) When the rate-of-change of the available power in the wind is less than the power gradient constraint. 2) When the available power in the wind only changes from one point to another, where both are above rated power. 3) When sufficiently large amounts of storage is added and its charge/discharge is controlled by some kind of predictive control with knowledge of the future wind speed.

III. MPC CONTROLLER

The WPP is influenced by disturbances from the wind speed which we can predict (with some uncertainty) over a time horizon into the future. The controller must obey certain constraints, while maximizing the power supply and limiting additional costs for storage. Economic MPC can address all these concerns. Whereas the cost function in MPC traditionally penalizes a deviation from a set-point, the proposed economic MPC directly reflects the actual costs of operating the plant. Like in traditional MPC, we implement the controller in a receding horizon manner, where an optimization problem over N_p time steps (the control and prediction horizon) is solved at each step. The

result is an optimal input sequence for the entire horizon, out of which only the first step is implemented. The optimization problem is thus formulated as

$$\begin{aligned} & \text{maximize} && E - S, \\ & \text{subject to} && (\mathbf{T}_{g,\text{ref}}, \boldsymbol{\beta}_{\text{ref}}, \mathbf{c}) \in \Omega, \end{aligned} \quad (12)$$

where the variables are $T_{g,\text{ref}}$, β_{ref} and c (all functions of time). Instead of (12) we solve a discretized version with N_p steps over the time interval $[T_0, T_{\text{final}}]$,

$$\{\mathbf{T}_{g,\text{ref}}, \boldsymbol{\beta}_{\text{ref}}, \mathbf{c}\} = \{T_{g,\text{ref}}^k, \beta_{\text{ref}}^k, c^k\}_{k=0}^{N_p-1}. \quad (13)$$

The MPC feedback law is the first move in (13). The controller uses the initial state as well as predictions of the wind speed for the time interval. The predictions could come from any good sources available, see *e.g.*, [29] where 10-minute ahead predictions are implemented.

A. Sequential convex programming method

As we saw in §II, neither the feasible set Ω nor the cost function term P are fully convex. Instead of using a generic nonlinear optimization tool, we choose to solve the optimization problem iteratively using convex programming, replacing the nonconvex terms with convex approximations. In each iteration, i , we perform a first-order Taylor expansion of the nonconvex parts around the operating point found in iteration $i - 1$, estimating the derivatives that involve table look-ups by perturbing the parameters. As the wind speed v is predicted we can use v^3 as input to our model instead. We establish the following linear approximations

$$\begin{aligned} \hat{T}_r^i &= T_r^{i-1} + \left[\frac{\Delta T_r}{\Delta \omega_r}, \frac{\Delta T_r}{\Delta C_P} \right]^{i-1} \begin{bmatrix} \omega_r^i - \omega_r^{i-1} \\ C_P^i - C_P^{i-1} \end{bmatrix}, \\ \hat{P}_g^i &= P_g^{i-1} + \left[\frac{\Delta P_g}{\Delta \omega_g}, \frac{\Delta P_g}{\Delta T_g} \right]^{i-1} \begin{bmatrix} \omega_g^i - \omega_g^{i-1} \\ T_g^i - T_g^{i-1} \end{bmatrix}. \end{aligned}$$

Thus, in each iteration we solve a convex optimization problem, which can be done very reliably and extremely quickly [30]. While our proposed method gives no theoretical guarantee on the performance, we must remember that the optimization problem is nothing but a heuristic for computing a good control and that the quality of closed-loop control with MPC is generally good without solving each problem accurately.

B. Regularization

We use two different types of regularization in the optimization problem. To avoid oscillations from iteration to iteration, we add proximal regularization of the form

$$\varphi_{\text{prox}} = \rho_{\text{prox}} \sum_{k=0}^{N-1} \|X^k - X^{k,\text{prev}}\|_2^2, \quad (14)$$

for each of the control variables $X = \{T_{g,\text{ref}}, \beta_{\text{ref}}, c\}$. The superscript ‘prev’ indicates that it is the solution from the previous iteration and ρ_{prox} is a constant weight chosen to damp large steps in each iteration. In addition, we add a

quadratic penalty on the rate-of-change (roc) of the manipulable variables,

$$\varphi_{\text{roc}} = \rho_{\text{roc}} \sum_{k=1}^{N-1} \|X^k - X^{k-1}\|_2^2. \quad (15)$$

This regularization term serves two purposes: It improves the convergence of the sequential programming method, and also discourages rapid changes, which helps reduce oscillations and fatigue loads.

C. Algorithm

Algorithm 1 outlines the method. The term *nominal* refers to the solution obtained from the nominal controller. We use this as a baseline for initializing the algorithm. In MPC, the solution from the previous time step is usually well suited for warm-starting the algorithm.

Algorithm 1 Iterative optimization.

Initialize

$$\{\mathbf{T}_g^0, \mathbf{T}_r^0, \boldsymbol{\omega}_g^0, \boldsymbol{\omega}_r^0, \mathbf{C}_P^0, \mathbf{T}_{g,\text{ref}}^0, \boldsymbol{\beta}_{\text{ref}}^0\} = \{\text{nominal}(v_k)\}_{k=1}^{N_p},$$

$i = 1$.

Compute

$$\hat{\mathbf{T}}_g^i, \hat{\mathbf{P}}_g^i \text{ and } \hat{\mathbf{C}}_P^i, \text{ from } \{\mathbf{T}_g, \mathbf{T}_r, \boldsymbol{\omega}_g, \boldsymbol{\omega}_r, \mathbf{C}_P, \mathbf{T}_{g,\text{ref}}, \boldsymbol{\beta}_{\text{ref}}\}^{i-1}$$

and v .

Solve

$$\begin{aligned} & \text{maximize} && E^i - S^i + \varphi_{\text{prox}} + \varphi_{\text{roc}}, \\ & \text{subject to} && (\mathbf{T}_{g,\text{ref}}^i, \boldsymbol{\beta}_{\text{ref}}^i, \mathbf{c}) \in \hat{\Omega}, \end{aligned}$$

Update

$$\{\mathbf{T}_g^i, \mathbf{T}_r^i, \boldsymbol{\omega}_g^i, \boldsymbol{\omega}_r^i, \mathbf{C}_P^i, \mathbf{T}_{g,\text{ref}}^i, \boldsymbol{\beta}_{\text{ref}}^i\}, \text{ and } i = i + 1$$

Repeat until convergence.

IV. RESULTS

In this paper, we apply the proposed method to a conceptual study limited to only one wind turbine. We implement and solve our controller for different scenarios using CVX [31], [32]. In this section, we report on results with a power gradient constraint as low as 3% of the rated power per minute ($\Delta p u = 5 \cdot 10^{-4} \text{pu/s}$) and with an allowed overspeed of 50% above rated speed for short time intervals. We sample with $T_s = 1 \text{s}$ intervals and use a horizon of 5 minutes ($N_p = 300$) in this case. Obviously, a wide range of solutions can be obtained depending on the specific ramp rate of the wind speed, the wind speeds before and after the change occurs, the allowed amount of overspeed and the definition of storage price versus power sales price. In this paper, we give proof-of-concept of the method, using a few selected trajectories, and for the next version of this work, we will derive a more generalized measure of the relation between wind ramp rates, overspeed ratio, power constraint, and storage capacity.

Fig. 2 shows examples of how our proposed method performs in different cases, while satisfying the power gradient constraint. For all four cases shown in the figures, we can calculate the total power delivered to the grid from $t = 0 \dots 800 \text{s}$. For the scenario in figures 2(a)–2(b) (wind speed changes from 10m/s to 8m/s), the available power in the wind is below the rated power for the entire interval. Thus,




Fig. 2: Test of power gradient satisfaction. We use pu as the unit for all quantities, except the state-of-charge (Q) which is normalized first with respect to maximum storage capacity. In all scenarios we let the wind speed drop from v_1 to v_2 linearly from $t = 400$ s to $t = 430$ s, and we show cases with high and low storage cost. $P_{g,nom}$ is the power output from the nominal controller.

no extra power exists for accelerating the rotor beyond rated speed. If central energy storage is relatively cheap (Fig. 2(a)) this is used entirely as a buffer for achieving the commanded power gradient while the turbine behaves exactly as with the nominal controller. In this case, the total amount of energy delivered to the grid is equal to the nominal case too. As the price of energy storage increases the controller trades off the power production that is lost during the phase where the rotor is accelerated, in order to use that kinetic energy during the power ramp to reduce the peak of needed storage capacity. In Fig. 2(b) the storage capacity is reduced by 26.5% at the cost of 1.3% of the energy delivered to the grid, compared to Fig. 2(a). For the scenario in figures 2(c)–2(d) (wind speed changes from 12m/s to 10m/s), the available power in the wind goes from above rated to below rated power. In this case, the rotor can be accelerated to reach the maximum

allowed speed just when the available power in the wind begins to drop. This kinetic energy is used during the power ramp no matter how cheap storage is, as it only adds to the total delivered energy. In Fig. 2(c) the amount of energy delivered to the grid is 1.6% higher than with the nominal controller. When storage cost is increased, the utilization of stored inertia is shifted towards the time when the storage needs peak, to reduce the required additional capacity, and the extra production gained otherwise is now traded off with storage cost. In Fig. 2(d) the energy delivered is just 0.03% less than with the nominal controller while the storage need is reduced by almost 7%.

A. Convergence and Computation

When initialized with the trajectory from the nominal controller, the proposed method generally converges in 5–10 iterations. In MPC, however, the open-loop trajectory

from the previous run of the optimizer, shifted one time-step, is an excellent guess on the next outcome and is well-suited for warm-starting the algorithm. Using this warm start initialization, the method generally just need a couple of iterations to converge.

V. CONCLUSION

In this paper, we present an approach to power gradient reduction for fulfilling future, tighter grid codes and for improving the quality of power delivered to the grid from wind power plants. We utilize turbine inertia as a resource of distributed energy storage, limited by the rotational speed, in addition to a central storage unit which is associated with an extra cost. Our method is based on convex optimization, solved iteratively to handle the nonconvexity of the aerodynamics. Simulation on realistic models reveal a significant ability to reject the disturbances from fast changes in wind speed, ensuring certain power gradients, while keeping the amount of produced power close to nominal. We can easily trade off lost production versus price of extra energy storage.

VI. ACKNOWLEDGMENTS

We thank Martin Ansbjerg Kjær from Vestas Technology R&D for helpful discussions, ideas and suggestions.

REFERENCES

- [1] Danish Ministry of Climate Energy and Building, "Energy policy report 2012," Tech. Rep., May 2012. [Online]. Available: <http://www.ens.dk/en-US/policy/danish-climate-and-energy-policy/Sider/danish-climate-and-energy-policy.aspx>
- [2] J. Morren, J. Pierik, and S. W. H. de Haan, "Inertial response of variable speed wind turbines," *Electric Power Systems Research*, vol. 76, no. 11, pp. 980–987, 2006.
- [3] J. F. Conroy and R. Watson, "Frequency Response Capability of Full Converter Wind Turbine Generators in Comparison to Conventional Generation," *IEEE Transactions on Power Systems*, vol. 23, no. 2, pp. 649–656, may 2008.
- [4] F. Iov, A. Hansen, P. Sørensen, and N. Cutululis, "A survey of interconnection requirements for wind power," in *Proc. of the Nordic wind power conference (NWPC)*. Risø National Laboratory, 2007.
- [5] B. Singh and S. Singh, "Wind Power Interconnection into the Power System: A Review of Grid Code Requirements," *The Electricity Journal*, vol. 22, no. 5, pp. 54–63, 2009.
- [6] Eltra/Elkraft/Energinet.dk, "Regulation TF 3.2.5, Wind turbines connected to grids with voltages above 100 kV — Technical regulation for the properties and the regulation of wind turbines," <https://selvbetjening.preprod.energinet.dk/NR/rdonlyres/E4E7A0BA-884F-4E63-A2F0-98EB5BD8D4B4/0/WindTurbinesConnectedtoGridswithVoltageabove100kV.pdf>.
- [7] M. Korpås and A. T. Holen, "Operation planning of hydrogen storage connected to wind power operating in a power market," *IEEE Transactions on Energy Conversion*, vol. 21, no. 3, pp. 742–749, 2006.
- [8] M. Black and G. Strbac, "Value of Bulk Energy Storage for Managing Wind Power Fluctuations," *IEEE Transactions on Energy Conversion*, vol. 22, no. 1, pp. 197–205, march 2007.
- [9] D. I. Stroe, A. I. Stan, R. Diosi, R. Teodorescu, and S. J. Andreasen, "Short term energy storage for grid support in wind power applications," in *Proc. of the 13th International Conference on Optimization of Electrical and Electronic Equipment*, 2012, pp. 1012–1021.
- [10] C. Budischak, D. A. Sewell, H. Thomson, L. Mach, D. E. Veron, and W. Kempton, "Cost-Minimized Combinations of Wind Power, Solar Power and Electrochemical Storage, Powering the Grid up to 99.9% of the Time," *Journal of Power Sources*, vol. 225, pp. 60–74, 2013.
- [11] M. Swierczynski, R. Teodorescu, and P. Rodriguez, "Lifetime investigations of a lithium iron phosphate (LFP) battery system connected to a wind turbine for forecast improvement and output power gradient reduction," in *Proc. of the 15th Battcon Stationary Battery Conference and Trade Show*, 2011, pp. 20.1–20.8.
- [12] T. G. Hovgaard, L. F. S. Larsen, K. Edlund, and J. B. Jørgensen, "Model predictive control technologies for efficient and flexible power consumption in refrigeration systems," *Energy*, vol. 44, no. 1, pp. 105 – 116, 2012.
- [13] T. G. Hovgaard, L. F. S. Larsen, J. B. Jørgensen, and S. Boyd, "Nonconvex model predictive control for commercial refrigeration," *International Journal of Control*, p. In Press, 2012.
- [14] T. Knuppel, J. N. Nielsen, K. H. Jensen, A. Dixon, and J. Østergaard, "Power oscillation damping controller for wind power plant utilizing wind turbine inertia as energy storage," in *Proc. of the IEEE Power and Energy Society General Meeting*, 2011, pp. 1–8.
- [15] G. C. Tarnowski, "Coordinated Frequency Control of Wind Turbines in Power Systems with High Wind Power Penetration," Ph.D. dissertation, Technical University of Denmark, 2012.
- [16] K. Hammerum, P. Brath, and N. K. Poulsen, "A fatigue approach to wind turbine control," *Journal of Physics: Conference Series*, vol. 75, 2007.
- [17] D. Q. Dang, S. Wu, Y. Wang, and W. Cai, "Model Predictive Control for maximum power capture of variable speed wind turbines," in *International Power Electronics Conference*, 2010, pp. 274–279.
- [18] L. C. Henriksen, N. K. Poulsen, and M. H. Hansen, "Nonlinear Model Predictive Control of a Simplified Wind Turbine," in *Proc. of the 18th IFAC World Congress*, 2011, pp. 551–556.
- [19] M. Soleimanzadeh, R. Wisniewski, and S. Kanev, "An optimization framework for load and power distribution in wind farms," *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 107–108, pp. 256–262, 2012.
- [20] V. Spudić, M. Jelavić, and M. Baotić, "Wind turbine power references in coordinated control of wind farms," *Automatika—Journal for Control, Measurement, Electronics, Computing and Communications*, vol. 52, no. 2, pp. 82–94, 2011.
- [21] T. G. Hovgaard, L. F. S. Larsen, J. B. Jørgensen, and S. Boyd, "Sequential Convex Programming for Power Set-point Optimization in a Wind Farm using Black-box Models, Simple Turbine Interactions, and Integer Variables," in *Proc. of the 10th European Workshop on Advanced Control and Diagnosis (ACD)*, 2012, pp. 1–8.
- [22] S. Boyd, "EE364b, Lecture Slides and Notes: Sequential convex programming," http://www.stanford.edu/class/ee364b/lectures/seq_slides.pdf.
- [23] P. T. Boggs and J. W. Tolle, "Sequential Quadratic Programming," *Acta Numerica*, vol. 4, pp. 1–51, 1995.
- [24] Q. T. Dinh, C. Savorgnan, and M. Diehl, "Real-time sequential convex programming for nonlinear model predictive control and application to a hydro-power plant," in *Proc. of the 50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*, 2011, pp. 5905–5910.
- [25] M. Kraning, Y. Wang, E. Akuiyibo, and S. Boyd, "Operation and Configuration of a Storage Portfolio via Convex Optimization," in *Proc. of the 18th IFAC World Congress*, 2011, pp. 10487–10492.
- [26] T. G. Hovgaard, S. Boyd, and J. B. Jørgensen, "Model predictive control for wind power gradients," http://www.stanford.edu/~boyd/papers/wind_gradients_cvx.html, 2013.
- [27] J. M. Jonkman, S. Butterfield, W. Musial, and G. Scott, *Definition of a 5-MW reference wind turbine for offshore system development*. National Renewable Energy Laboratory, Feb 2009.
- [28] J. D. Grunnet, M. Soltani, T. Knudsen, M. N. Kragelund, and T. Bak, "Aeolus Toolbox for Dynamics Wind Farm Model, Simulation and Control," in *Proc. of the European Wind Energy Conference and Exhibition, EWEC*, 2010.
- [29] P. Pinson, "Very-short-term probabilistic forecasting of wind power with generalized logit-normal distributions," *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, vol. 61, no. 4, pp. 555–576, 2012.
- [30] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [31] CVX Research, Inc., "CVX: Matlab Software for Disciplined Convex Programming, version 2.0 beta," <http://cvxr.com/cvx>, Sept. 2012.
- [32] M. Grant and S. Boyd, "Graph implementations for nonsmooth convex programs," in *Recent Advances in Learning and Control*, ser. Lecture Notes in Control and Information Sciences, V. Blondel and S. Boyd and H. Kimura, Ed. Springer-Verlag Limited, 2008, pp. 95–110.