Value of Storage for Wind Power Producers in Forward Power Markets

Milind Rao, Mainak Chowdhury, Yue Zhao, Tara Javidi, Andrea Goldsmith

Abstract—**Wind power producers (WPPs) that sell power in forward power markets would like to minimize their operating costs which increase with generation uncertainty. In this work, the value of energy storage for reducing such costs is studied. In particular, profit maximization is considered for a WPP who participates in a two-settlement (forward and real time) market and utilizes energy storage by charging/discharging it strategically. An infinite horizon discounted cost minimization problem for the optimal use of energy storage is formulated as a dynamic programming (DP) problem that includes the past unfulfilled forward contracts in the state space. The optimal storage operation policy is shown to have a structure with two thresholds: after delivering its contracted power, if a WPP's energy falls below a lower threshold, it buys energy and charges its storage up to this threshold; if its energy exceeds a higher threshold, it sells the excess energy and maintains its storage level at this threshold. Several heuristics for solving the DP are derived based on approximating the problem model: a) a discrete policy based on discretizing the state and action space, and b) affine and look ahead policies derived by solving a Linear Quadratic (LQ) controller whose parameters are fit from the DP. The heuristics are tested both with simulated and real world wind and price data. It is observed that while the discrete optimal policy performs better on simulated data than either the look ahead or the affine policies (except with a very high battery capacity), the look ahead policy performs much better with real world data. This suggests that the performance of look ahead approximate optimal policy is more robust to the modeling errors and mismatch between analytic models and real data traces. The appropriate heuristic to use thus depends on modeling fidelity, available computational resources and variability of wind and price forecasts.**

I. Introduction

Driven by greenhouse gas reduction goals, many countries and areas have set ambitious target levels for incorporating renewable energy into the electric grid. Wind energies, in particular, have experienced significant growth in their installation, as their capital costs are rapidly decreasing [1]. While wind energies are clean and have low variable cost in their operation, they are also intrinsically uncertain and variable. Hence, conventional generation assets must be used to compensate for the uncertainty and variability of the wind energies in the electric grid.

A primary approach for integrating wind energy into the electric grid is to let wind power producers (WPPs) participate in conventional power markets to sell their power generation [2]. Typically, a vast majority of the power supply is traded in forward markets (e.g., day-ahead wholesale market), with the remaining supply and demand handled by the real time markets. For a WPP, however, selling power in forward markets critically depends on knowledge of its future wind power generation, which is inherently uncertain due to difficulties of wind forecast [3]. Any shortfall between a WPP's commitment in a forward market and its actual generation at the delivery time has to be made up by buying in the real time market, which is typically costly. As wind power generation has very low variable cost, the cost due to its generation uncertainty is thus a WPP's major operating cost. Consequently, there are strong incentives for WPPs to reduce their uncertainty.

A variety of approaches for WPPs to reduce generation uncertainty have been studied. These include exploiting statistical characteristics of wind generation, such as improving wind forecast [3] and aggregating diverse wind sources [4]. In this context, energy storage promises to emerge as an increasingly viable technology as its cost continues to decrease [5]. For instance, co-located energy storage enables a WPP to shift energy across time, and charge/discharge the stored energy to compensate any generation shortfall required to meet forward contracts. Nonetheless, as energy storage is currently still expensive, assessing its value in reducing operating cost and increasing profit is of primary interest to WPPs.

We consider a WPP participating in a dynamically evolving conventional two-settlement (day-ahead and real time) market, with the help of co-located energy storage (i.e., a battery). The problem of a WPP participating in a two-settlement market without storage has been studied in [6], [7] and [8], and optimal forward contracts based on the statistics of wind generation have been developed. The problem of using energy storage (in the absence of renewable energy) to arbitrage in a real time power market has been studied in [9] and an approximate dynamic programming solution was proposed. Co-located with wind power participating in a *real time* market only, the value of storage has been studied in [10] and [11], for which dynamic programming (DP) solutions and online algorithms have been developed, respectively. Day-ahead market is considered in assessing the role of storage

This research was supported in part by 3Com Corporation Stanford Graduate Fellowship, and in part by the NSF under CPS Synergy grant 1330081.

M. Rao, M. Chowdhury, and A. Goldsmith are with the Dept. of Electrical Engineering, Stanford University, Stanford, CA 94305, USA (e-mail: mainakch@stanford.edu, milind@stanford.edu, andrea@ee.stanford.edu).

Y. Zhao is with the Dept. of Electrical and Computer Engineering, Stony Brook University, Stony Brook, NY 11794 USA (e-mail: yue.zhao.2@stonybrook.edu).

T. Javidi is with the Dept. of Electrical and Computer Engineering, University of California, San Diego, La Jolla, CA 92093, USA (e-mail: tjavidi@ucsd.edu).

The first three authors contributed equally to this work.

co-located with wind power in [12], where the optimal day-ahead contracts and storage operation are solved for a single day period. Finally, a dynamically evolving two-settlement market has been considered in [13] and [14] in which optimal operation of co-located storage and wind power are studied. There, the model is limited to one in which the forward market trades power *one time slot ahead*. As will be clarified next, this corresponds to the case of $D = 1$ in this paper. We note that, when hourly day-ahead and hourly real-time markets are considered, D can be as large as 24. This poses great challenges in computing the optimal policy as the computational complexity grows exponentially with D.

In this work, we study a WPP that has co-located storage in a general dynamically evolving two-settlement market, where forward contracts are made D (≥ 1) time slots ahead of delivery. We consider independent periodic wind and price processes, and find optimal control policies for forward contracts as well as battery charging/discharging levels. We formulate the problem as a DP by including the committed forward contracts for the past D time slots in the state space. Based on the convexity of the value function of this DP, we prove that the optimal energy storage control policy exhibits a structure associated with two battery thresholds $\underline{b}$ and $\bar{b}$: a) if a WPP has remaining energy (after fulfilling its current delivery commitment) below $\underline{b}$, it buys energy from the real time market to charge the battery up to $\underline{b}$, b) if it has remaining energy above $\bar{b}$, it sells all its excess energy in the real time market to maintain the battery level as $\bar{b}$, and c) otherwise, it stores the remaining energy in the battery without buying or selling in the real time market. Even with such characterization of the optimal policy, it is not clear how the WPP should optimally grant contracts in the forward market. Furthermore, the DP characterization suffers from the curse of dimensionality as well as the infiniteness of state and action space.

To address the computational challenge of solving the proposed DP formulation, two approximations of the problem are studied and optimal policies for these approximations are found. In particular, we determine an optimal discrete policy based on discretizing the state and action space. We also develop affine and look ahead policies by solving a Linear Quadratic (LQ) stochastic control problem with quadratic stage cost functions approximating costs in the original DP formulation and uncertainty in price process not considered. We test the approximate optimal policies first with simulated wind and price processes whose statistics are learned from the real world data, and then with real world wind and price data directly. As one would expect, discrepancies between real world wind and price processes and the ones derived from a parameterized model can be significant. As a result, we observe a notable performance penalty gap when the proposed algorithm is applied to the real world data instead of simulated data. An analytical characterization of this gap, in terms of the statistical properties of the wind and price processes, is left for future work.

The remainder of the paper is organized as follows. Section II establishes the system model and formulates the DP problem. Section III proves the threshold structure of the energy storage operation policy. Section IV develops computation methods for approximate optimal policies. Simulations based on simulated and real world data are presented in Section V. Conclusions are drawn in Section VI.

II. Problem Formulation

A. System Model

We consider a WPP with battery energy storage of capacity B. We study the infinite horizon problem of maximizing the expected discounted profits by selling wind power in a two-settlement market accompanied by storage operation. The system model is described as follows. At time $t \geq 0,\ t \in \mathcal{Z}$,

- The WPP receives wind energy w_t, which arises independently from a distribution with a time dependent probability density function (pdf) $p_t(w_t)$.
- The battery level is b_t.
- The WPP participates in the forward market by committing to supplying s_t units of energy, to be delivered D time slots later. It earns $p_t^f s_t$, where p_t^f is the price of unit power in this forward market.
- The previously committed power delivery s_{t-D} is fulfilled. The energy surplus (which is negative if there is a deficit) after using the available energy (from wind and battery) to fulfil s_{t-D} is denoted by $e_t = b_t + w_t - s_{t-D}$.
- The WPP decides the target battery level at the next time slot, from which the charging/discharging amount is implied. We denote this action by b_{t+1}.
- If $e_t > b_{t+1}$, there is remaining energy after storing the next time slot's battery amount b_{t+1}, and this remaining energy $e_t - b_{t+1}$ is sold in the real time market. If $e_t < b_{t+1}$, the WPP must buy $b_{t+1} - e_t$ in the real time market to charge the battery to b_{t+1}. The selling price in the real time market is denoted by p_t^s, and the buying price by p_t^b. We assume $p_t^b > p_t^s$, and $\mathbb{E}\left[p_t^b\right] > p_t^f > \mathbb{E}\left[p_t^s\right]$ to prevent arbitrage opportunities where the WPP can purchase power from the real time markets at low rates to satisfy past contracts or is motivated to sell power in the real time market instead of contracting to sell this power in the forward market.
- The price processes (p_t^f, p_t^b, p_t^s) arise independently from a time dependent joint distribution.
- The wind and (joint) price processes are independent and arise from distributions that are periodic with a period D.

We define a policy decision at time t to be $u_t \triangleq (s_t, b_{t+1}) \in \mathcal{A}$, where $\mathcal{A}$ is the action space. Specifically, as wind generation is always non-negative,

$$\mathcal{A} = \{(s_t, b_{t+1}) : 0 \leq b_{t+1} \leq B, s_t \geq 0\}. \quad (1)$$

We observe that the state of the system at time t consists of the following tuple $x_t \triangleq \left(t \bmod D, p_t^f, p_t^s, p_t^b, w_t, s_{t-D}^{t-1}, b_t\right)$. We denote the state space by $\mathcal{X}$. In particular, the past unfulfilled contracts $s_{t-D}^{t-1} \triangleq \{s_{t-D}, \ldots, s_{t-1}\}$ need to be included in the state space in order to formulate a dynamic program as follows.

B. Dynamic program

We consider expected discounted profit maximization as the objective function, and formulate the following optimal control problem:

$$V(x) \triangleq \min_{u_t(\cdot)} \quad \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t g(x_t, u_t(x_t)) \mid x_0 = x\right] \quad \text{s.t.} \quad x_{t+1} = f(x_t, u_t(x_t), z_t) \tag{2}$$

where

- $\beta \in (0,1)$ is some discount factor,
- Stage cost $g(\cdot,\cdot)$ is defined as
$$g(x_t,u_t) = g_f(p_t^f, s_t) + g_r(p_t^b, p_t^s, b_{t+1} - e_t) + \mathbf{I}_{\mathcal{B}_t}(b_{t+1}), \tag{3}$$
where $g_f(\cdot,\cdot)$ represents the forward market cost function, and $g_r(\cdot,\cdot,\cdot)$ represents the real time market cost function. $g_f(\cdot,\cdot)$ and $g_r(\cdot,\cdot,\cdot)$ are both convex in their individual arguments. The third indicator function ensures that the battery energy level satisfies the constraints (1), and is defined for an appropriate set $\mathcal{B}$ as
$$\mathbf{I}_{\mathcal{B}}(x) = \begin{cases} 0 & \text{if } x \in \mathcal{B} \\ \infty & \text{otherwise.} \end{cases}$$
The battery operation is modeled to be efficient such that associated loss is neglected and the storage ramping constraint is not considered. In other words, the battery can be charged or discharged to any amount in the subsequent stage with no losses [1]. Cost functions are then given by $((x)_+ \triangleq \max(0,x))$,
$$\mathcal{B}_t = [0,B], g_f(p_t^f, s_t) = -p_t^f \cdot s_t, \tag{4}$$
$$g_r(p_t^b, p_t^s, b_{t+1} - e_t) = p_t^b (b_{t+1} - e_t)_+ - p_t^s (e_t - b_{t+1})_+, \tag{5}$$
- $z_t = [w_{t+1}, p_{t+1}^f, p_{t+1}^s, p_{t+1}^b]$ captures the random wind and price processes at time $t+1$.
- $f(x,u,z)$ is the state transition function, and it is immediate to see that x_{t+1} can be simply determined by copying appropriate entries from x_t, u_t, z_t. Thus, $f(x,u,z)$ is a linear function in x, u and z.

Note that the optimization variables here are the (infinite-dimensional) policies $\{u_t(x_t)\}_{t=0}^{\infty}$. Each u_t maps the state x_t to (s_t, b_{t+1}). Thus, we have $u_t : \mathbf{R}^{D+6} \to \mathbf{R}^2$. The information pattern is such that at time t, the following quantities are known exactly:

- Wind realization $\{w_\tau\}_{\tau=-\infty}^{t}$.
- Price processes $\{p_\tau^f, p_\tau^b, p_\tau^s\}_{\tau=-\infty}^{t}$.

The above formulation provides a dynamic programming problem. In particular, the Bellman operator is

$$\mathcal{T}(V)(x_t)(\cdot) \triangleq \inf_{u\in\mathcal{A}} \left(g(x_t,u) + \beta\mathbb{E}\left[V(f(x_t,u,z))\right] \right).$$

The expectation is over randomness in z. The time at each step is considered to be known. The optimality equation of the dynamic program is given by the fixed point equation

$$V = \mathcal{T}(V). \tag{6}$$

[1]The model can accommodate storage ramping constraints and the only change that has to be made is $\mathcal{B}_t = [0,B] \cap [b_t - B_r, b_t + B_r]$, where B_r is the ramping constraint at each stage.

We observe that the stage cost $g(\cdot,u)$ is uniformly bounded above by a constant as the maximum wind, price and contract levels are finite. $\mathcal{T}$ defines a contraction over functions, i.e. $\|\mathcal{T}(V_1) - \mathcal{T}(V_2)\|_\infty \leq \beta \|V_1 - V_2\|_\infty$, where $\|\cdot\|_\infty$ represents the max norm. We thus have the result that a value iteration converges to a fixed point [15], i.e. $\lim_{n\to\infty} \mathcal{T}^n(V_0) = V^*$.

III. Structural Properties of Value Function and Optimal Policy

We now investigate structural properties of the value function and the optimal policy of the dynamic program (6). Throughout this section, we use b_+ to denote the next time slot's battery level as a result of the control action, b_- the current battery level, and s_{-D} the previous commitment that needs to be fulfilled in the current time slot.

We begin with the following fact [16].

Theorem 1: $V(x)$ is a convex function in the state variables w_t, b_t, s_{t-D}^{t-1}.

Next, we derive structural properties of the optimal policy $u^*(x) = [b_+^*(x), s^*(x)]^T$. By definition,

$$u^*(x) = \operatorname{argmin}_u g(x,u) + \mathbb{E}\left[\beta V(f(x,u,z))\right]. \tag{7}$$

Furthermore, the optimal contract *given an action* b_+, denoted by $\tilde{s}^*$, is

$$\tilde{s}^*(x,b_+) = \operatorname{argmin}_s g(x,b_+,s) + \mathbb{E}\left[\beta V\left(f(x,b_+,s,z)\right)\right].$$

It is clear that $\tilde{s}^*(x,b_+)$ depends only on b_+ and a *reduced set* of state variables $\tilde{x} \triangleq [p^f, s_{-D+1}^{-1}, t \bmod D]^T$, and can be equivalently written as $\tilde{s}^*(\tilde{x},b_+) = \operatorname{argmin}_s \mathbb{E}\left[g_f(p^f,s) + \beta V\left(f(\tilde{x},b_+,s,z)\right)\right]$. In addition, the optimal actions b_+^* and s^* satisfy the following equations:

$$b_+^*(x) = \operatorname{argmin}_{b_+} \mathbb{E}\big[g\left(x,b_+,\tilde{s}^*(x,b_+)\right) + \beta V\left(f\left(x,b_+,\tilde{s}^*(x,b_+),z\right)\right)\big], \tag{8}$$

$$s^*(x) = \tilde{s}^*(x, b_+^*(x)). \tag{9}$$

Furthermore, we will make use of the following notation:

$$\tilde{V}(\tilde{x},b_+) \triangleq \min_s g_f(p^f,s) + \mathbb{E}\left[\beta V\left(f(\tilde{x},b_+,s,z)\right)\right].$$

We have the following lemma.

Lemma 1: $\tilde{V}(\tilde{x},b_+)$ is a non-increasing convex function of b_+.

Proof: The non-increasing property follows from the observation that storing more energy cannot decrease expected profit (and the value function captures the negative of profit). Next, we observe that $g_f(p^f,s) + \mathbb{E}\left[\beta V\left(f(\tilde{x},b_+,s,z)\right)\right]$ is a jointly convex function in b_+, s_{-D}^0 because a) $V(\cdot)$ is convex in b_t, s_{t-D}^{t-1} (cf. Theorem 1), and b) $f(\cdot)$ is an affine mapping and the entries corresponding to b_{t+1}, s_{t-D+1}^t only depend on b_+, s, s_{t-D}^{t-1}. Finally, expectation over z and minimization over s both preserve convexity [17].

Employing (4) and (5), $b_+^*(x)$ can be rewritten as follows,

$$b_+^*(x) = \operatorname{argmin}_{b_+} p^b(b_+ - e)_+ - p^s(e - b_+)_+ + \tilde{V}(\tilde{x},b_+), \tag{10}$$

where $e = b_- + w - s_{-D}$ is the remaining energy (can be negative) after using the available energy to fulfill the contract s_{-D}. Based on Lemma 1, we now establish the following threshold structure of $b_+^*(x)$:

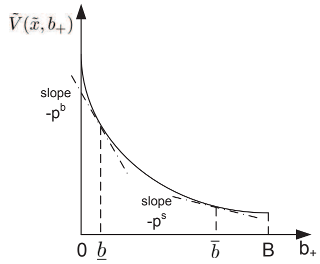

Fig. 1. Illustration of the thresholds $\underline{b}(\tilde{x})$ and $\bar{b}(\tilde{x})$, determined by p^b, p^s and $\tilde{V}(\tilde{x}, b_+)$. Note that $\tilde{V}(\tilde{x}, b_+)$ is typically negative as expected profit is typically positive.

Theorem 2: Let g_f and g_r be given by (4) and (5). There exists two thresholds $\underline{b}(\tilde{x})$ and $\bar{b}(\tilde{x})$, such that

$$\begin{cases} b_+^*(x) = \underline{b}(\tilde{x}), & \text{if } e \le \underline{b}(\tilde{x}), \\ b_+^*(x) = e, & \text{if } \underline{b}(\tilde{x}) \le e \le \bar{b}(\tilde{x}), \\ b_+^*(x) = \bar{b}(\tilde{x}), & \text{if } e \ge \bar{b}(\tilde{x}), \end{cases} \tag{11}$$

where $\underline{b}(\tilde{x})$ and $\bar{b}(\tilde{x})$ satisfy

$$0 \le \underline{b}(\tilde{x}) \le \bar{b}(\tilde{x}) \le B, \tag{12}$$

$$\forall b_+ \in [0, B], \tilde{V}(\tilde{x}, b_+) \ge \tilde{V}(\underline{b}(\tilde{x})) + p^b(b_+ - \underline{b}(\tilde{x})), \tag{13}$$

$$\forall b_+ \in [0, B], \tilde{V}(\tilde{x}, b_+) \ge \tilde{V}(\bar{b}(\tilde{x})) + p^s(b_+ - \bar{b}(\tilde{x})). \tag{14}$$

Proof: We prove the case when $\tilde{V}(\tilde{x}, b_+)$ is differentiable with respect to b_+. The proof can be straightforwardly extended to the non-differentiable case.

From Lemma 1, $\frac{\partial \tilde{V}(\tilde{x}, b_+)}{\partial b_+}$ is a non-positive and non-decreasing function of b_+. Let

$$\underline{b}(\tilde{x}) = \begin{cases} 0, \text{ if } \forall b_+ \in [0, B], \frac{\partial \tilde{V}(\tilde{x}, b_+)}{\partial b_+} + p_b > 0, \\ B, \text{ if } \forall b_+ \in [0, B], \frac{\partial \tilde{V}(\tilde{x}, b_+)}{\partial b_+} + p_b < 0, \\ b_+ \text{ such that } \frac{\partial \tilde{V}(\tilde{x}, b_+)}{\partial b_+} + p_b = 0, \text{ otherwise,} \end{cases} \tag{15}$$

and

$$\bar{b}(\tilde{x}) = \begin{cases} 0, \text{ if } \forall b_+ \in [0, B], \frac{\partial \tilde{V}(\tilde{x}, b_+)}{\partial b_+} + p_s > 0, \\ B, \text{ if } \forall b_+ \in [0, B], \frac{\partial \tilde{V}(\tilde{x}, b_+)}{\partial b_+} + p_s < 0, \\ b_+ \text{ such that } \frac{\partial \tilde{V}(\tilde{x}, b_+)}{\partial b_+} + p_s = 0, \text{ otherwise.} \end{cases} \tag{16}$$

With the above $\underline{b}(\tilde{x})$ and $\bar{b}(\tilde{x})$, the optimality of (11) can be verified by computing the first order condition of (10). Furthermore, because of the convexity of $\tilde{V}(\tilde{x}, b_+)$ in b_+, (12), (13) and (14) are supporting hyperplane conditions that are equivalent to the first order conditions (15) and (16).

An illustrative example is depicted in Figure 1. In other words, a WPP always projects its remaining energy e (after fulfilling current delivery commitment s_{-D}) into the interval $[\underline{b}, \bar{b}]$ as the next battery level. The two thresholds have a clear intuition as follows. If we have less than $\underline{b}$ energy, the marginal benefit of getting more energy to store in the battery is higher than the cost of buying it from the real time market, thus justifying buying energy to fill the battery up to $\underline{b}$. If we have greater than $\bar{b}$ energy, the marginal benefit of selling energy to the real time market is higher than storing it in the battery, thus justifying selling any excess energy above the battery level $\bar{b}$.

IV. Computing Approximate Optimal Policies

The optimality equation (6) is in general hard to solve. In this section we look at methods of solving the optimal policies approximately.

A. Discrete optimal policies

This approximation involves discretizing both the state and action space and the support of the random price and wind processes into a finite number of levels, and solving the optimality equation in (6) for the resulting finite discrete system. The complexity of this approach for a fixed $\beta < 1$ is $\Theta(|\mathcal{X}||\mathcal{A}|)$. Note that once we solve for the optimal value function V^*, we can find out the optimal policy $u^*(x)$.

B. Affine and Look Ahead Policies

Solving the discrete DP exactly has a high complexity if the number of the discretization levels and the forward contract horizon D are large, as the state space grows exponentially with D. For these reasons, we look at affine and look ahead policies which are much simpler to compute. Specifically, if the state is represented by a vector as $X_t = [w_t, b_t, s_{t-D}^{t-1}]$ and control inputs $U_t = [b_{t+1}, s_t]$, the Affine Policy (AP) heuristic is of the form

$$U_t \quad = \quad K_t X_t + k_t, \tag{17}$$

where $K_t \in \mathbf{R}^{2\times(2+D)}$ and $k_t \in \mathbf{R}^{2\times 1}$.

We consider the following approach to obtain such an affine policy via finding the optimal policy of a Linear Quadratic (LQ) stochastic control problem. Specifically, in the dynamic program defined in (2), consider the following cost functions associated with the stage cost in (3):

$$g_f(p_t^f, s_t) = h_{f,1}(\mathbb{E}[p_t^f], B, p_t(w_t)) \cdot (s_t + h_{f,2}(\mathbb{E}[p_t^f], B, p_t(w_t)))^2 \tag{18}$$

$$g_r(p_t^b, p_t^s, x) = h_{r,1}(\mathbb{E}[p_t^b], \mathbb{E}[p_t^s], B, p_t(w_t)) \cdot (x + h_{r,2}(\mathbb{E}[p_t^b], \mathbb{E}E[p_t^s], B, p_t(w_t)))^2 \tag{19}$$

$$\mathbf{I}_{[0,B]}(b_{t+1}) = \gamma \left(b_{t+1} - \frac{B}{2} \right)^2. \tag{20}$$

Here convex quadratic cost functions are employed whose coefficients are chosen as functions of the expectations of the price processes, the battery capacity and the statistics of the wind process. Furthermore, the indicator function is approximated by a quadratic function. The only randomness now considered by the system is wind z_t which is assumed to arise from the same independent periodic distribution. Only the mean of the z_t in future time steps has an impact.

The dynamics of the system is linear as mentioned in Section II-B. The convex quadratic stage cost and linear dynamics of the modified program can then be written as:

$$g_t^{\text{quad}}(X_t, U_t) = \begin{bmatrix} X_t \\ U_t \\ 1 \end{bmatrix}^\intercal \begin{bmatrix} Q_t & q_t \\ q_t^\intercal & r_t \end{bmatrix} \begin{bmatrix} X_t \\ U_t \\ 1 \end{bmatrix}$$

$$X_{t+1} = A_1 X_t + A_2 U_t + A_3 z_t, \tag{21}$$

with appropriate Q_t, q_t, and A_i

This modified DP is an LQ problem with a convex quadratic value function, denoted by $\bar{V}(x)$, and the optimal policy is affine. The Bellman optimality equation results in a set of algebraic Ricatti equations which can be solved in closed form to derive the parameters K_t and k_t in the affine policy (17). The policies derived above are optimal when the following model assumptions hold:

- The price processes are deterministic,
- The forward and real time market costs are quadratic.
- The constraint that the battery should strictly be within capacity is replaced by a quadratic penalty for battery energy levels.

In practice, the above assumptions may not hold. In these cases, the developed affine policy serves as a heuristic to solve the original problem. Furthermore, we can also use the LQ quadratic value function $\bar{V}(x)$ to obtain a look ahead (LA) policy as follows [18]:

$$u_t = \min_u g(x_t, u) + \beta\theta\bar{V}(f(x_t, u, \mathbb{E}z_t)), \quad (22)$$

where θ is a parameter that we can tune.

V. Simulation Results

In this section, we numerically evaluate the benefits of having a battery in a wind farm using the schemes outlined in the previous sections.

A. Simulation Setup

We consider a WPP participating in hourly day-ahead markets and hourly real time markets. The natural choice of D would thus be 24 since a day-ahead forward contract commits power to be delivered 24 hours later. In our simulation, to simplify the computation while still preserving the characteristics of the real world wind and price processes, we consider the case of $D = 4$ in the following way. One day is divided into four 6-hour blocks starting from 12am: early morning, morning, afternoon, and evening. At any given decision block (of 6 hours), the WPP decides on a fixed contract level $D = 4$ blocks later (the contract commitments within a decision block are fixed) . This reduction to $D = 4$ is done due to daily patterns of typical wind and price processes.

We investigate the total discounted profit that the WPP gets with a battery and a discount factor of $\beta = 0.99$. We simulate the developed policies with two types of wind and price data: a) simulated data generated with our model assumptions (cf. Section II-A) of the discrete model whose parameters are learned from real world data, and b) real world data. The idea is that, with simulated data, the assumed model is accurate for the discrete model, and we can derive an optimal policy corresponding to this, whereas with real world data there is always some mismatch between our assumed model and the actual data. Specifically, for the discrete optimal policy, we divide the wind, price processes each into three levels. The choices and statistics of these discrete levels are learned from real world wind and price data in the PJM interconnection [19]. In particular, the wind power w_t is quantized to be one of $17, 50, 83$ MW, and the forward power price p_t^f is quantized to one of $40, 80, 120$ \$/MWh. For simplicity, p_t^b is taken to be $2 \cdot p_{t-D}^f$, and p_t^s to be $\frac{p_{t-D}^f}{2}$.

Fig. 2(a) plots the total discount profit as a function of battery capacity for the different approximate policies described in Section IV, assuming both wind and price processes randomly vary. Fig. 2(b) shows the same total discount profit as a function of battery capacity but now assuming constant price processes, where the price levels are taken to be equal to the empirical means of the price data. These simulations are averaged over 100 realizations. Finally, Fig. 2(c) plots the total discount profit as a function of battery capacity with real world wind and price data from January to February 2004 [19]. The discrete optimal policy employs the same wind and price forecast parameters as it did for the simulated data scenario. The LQ models employ only the expected wind and price values.

In each figure, we evaluate the discrete, affine and look ahead policies developed in the previous section. In order to get a sense of the optimality of the policies, a genie upper bound on the maximum achievable discounted profit is developed by optimizing the actions with the following assumptions: a) all realizations of the random quantities are known ahead of time, and b) continuous battery and contract levels can be set. It is immediate to check that computing such a genie upper bound is a convex optimization problem, and can be solved efficiently.

B. Managing uncertainty and exploiting price variations

From all three figures, it is clear that having more battery capacity leads to a higher profit. We further see that the curves in Fig. 2(b) are evidently lower than the corresponding curves in Fig. 2(a). This demonstrates that battery capacity not only helps a WPP manage wind uncertainty, but can also help it take advantage of price differences over time. Note that the latter opportunity is eliminated if prices are constant, which is the assumption made in Fig. 2(b).

Furthermore, Fig. 2(a) demonstrates that the discrete optimal policy has a performance that is fairly close to the genie bound. It is also observed that the look ahead policy performs much closer to the genie bound than the affine policy. The look ahead policy also outperforms the discrete-approximation when the battery capacity becomes large. This is because our discretization error increases as the battery size grows large.

C. Performance with real data

Finally, we compare the performance of our policies with the genie bound on real world data in Fig. 2(c). In this figure, the performance gap between any given heuristic policy and the genie upper bound becomes significant. In particular, the discrete approximation sees the worst performance penalty. The gaps highlight the importance of model accuracy. Model mismatch is both due to the fact that our discrete policy has to quantize the continuous system state into rather coarse levels and also because the distributions on the real world data differ from our modeling assumptions (i.e., independent periodic processes). While the achieved discounted profit of the policies based on the LQ approximation is seen to be worse than that of the discrete optimal policy for Fig. 2(a) and

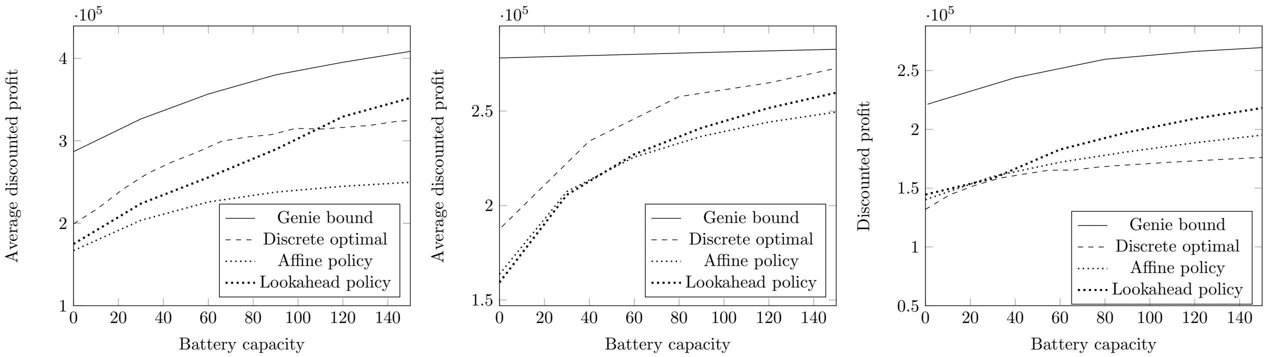

(a) (b) (c)

Fig. 2. Average discounted profit as a function of battery capacity. (a) Simulated periodic random wind and price processes. (b) Simulated periodic random wind process and constant price processes. (c) Real world wind and price data.

2(b), it is seen to be better than the discrete optimal policy on this trace of real world data, suggesting that it is more robust to modeling imperfections. Investigating ways to incorporate these insights into robust policies remains an interesting topic for future work.

VI. Conclusion

We have considered a WPP with co-located energy storage that participates in a conventional two-settlement market. To maximize its profit, the WPP optimizes its forward contract D time slots ahead of delivery as well as its storage operation simultaneously over time. An infinite horizon discounted cost minimization problem is formulated as a dynamic program that includes the past D time slots' forward contracts in the state space. We proved that the optimal battery operation policy admits a threshold structure: there exists two thresholds, $\underline{b}$ and $\overline{b}$, with which a WPP always projects its remaining energy (after fulfilling current delivery commitment) into the interval $[\underline{b}, \overline{b}]$ as the next battery level. The thresholds $\underline{b}$ and $\overline{b}$ are, however, difficult to compute numerically, motivating us to look at approximate solutions.

We developed two heuristics for policy computations based on model approximations. The first heuristic discretizes the state and action space of the original policy. The second heuristic is based on affine and look ahead policies that solve a Linear Quadratic (LQ) stochastic controller based on quadratic approximations of the objective and constraints of the DP. We evaluated the approximate policies with both simulated and real world wind and price data. A genie bound on the optimal policy is also computed for bounding the performance gaps of the approximate policies to optimality. With the simulated data, the discrete and look ahead policies have close-to-optimal performance. With real world data, as there is a discrepancy between the wind and price models we assumed and the actual data, the performance of the developed policies has a larger gap to the genie bound, with the look ahead policy performing reasonably well.

References

[1] U.S. Energy Information Administration, "Annual energy outlook," Tech. Rep., 2014.

[2] C. Klessmann, C. Nabe, and K. Burges, "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*, vol. 36, no. 10, pp. 3646–3661, Oct. 2008.

[3] P. Pinson, "Wind energy: Forecasting challenges for its operational management," *Statistical Science*, vol. 28, no. 4, pp. 564–585, 2013.

[4] Y. Zhao, J. Qin, R. Rajagopal, A. Goldsmith, and H. V. Poor, "Wind aggregation via risky power markets," *to appear in IEEE Transactions on Power Systems*, 2014.

[5] H. L. Ferreira, R. Garde, G. Fulli, W. Kling, and J. P. Lopes, "Characterisation of electrical energy storage technologies," *Energy*, vol. 53, pp. 288–298, 2013.

[6] E. Y. Bitar, R. Rajagopal, P. P. Khargonekar, K. Poolla, and P. Varaiya, "Bringing wind energy to market," *Power Systems, IEEE Transactions on*, vol. 27, no. 3, pp. 1225–1235, 2012.

[7] M. Zugno, T. Jónsson, and P. Pinson, "Trading wind energy on the basis of probabilistic forecasts both of wind generation and of market quantities," *Wind Energy*, vol. 16, no. 6, pp. 909–926, 2013.

[8] J. M. Morales, A. J. Conejo, and J. Pérez-Ruiz, "Short-term trading for a wind power producer," *Power Systems, IEEE Transactions on*, vol. 25, no. 1, pp. 554–564, 2010.

[9] D. R. Jiang and W. B. Powell, "Optimal hour ahead bidding in the real time electricity market with battery storage using approximate dynamic programming," *arXiv preprint arXiv:1402.3575*, 2014.

[10] P. Harsha and M. Dahleh, "Optimal management and sizing of energy storage under dynamic pricing for the efficient integration," *to appear in IEEE Transactions on Power Systems*, 2014.

[11] J. Qin, Y. Chow, J. Yang, and R. Rajagopal, "Online modified greedy algorithm for storage control under uncertainty," *arXiv preprint arXiv:1405.7789*, 2014.

[12] E. Bitar, R. Rajagopal, P. Khargonekar, and K. Poolla, "The role of co-located storage for wind power producers in conventional electricity markets," in *American Control Conference (ACC)*. IEEE, 2011, pp. 3886–3891.

[13] J. H. Kim and W. B. Powell, "Optimal energy commitments with storage and intermittent supply," *Operations research*, vol. 59, no. 6, pp. 1347–1360, 2011.

[14] S. Kurandwad, C. Subramanian, A. Vasan, V. Sarangan, V. Chellaboina, A. Sivasubramaniam *et al.*, "Windy with a chance of profit: bid strategy and analysis for wind integration," in *Proceedings of the 5th international conference on Future energy systems*. ACM, 2014, pp. 39–49.

[15] D. P. Bertsekas, D. P. Bertsekas, D. P. Bertsekas, and D. P. Bertsekas, *Dynamic programming and optimal control*. Athena Scientific Belmont, MA, 1995, vol. 1, no. 2.

[16] A. Shapiro, D. Dentcheva *et al.*, *Lectures on stochastic programming: modeling and theory*. SIAM, 2009, vol. 9.

[17] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge Press, 2004.

[18] W. B. Powell, "Clearing the jungle of stochastic optimization," *Informs Tutorials in Operations Research: Bridging Data and Decisions*, pp. 109–137, 2014.

[19] National Renewable Energy Laboratory, "Eastern wind integration and transmission study," Tech. Rep., 2010.