

Supplying Renewable Energy to Deferrable Loads: Algorithms and Economic Analysis

Anthony Papavasiliou
Department of Industrial Engineering
and Operations Research
4178 Etcheverry Hall
UC Berkeley
Berkeley, California 94709
Email: tonypap@berkeley.edu

Shmuel S. Oren
Department of Industrial Engineering
and Operations Research
4141 Etcheverry Hall
UC Berkeley
Berkeley, California 94709
Email: oren@ieor.berkeley.edu

Abstract—In this paper we propose a direct coupling of renewable generation with deferrable demand in order to mitigate the unpredictable and non-controllable fluctuation of renewable power supply. We cast our problem in the form of a stochastic dynamic program and we characterize the value function of the problem in order to develop efficient solution methods. We develop and compare two algorithms for optimally supplying renewable power to time-flexible electricity loads in the presence of a spot market, backward dynamic programming and approximate dynamic programming. We describe how our proposition compares to price responsive demand in terms capacity gains and energy market revenues for renewable generators, and we determine the optimal capacity of deferrable demand which can be reliably coupled to renewable generation.

I. INTRODUCTION

Renewable power is emerging as a mainstream source of energy supply in power systems. Various policy thrusts are promoting the advent of renewable power in the United States and around the world. Twenty four states and the district of Columbia have set renewable portfolio standards, which commit electric utilities to procure at least a certain percentage of their energy from renewable energy sources. California set the example in 2006 by establishing a renewable portfolio standard which now requires that the state cover 33% of its electricity demand from renewable energy sources. In an effort to coordinate efforts at a national level, the federal government voted for the American Clean Energy and Security Act in 2009. Among its various measures, the legislation sets a 20% renewable electricity standard at the national level, and requires that US emissions be reduced by 17% compared to their 2005 levels. These legislative measures signal the determination of the United States, the lead consumer of energy globally, to utilize renewable energy at an unprecedented scale.

Integrating large amounts of renewable energy in power systems presents a host of new technological challenges to power systems operations. In this paper we focus on exploiting flexibility in electricity consumption in order to balance renewable energy supply. Renewable energy sources such as wind power and solar power are supplied in unpredictable and highly variable rates. We cannot forecast renewable generation accurately, and even if we could these energy sources are

highly variable. These inherent characteristics of renewable energy sources place them in a competitive disadvantage compared to traditional fossil fuel energy sources.

A. The nature of the problem

The scheduling of power system operations is highly complicated by the requirement of maintaining a continuous balance between the supply and demand of electricity, in order to prevent instabilities in the grid. The scheduling of power system resources is already a highly complex task, and introducing renewable energy sources in significant amounts further complicates this task.

The unpredictability of renewable power supply may cause imbalances to the system which require expensive deviations from day-ahead dispatch schedules. Starting up or shutting down units to compensate for a sudden change in renewable power supply may take hours, lead to additional air pollution, result in wear and the need for frequent maintenance of startup units, and upset system dispatch due to the minimum generation capacity of startup units. Such large scale disturbances in renewable power supply can occur during storms in systems with large amounts of wind power.

The minute-by-minute variability of renewables imposes a requirement for primary control, generators which can rapidly adjust their power output in response to an unanticipated event. Moreover, secondary control units are necessary which can back up and dismiss primary control units. Since renewable supply also tends to vary rapidly and in great magnitude, an additional backup of ramping generators is necessary.

In a British report by the UK Energy Research Center [1] the authors assembled a variety of wind power integration studies with the objective of estimating the costs and impacts of intermittent generation on the UK electricity network. Over 80% of the studies that the authors examined concluded that for wind power energy penetration levels above 20% an investment in system backup in the range of 5-10% of installed wind capacity is required in order to balance the short term (seconds to tens of minutes) variability of wind power supply. The authors conclude that additional conventional capacity to

maintain system reliability during demand peaks amounts to 15-22% of installed wind power capacity.

The California Independent System Operator published a report recently [2], which analyzes the integration of 6700 MW in the California grid. According to the study the 3-hour morning ramp of will increase by 926 MW to 1529 MW due to the fluctuations of wind generation at the time when morning demand increases, and the evening ramp will decrease by 427 MW to 984 MW. The regulation capacity requirement will increase by 170 to 250 MW for regulation up and by 400 to 500 MW for regulation down. The regulation ramping will increase by ± 15 to ± 25 MW/min. The load following ramps will increase by ± 30 to ± 40 MW/min.

Various systems absorb large amounts of hydroelectric power. During the months that snow melts and hydroelectric power supply increases and must be absorbed, the additional generation of renewable energy causes an over-supply problem. It is also possible that renewable energy supply increase during the night and abate during daytime, hence renewable generation is negatively correlated with electricity consumption.

B. The Effects of the Problem

The costs associated with renewable power integration result from the offset of variability by stand-by generators and the requirement for investments on system backup. These costs are captured by market tariffs and may be allocated to the whole market or directly to renewable generators, depending on market regulations. Research and experience indicate that integration costs range between 0 and 7\$/MWh [3], [4]. The UK study mentioned above [1] placed an estimate of no more than 5 British pounds for wind power integration. Another recent study conducted by Enernex for wind power integration in Minnesota [5] concludes that the cost of additional reserves and costs related to variability and day-ahead forecast errors will result in an additional \$2.11 (15% penetration) to \$4.41 (25% penetration) per MWh of delivered wind power. In a similar vein, the CAISO report [2] has predicted an expected increase in 10-minute real time energy prices due to wind forecasting errors which become comparable to load forecasting errors.

Renewable energy may be discarded during hours of excess renewable power supply if power systems cannot reliably absorb this supply [3], [6]. During early spring the California system operator either spills water supplies from hydroelectric dams or discards wind power [2]. Wind power is also discarded under normal operating conditions in California whenever forecasting underestimates the amount of wind power supply to the system and the excess power cannot be sold. In Texas the system operator discards wind power during load pick-up for reliability reasons [7].

Though the integration of renewable energy is increasing, an integration level beyond 20% is not perceived as economical (integration levels count 20% in Denmark, 9% in Spain, 7% in Germany, and California is aiming for 33% by 2030). Assuming capital costs for renewable power will

continue to decline in the future, one of the major challenges for the large scale integration of renewable energy will be its variability. Currently renewable generators operate under favorable regulations in many markets. A number of system operators in Europe (Denmark, Greece) and the United States (PJM, NYISO, CAISO, Ontario IMO) accept wind generation on a priority basis [8]. It is clear that this preferential treatment has its limitations. Large scale renewable power integration cannot rely on regulatory support alone, but will also require technological improvements. The utilization of demand side flexibility creates an excellent opportunity for addressing this problem.

C. Demand side flexibility

Across the full spectrum of residential, commercial and industrial consumption, a significant proportion of the power that we generate is supplied to loads which are time flexible, deferrable for a few minutes or hours at little or no cost. Examples abound: electric vehicles, heating, ventilation, air conditioning, thermostats, refrigeration, agricultural pumping, controllable lighting. These time flexible demand side resources could adapt their energy consumption according to the fluctuation of renewable power supply in order to counter-balance the resulting supply variations and enable large scale integration of renewable energy without significant impacts on grid operations.

With the appropriate communications and control infrastructure in place, flexible loads can be manipulated as controllable resources by the system operator, in much the same way that generators are actively controlled today [9]. In the same fashion that generators communicate a set of operating characteristics to the system operator (marginal cost, minimum and maximum generation limits, minimum and maximum ramping limits, minimum up and down time) which then determines their dispatch such that the cost of operating the system is minimized, loads could also declare certain parameters to the system operator which characterize their flexibility, such as their required energy demand and a deadline by which this demand should be met. These load resources can then be dispatched in a least-cost fashion, such that their demand is met within their designated deadline. Due to the fact that resources such as wind and solar power operate at near zero marginal cost, renewable energy sources are an excellent candidate for fulfilling such flexible energy requests.

In order to address the problems raised above, we propose a paradigm whereby the demand of flexible consumers is regulated by a central scheduler which receives requests for energy consumption within a certain deadline, and decides how the available renewable resources are allocated to consumers. The scheduler commits to satisfy consumers by their deadlines, if necessary by resorting to an electricity spot market in order to procure energy at the last minute. This gives rise to a stochastic optimal control problem.

II. ALGORITHMS

In this section we define the stochastic dynamic programming problem at hand, certain properties of the value function and two dynamic programming algorithms for solving the problem. We then compare the performance of these algorithms in terms of computation time and performance.

A. Problem Formulation

Consider a renewable power supplier which has entered an agreement to supply a certain amount of energy to a customer within a certain deadline. The renewable resource supplier has a contractual obligation to fully satisfy customer demand, either through renewable energy supply or through spot market purchases. The objective is to determine the optimal spot market strategy for the supplier, i.e. when it is worth procuring energy from the spot market in order to satisfy residual demand, and how much energy should be procured at each period.

Our problem has a three-dimensional state vector, $x_t = (\lambda_t, s_t, r_t)$, where λ_t is the spot price of the resource, s_t is the amount of resource which is freely available and r_t represents the remaining quantity of demand. In what follows we will also use the notation $x_t(i)$ to denote the i -th coordinate of the state vector for period t . Hence, $\lambda_t = x_t(1)$, $s_t = x_t(2)$ and $r_t = x_t(3)$. The residual energy r_t evolves according to $r_{t+1} = r_t - u_t$, where u_t , our control, is the amount of power supplied to the consumer in period t . We assume that the two-dimensional stochastic process (λ_t, w_t) can be described by a non-stationary Markov transition probability matrix, $\mathbb{P}_t[\lambda_{t+1} = \lambda', s_{t+1} = s' | \lambda_t = \lambda, s_t = s]$, which we denote generically as $p_t(x, x')$, where x denotes the current period state vector and x' denotes the state vector of the next period.

The objective is to minimize the following expected cost:

$$\min_{\mu_t(x_t)} \mathbb{E} \left[\sum_{t=1}^{N-1} \lambda_t (\mu_t(x_t) - s_t)^+ \right] \Delta t, \quad (1)$$

where $\mu_t(x)$ represents the rate at which the resource is supplied and N is the number of periods. The state vector is associated with the following initial and terminal conditions: $r_1 = R$, $r_N = 0$, where R is the amount of demand to be satisfied. The control u_t cannot exceed an upper bound on the rate of supply, $u_t \leq C$. Each interval of the problem has a duration of Δt units of time.

Although the problem has been cast from the point of view of a renewable power supplier, one can consider the same problem facing a smart switch installed in any flexible energy consuming device (e.g. a pool pump or a refrigerator) which responds to renewable energy signals. Moreover, we present here the problem of a single customer, but the same problem applies for the case of multiple customers with identical quantities of energy demand R , deadlines N , and rate constraints C .

B. Backward Dynamic Programming

The backward dynamic programming algorithm is given by the following equation:

$$J_t(x) = \max_{u \in U_t(x)} \{g(x, u) + \sum_{y \in S_{t+1}} p_t(x, y) J_{t+1}(y)\} \quad (2)$$

where $J_t(x)$ is the value function of period t , $g(x, u) = \lambda(u - s)^+ \Delta t$ is the cost incurred at each period, $U_t(x)$ is the feasible region of actions for period t , and S_t is the feasible region of the state vector at period t . Backward dynamic programming can yield the optimal policy in principle, however in the worst case it can require as many as $N \cdot |S|^2 \cdot |U|$ operations [10], where $|S|$ is the cardinality of the state space, and $|U|$ is the cardinality of the action space, assuming the state and action spaces have equal cardinality for each period.

C. Structure of the Value Function

We now present two results about the structure of the value function in the case where the action and state space are continuous, which assist us in the development of an approximate dynamic programming algorithm for this problem. For the case of continuous action and state spaces, our actions are constrained in the following interval: $U_t(x_t) = [a_t(x_t), b_t(x_t)]$, where $b_t(x_t) = C \wedge r_t$ and $a_t(x_t) = b_t(x_t) \wedge (r_t - C(N - t - 1))^+$.

Proposition 1: The value function $J_t(x)$ is convex in $r = x(3)$ for all t .

Proof:

We will prove the argument by induction, starting from period $N - 1$. We know that

$$J_{N-1}(x_{N-1}) = \lambda_{N-1}(r_{N-1} - s_{N-1})^+ \Delta t, \quad (3)$$

so the hypothesis holds for $N - 1$.

Now suppose that the hypothesis is true for all k up to $k \geq t + 1$. Consider the Q -factor at period t , $Q_t(x_t, u)$:

$$Q_t(x_t, u) = \lambda_t(u - s_t)^+ \Delta t + \mathbb{E} J_{t+1}(\lambda_{t+1}, s_{t+1}, r_t - u). \quad (4)$$

We now use various convexity preservation arguments from [11]. Since $J_{t+1}(\lambda, s, r)$ is convex in r and $r - u$ is convex in (r, u) , $V_{t+1}(\lambda, s, r - u)$ is convex in (r, u) . The expectation operator preserves convexity so $\mathbb{E} J_{t+1}(\lambda_{t+1}, s_{t+1}, r_t - u)$ is convex in (r, u) . $\lambda(u - s)^+ \Delta t$ is also convex in (r, u) . Since we are minimizing a convex function in (r, u) over u , with u constrained in the convex set $U_t(x_t)$, the resulting function $J_t(x_t)$ is convex in r and the desired result follows. ■

Proposition 2: Suppose that there are finitely many random outcomes in each period. Then the value function in period t , $J_t(x)$, is piecewise affine convex in r :

$$J_t(\lambda, s, r) = \sum_{i=1}^{n(t)} a_i(t) r + b_i(t) \quad (5)$$

for some constants $a_i(t)$, $b_i(t)$, $i = 1, \dots, n(t)$.

Proof:

Again, we will prove the argument by induction. From equation 3 the induction hypothesis holds for period $N - 1$. Suppose that the hypothesis holds for all periods up to period $t + 1$ and denote Ω_t as the set of random outcomes in period t . Then for period t we have from the Bellman equation:

$$\begin{aligned}
J_t(\lambda, s, r) &= \min_{u \in U_t(x_t)} \{ \lambda(u - s)^+ \Delta t \\
&\quad + \sum_{\omega \in \Omega_t} p_t(\omega) \sum_{i=1}^{n(t)} (a_i(t)(r - u) + b_i(t)) \} \\
&= \lambda(u^* - s)^+ \Delta t \\
&\quad + \sum_{\omega \in \Omega_t} p_t(\omega) \sum_{i=1}^{n(t)} (a_i(t)(r - u^*) + b_i(t)) \\
&= \sum_{\omega \in \Omega_t} p_t(\omega) \sum_{i=1}^{n(t)} (a_i(t)(r - u^*) + b_i(t)) \\
&\quad + \lambda(u^* - s)^+ \Delta t
\end{aligned} \tag{6}$$

where u^* is the minimizer of the right hand side, which must exist since we are minimizing a convex function over a line segment. Now the result follows from the fact that the weighted sum of piecewise linear convex curves is a piecewise linear convex curve. ■

D. Approximate Dynamic Programming

In order to scale our control policy to more complex conditions with multiple loads, random arrival and departure times, and random energy requests, we develop an approximate dynamic programming algorithm which can perform adequately in this simple problem. Although backward dynamic programming is perfectly adequate for solving this problem, it will become computationally intractable to consider the more complex conditions which we described, but backward dynamic can still be useful in providing a benchmark to compare our approximate algorithm for this simple instance of the problem.

We will approximate the Q -factors of the problem [12] with a set of basis functions, $Q(x, u) = \sum_{k \in K} r_k \phi_k(x, u)$, where K is the set of basis functions, $\phi_k(x, u)$ are the basis functions and r_k are the weights of the bases.

We have selected the basis functions for this problem as follows: $\phi_1(x, u) = 1$ is the constant function. $\phi_2(x, u) = (x(3) - u - (N - t)C)^+$ is the amount of remaining energy that would need to be fulfilled if we were to charge at a rate u for the current period and charge at full rate after that. $\phi_3(x, u) = x(1)(u - x(2))^+ \Delta t$ is the cost incurred at the current period, $g(x, u)$. We have also created one basis function for each state and the time index, $\phi_4(x, u) = k$, $\phi_5(x, u) = x(1)$, $\phi_6(x, u) = x(2)$, $\phi_7(x, u) = x(3)$. Finally, we have created a set of basis functions which can be used to create a piecewise linear approximation of the Q -factors in the third coordinate of the state space, $\phi_{k+8}(x, u) = (r - u - kC)^+$, $k = \{0, \dots, 14\}$. This is motivated by the result which we have proven about the structure of the value function.

We will implement the SARSA algorithm [12] for this problem with a linear approximation for the Q -factors and a

temporal difference algorithm for the update of the basis function weights. For the learning rate of the basis function weights we have chosen a step-size of $\gamma(t) = 10,000/(10,000 + t - 1)$, scaled appropriately [13]. We are also using Boltzmann exploration in order to facilitate the exploration of the state space, with a Boltz exploration constant of $\beta = 10^{-4}$. The SARSA algorithm consists of the following steps:

- Select control u_t with probability
$$\frac{\exp(-\beta \sum_k r_k \phi_k(x_t, u_t))}{\sum_{u \in U_t(x_t)} \exp(-\beta \sum_k r_k \phi_k(x_t, u_t))}$$
- Given a state/action pair (x_t, u_t) , generate a new state x_{t+1} according to the transition distribution $p_t(x_t, x_{t+1})$.
- Update the basis weights,

$$\begin{aligned}
r_k &= r_k + \gamma(t) \phi_k(x_t, u_t) (g(x_t, u_t) \\
&\quad + \sum_{k \in K} r_k \phi_k(x_{t+1}, u_{t+1}) - \sum_{k \in K} r_k \phi_k(x_t, u_t))
\end{aligned} \tag{7}$$

E. Data

In order to calibrate our stochastic model, we have used wind data from the National Renewable Energy Laboratory database for 10-minute wind generation at a typical site in the Techahapi region, and 10-minute settlement average prices from the Oasis database for the dates between September 1, 2006, and November 30, 2006.

The wind generation data is based on the output of a park of ten Vestas V90 3 MW wind generators. In particular, we used the hysteresis-corrected SCORE data, as explained in the NREL website¹.

The wind generation data ranges from 0 to 30 MW, and the price data ranges from -20\$/MWh to 180\$/MWh. For the derivation of the probability transition matrix, price values below -20\$/MWh and above 180\$/MWh are assumed to be equal to -20\$/MWh and 180\$/MWh, respectively. Given that such prices did not occur frequently in our dataset, this modeling assumption does not affect the derived stochastic model significantly. The transition probability matrix is derived by sampling the conditional probability of transitioning from one price-wind combination to any other price-wind combination for each 10-minute interval of the day. For the greatest possible resolution that we explored, 10 wind states and 10 price states, there are 100 wind-price combinations, and a sample of 91 transitions.

F. Aggregating States and Actions

As we described in section II-B, computation time is sensitive to the size of the state and action spaces. There is a tradeoff involved in aggregating states and actions: by aggregating we sacrifice in terms of the performance of the optimal solution in order to achieve faster computation of the optimal policy. Aggregation implies that we may be making decisions of lower quality due to the fact that we are not able to finely distinguish between different values of wind power and

¹<http://www.nrel.gov/wind/integrationdatasets/western/methodology.html>
#output

TABLE I

COST PERFORMANCE AND COMPUTATION TIME FOR VARIOUS LEVELS OF STATE AND ACTION SPACE AGGREGATIONS.

$ U $	N_w	N_p	$ S $	Sol. time (s)	Cost (\$)	St. dev. (\$)
10	10	10	89,200	812.365	11,600	8,251
5	5	5	9,925	199.399	12,137	8,365
2	5	5	2,500	72.915	12,601	8,131
2	3	3	900	42.904	12,791	8,228

market prices, or even if we are able to distinguish between states, we may not be able to differentiate our control. This tradeoff is presented in table I. In the experiments that we have run to derive this table, we assume that the maximum control is 30 MW, and we are set to fulfill a total demand of 2970 MW-10min. Wind ranges between 0 and 30 MW, and prices range between -20 and 180\$/MWh. The results derived in the table are derived from testing the algorithm against the actual 91-day sample of data.

As expected, we observe that as we increase the granularity of our state and action space the computation time increases and the performance of the algorithm improves. The variance in the performance of the algorithm does not vary significantly and is not monotonically increasing in the level of aggregation, as one might expect. In fact, the maximum and minimum of cost performance varies between \$200 and \$36,138 and is similar for all levels of aggregation. This very large difference stems from the very different nature of the wind and price outcomes in different days, and increasing the granularity of the state and action spaces does not help improve performance for a given sample outcome of wind and prices. The most surprising and practically relevant conclusion from this table is that aggregation results in a surprisingly low deterioration of performance. Comparing the first and last row of the table we observe that the solution time increases 19-fold from 43 seconds to about 13.5 minutes, yet the cost decreases by merely 9.31%.

An action state consisting of 2 elements, $U = \{0, C\}$, reduces to a switching control where we either turn the load on at its nominal consumption level or turn it off. Most loads can only be operated at an on-off state anyways, and we can conclude from this table that there are minor impacts on our performance by restricting attention to a switching control.

G. Comparison of Algorithm Performance

In table II we present the performance of the approximate dynamic programming algorithm for the same problem instances as those which are presented in table I. The approximate dynamic programming algorithm requires the same amount of time for all cases, since we are keeping the number of iterations fixed at 50,000. For the case where we are solving for 10 actions and 100 combinations of price and wind, the approximate dynamic programming algorithm is 16.7% suboptimal, although it runs in less than 1/3 of the time that the backward dynamic programming algorithm requires. It turns out that increasing the number of iterations does not improve the performance of the algorithm, and the loss in

TABLE II

PERFORMANCE OF SARSA ALGORITHM FOR VARIOUS LEVELS OF STATE AND ACTION SPACE AGGREGATIONS.

$ U $	N_w	N_p	Sol. time (s)	Cost (\$)	St. dev. (\$)
10	10	10	242.576	13,539	10,649
5	5	5	239.944	12,368	8,162
2	5	5	238.224	12,703	8,190
2	3	3	238.109	13,114	9,345

performance can therefore be attributed to the selection of the basis functions. Nevertheless, for higher levels of aggregation the algorithm performs very close to optimal. This reassures us that if we are willing to aggregate the action and state space sufficiently, then we have made a good selection of basis functions, which provides starting ground for scaling the algorithm to more complex versions of the problem, e.g. with random arrival and departure times or random quantities of energy demand. Such variations of the problem cannot, in general, be solved by the backward dynamic programming algorithm efficiently, and are best dealt with by approximate dynamic programming techniques. This is work which we wish to explore further in future research.

III. ECONOMIC ANALYSIS

In this section we discuss the economic implications of directly coupling renewable generation with deferrable demand by comparing it to a baseline scenario whereby renewable generators participate in the market with a reduced capacity credit and deferrable loads participate in the market with the objective of minimizing their expenditures.

In order to assess the value of coupling renewable generation with deferrable loads, we compare it to a case where both resources respond to price signals without actively coordinating. In the case where renewables and flexible consumers coordinate their operations wind appears 'behind the meter' for the system operator, and the spot market is utilized in order to correct for unanticipated deviation in the supply of renewable energy in advance of an emergency shortage in supply, which is an economic alternative to resorting to backup generation for supporting renewable power variability.

A. Capacity Credit

As we discussed in the introduction, renewable generators which funnel their entire production to the grid cause various operational problems due to the fluctuation of their supply. In California this deterioration in performance is penalized by charging renewable generators penalties for deviating from their forecast production. Such deviation penalties can capture the impact of renewable power variability on the economic performance of wind generators, as described also in the literature by Bathurst et al. [14], Matevosyan and Soder [15] and Pinson et al. [16]. In other markets, such as PJM, renewable generators receive partial credit for their available capacity, reflecting the fact that the system operator cannot securely rely on the entire capacity of these resources at critical hours of grid operations, but must instead procure standby capacity. In this paper we use capacity credit as a metric for

the impact of renewable power variability on the economic performance of renewable generators.

We assume that capacity earns a credit of 1440\$/MW, which we have calculated by continuous compounding of a \$2 million investment in 1 MW of new capacity at 8% interest rate for a payback period of 10 years. In the baseline scenario we assume that renewable generators receive a 30% capacity credit. In the case where the resources are coupled, renewable generators earn full capacity credit for the average capacity of the deferrable loads which they serve. However, it is natural to expect that the incremental ability of renewable generators to serve deferrable loads deteriorates as more and more deferrable capacity is reserved by renewable suppliers, since renewable generators will need to resort to the spot market more and more frequently. We model this as a deterioration of the capacity credit for the coupled system. In particular, we simulate the performance of the smart charging algorithm for the entire horizon of the 91 days of sample data, and we count the average amount of power procured by the algorithm for the periods of peak demand. Peak demand periods are identified as those periods of the sample data during which the spot price of electricity exceeded 350\$/MWh. Such an event occurred 45 times in our 91-day data. This average quantity is then subtracted from the average capacity of the deferrable loads, and represents the derated capacity of the coupled system.

B. The Value of Coupling

The value of coupling renewable generation to flexible demand arises from the fact that renewable generators earn capacity credit in proportion to the capacity of the flexible loads which are served by the renewable power provider. In return for this increased credit, renewable generators reserve their entire supply for deferrable loads, and assume the responsibility of providing capacity service to deferrable loads. The tradeoff, then, for renewable generators is to gain capacity credit by contracting with deferrable loads which are sufficiently flexible to be served reliably by smart charging, versus earning energy revenues by supplying electricity to the market at reduced capacity credit. The total demand of deferrable loads which are served determines the balance between the costs and benefits of coupling. In this section we determine the total demand of deferrable loads which maximizes the value of coupling.

By testing our proposition against price responsive demand, we are essentially comparing to the alternative of pooling all resources in the market. By coupling the aggregator receives economic gains by relieving risk from the system operator. Entering a contract to supply power within a deadline is much more flexible than the strict reliability criteria applied by the system operator, which will match any amount of load capacity, regardless of its flexibility, with a corresponding capacity of backup.

In figure 1 we present the net supply of power for coupling versus utilizing price-responsive demand for one sample outcome of prices and wind power supply. In the third frame we show the remaining energy to be fulfilled for both the

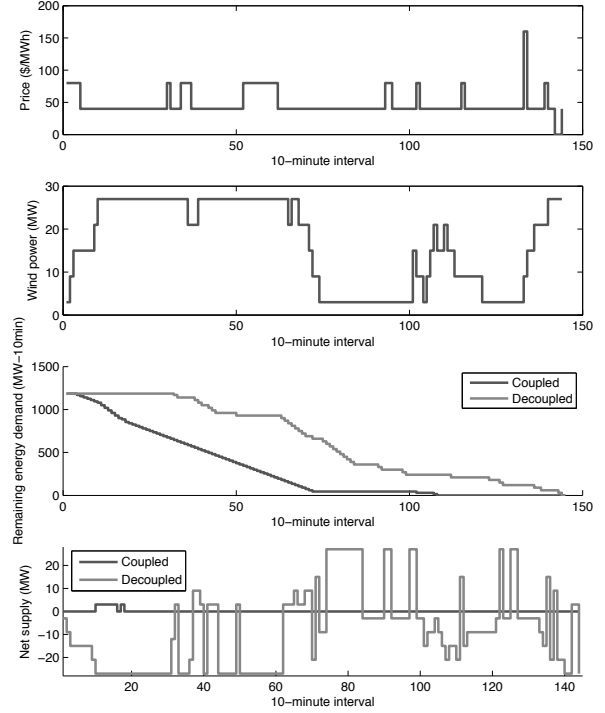


Fig. 1. Coupled operations versus price response.

coupled and decoupled case, and in the fourth frame we show the demand of power from deferrable loads net of the wind power supply. From the fourth frame we observe that the net demand profile of the coupled system is much smoother than the one resulting from price-based response. By observing the difference in the energy supply patterns in the lower right frame, we can conclude that the consumption patterns of the loads differ significantly. In the first periods the coupled system utilizes wind power, which is abundantly available, whereas the deferrable loads which simply respond to spot market prices fail to absorb the excess wind power. Towards the end of the horizon, the wind generators continue to supply power in the decoupled case, whereas in the coupled case this wind is discarded because loads have been fully served.

C. Simulation Results

In table III we see the capacity credit earned by the coupled system, according to the methodology described in section III-A. We are simulating a 30 MW wind generator which is serving varying amounts of deferrable demand. Rather than specifying the total energy demand, we specify slack,

$$S = \frac{C \cdot N}{R}, \quad (8)$$

which is a metric of how flexible we are in terms of postponing charge. Slack is the ratio of the total energy that would be

TABLE III
CAPACITY CREDIT FOR VARIOUS LEVELS OF SLACK.

Slack	2.4	1.8	1.44	1.2	1.03
Capacity served (MW)	12.40	16.56	20.73	24.90	29.06
Capacity required (MW)	0.91	2.43	3.22	4.18	9.88
Capacity credit (MW)	11.49	14.13	17.51	20.72	19.18

TABLE IV
VALUE OF COUPLING (ALL FIGURES IN \$/DAY).

Slack	2.4	1.8	1.44	1.2	1.03
Energy rev (B)	12,404	12,404	12,404	12,404	12,404
Energy cost (B)	11,435	15,578	20,054	24,667	31,358
Capacity credit (B)	10,975	10,975	10,975	10,975	10,975
Energy cost (C)	4,114	6,811	9,519	13,479	19,195
Capacity credit (C)	14,007	17,234	21,375	25,264	23,390
Coupling value	-2,146	2,884	8,082	12,948	12,147

absorbed by loads if they were to consume at nominal capacity, over the energy demand that these loads have specified within the horizon. The closer slack is to one, the less flexibility in postponing charging, and the opposite is true as slack increases.

The first line of the table refers to the average capacity of the deferrable loads which are served and therefore increases in proportion to R , the second line is the amount of capacity utilized from the spot market during periods of peak demand - hence the capacity derating - and the third line is the difference of the two. We observe that for $S = 1.2$ the amount of capacity credit peaks. As deferrable demand increases, the amount of capacity required from the spot market exceeds the average amount of deferrable demand which is served.

In table IV we calculate the value of coupling versus pooling resources in the market. The results are average values from running a stochastic model with 5 wind states, 5 price states and 3 actions. We have run 1,000 simulations for each level of slack.

The first three line of the table refer to the baseline (B) scenario where we are pooling all resources in the market. The first line is the revenue earned by wind generators which supply all of their energy to the market and earn market prices. This value does not depend on the demand of deferrable loads, and therefore remains constant. The second line is the procurement cost of flexible loads, which fulfill their demand by buying from the spot market during periods of low market prices. We have derived these figures by solving a stochastic dynamic program for the loads, which is similar to the problem formulated in section II-A, with the difference that the cost per period is $g(x, u) = \lambda u \Delta t$, i.e. the cost of procuring energy from the spot market. The third line is the daily dollar amount of capacity income earned by wind generators for 30% capacity credit, at 1,440\$/MW. The net value of operating wind and load resources independently is equal to the sum of the first and third line net the second line.

The fourth and fifth line of table IV refer to the case of coupled (C) resources. The fourth line is equal to the solution of the dynamic program of section II-A, and the fifth line is equal to the capacity credit which is earned for the amount of

capacity calculated in the third line of table III. The net value of operating the coupled system is calculated by subtracting the fourth line from the fifth line. Finally, the value of coupling which is presented in the last line of table IV is calculated by subtracting the net value of coupling from the net value of the decoupled system.

We note that the value of coupling peaks at $S = 1.2$, which is driven by the fact that capacity credit peaks at the same level of slack. For the baseline scenario, energy revenues from selling wind power are offset beyond $S = 1.8$ by the expenditures for charging deferrable loads. In the case of coupling, energy revenues are lost when wind is discarded, compared to the decoupled case where all wind is supplied to the market. Finally, we observe that for $S = 2.4$ the value of coupling is negative, which implies that the system is better off by leaving the resources decoupled. This is due to the fact that coupling results in discarding excessive amounts of wind power, since deferrable demand is relatively low, and the resulting capacity gains are not sufficiently large to offset the opportunity cost of directly supplying energy to the market.

D. Discussion

In this section we point out various assumptions regarding our analysis and comment on these assumptions. We also describe future directions of research on this topic.

Firstly, it is questionable whether capacity credit is a sufficient metric to identify the gains of coupling. Coupling renewable energy with deferrable loads not only makes renewable capacity appear more reliable but also mitigates a variety of other problems resulting from renewable generation profiles which were described in the introduction, such as ramping and load following requirements. Coupling addresses these problems too, but this benefit cannot be captured by capacity credit. Nevertheless, capacity credit is used to reward generators for their capacity contributions in various power markets, and therefore seems like an appropriate metric to be used for our analysis. Using other methods for specifying the gains of smoothing renewable power output profiles, such as deviation penalties, is possible, but involves a series of additional assumptions (e.g. the specific rules pertaining to deviation penalties) and possibly constrains the analysis to the cases of specific markets, thereby weakening the generality of the analysis.

Another strong assumption of our analysis is that we are rewarding the coupled system the average capacity of the deferrable loads. This tends to overestimate the value of coupling, since deferrable loads which respond to spot market prices tend to consume off peak, rather than spread their consumption uniformly throughout their consumption horizon. Hence, in our analysis in section III-B we are capturing not only the capacity value of coupling but also the capacity value of load flexibility. This problem could be circumvented by derating the capacity credit for coupling operations with deferrable loads, or by estimating the capacity requirements of deferrable loads through simulation.

In future work we are interested to determine the sensitivity of our results on the degree of correlation between electricity prices and wind power availability. In particular, in systems where wind power is negatively correlated to spot market prices (which is the case in California), coupling adds value since deferrable loads absorb power which would otherwise not be needed, and this can be accurately captured by our model. The opposite holds true for the opposite case where renewable generation is positively correlated with spot market prices. In order to develop a computationally tractable model, we have been working with recombinant lattices of market prices and wind power supply, as in [17]. In particular, we have developed a recombinant model for simulating geometric Brownian motion which has been calibrated to the data used in this study, however geometric Brownian motion appears not to accurately capture the evolution of the processes in our data set. In future work we will develop a mean-reverting model such as the one described in [17] which we hope will produce more accurate results.

IV. CONCLUSION

In this paper we have described a stochastic dynamic program for optimally supplying energy to deferrable loads in the presence of a spot market. We have derived properties of the value function which have guided us in the development of an approximate dynamic programming algorithm for solving the problem. We have compared the approximate dynamic programming algorithm to a backward dynamic programming algorithm and we have found that the performance of the approximate algorithm is near optimal for when state and action spaces are aggregated sufficiently. We have validated that aggregation does not result in significant performance losses, and in future work we wish to scale the approximate algorithm in more complex versions of the problem which cannot be addressed by backward dynamic programming.

We have determined the economic value of coupling renewable generation to flexible demand by comparing the performance of the coupled system to the alternative of pooling all resources in the market and using prices as a signal for coordinating renewable generation and flexible demand. We have estimated the value of direct coupling by estimating the capacity credit and energy market expenditures of the coupled system and comparing them to the case where all resources are pooled in the market. We have determined the quantity of deferrable demand which matches optimally to a given level of renewable generation capacity by maximizing the value of coupling. In future work we wish to explore the sensitivity of our results on various assumptions of our study such as the method for determining capacity credit, as well as certain model parameters such as the degree of correlation between market prices and wind power supply.

REFERENCES

[1] R. Gross, P. Heptonstall, D. Anderson, T. Green, M. Leach, and J. Skea, "The cost and impacts of intermittency: An assesment on the evidence on the costs and impacts of intermittent generation on the british electricity network," UK Energy Research Center, Tech. Rep., March 2006.

[2] C. Loutan and D. Hawkins, "Integration of renewable resources. transmission and operating issues and recommendations for integrating renewable resources on the california iso-controlled grid," California Independent System Operator, Tech. Rep., November 2007.

[3] T. Ackermann, *Wind Power in Power Systems*. John Wiley and Sons, 2005.

[4] E. Hirst and J. Hild, "Integrating large amounts of wind energy with a small electric system," Tech. Rep., April 2004.

[5] R. Zavadil, "2006 minnesota wind integration study, volume i," Enernex Corporation, Tech. Rep., November 30 2006.

[6] F. van Hulle, "Large scale integration of wind energy in the european power supply: Analysis, recommendations and issues," European Wind Energy Association, Tech. Rep., 2005.

[7] R. Sioshansi and W. Short, "Evaluating the impacts of real time pricing on the usage of wind power generation," in *The Economics of Energy Markets*, Toulouse, France, June 2008.

[8] E. P. DeMeo, "Wind power in electricity markets," Utility Wind Interest Group, Tech. Rep., December 2004.

[9] A. Papavasiliou, S. Oren, M. Junca, A. Dimakis, and T. Dickhoff, "Coupling wind generators with deferrable loads," CITRIS 2008 White Paper Competition, Tech. Rep., 2008.

[10] D. Bertsekas, *Dynamic Programming and Optimal Control*. Athena Scientific, 2000.

[11] S. Boyd and L. Vanderberghe, *Convex Optimization*. Cambridge University Press, 2004.

[12] D. Bertsekas and J. N. Tsitsiklis, *Neuro-Dynamic Programming*. Athena Scientific, 1996.

[13] W. B. Powell, *Approximate Dynamic Programming*. Hoboken, New Jersey: Wiley Series in Probability and Statistics, 2007.

[14] G. N. Bathurst, "Trading wind generation in short term energy markets," *IEEE Transactions on Power Systems*, vol. 17, no. 3, August 2002.

[15] J. Matevosyan and L. Soder, "Minimization of imbalance cost trading wind power on the short-term power market," *IEEE Transactions on Power Systems*, vol. 21, no. 3, August 2006.

[16] P. Pinson, C. Chevallier, and G. N. Kariniotakis, "Trading wind generation from short-term probabilistic forecasts of wind power," *IEEE Transactions on Power Systems*, vol. 22, no. 3, August 2007.

[17] S.-J. Deng and S. S. Oren, "Incorporating operational characteristics and start-up costs in option-based valuation of power capacity," *Probability in the Engineering and Informational Sciences*, vol. 17, pp. 155–181, 2003.