Optimal Planning of Slow-Ramping Power Production in Energy Systems with Renewables Forecasts and Limited Storage

Nathaniel Richmond[†§]   Peter Jacko[*§]   Armand M. Makowski[‡§]

[†]Applied Mathematical & Computational Sciences, University of Iowa, USA, Email: nathaniel-richmond@uiowa.edu
[*]Lancaster University Management School, Lancaster, LA1 4YX, UK, Email: p.jacko@lancaster.ac.uk
[‡]Department of Electrical & Computer Engineering and Institute for Systems Research,
University of Maryland, College Park, USA, Email: armand@isr.umd.edu
[§]BCAM — Basque Center for Applied Mathematics, 48009 Bilbao (Basque Country), Spain

Abstract—We address the cost-efficient operation of an energy production system under renewables uncertainty. We develop an MDP model for an idealized system with the following features: (1) perfectly predictable power demand, (2) a renewable power source subject to uncertain forecast, (3) limited energy storage, (4) an unlimited fast-ramping power source, and (5) a slow-ramping power source which requires (optimal) planning. A finite-horizon stochastic optimization problem is introduced to minimize the overall cost of operating the system, and then solved numerically using standard approaches (based on backward induction) and available data. In contrast with the unit commitment problem which is traditionally optimized for a single planning frame, we show in simple scenarios that it may be beneficial to optimize over a few planning frames, and that there is no benefit to considering longer (e.g., infinite) horizons. We discretize the state space in an attempt to mitigate the curse of dimensionality usually associated with numerically solving MDPs. We note that few discretization states already yield a significant decrease in the total cost.

I. Introduction

There is considerable pressure by governments to increase the proportion of renewables (such as wind and solar) in the energy mix. In the EU the well-known "20-20-20 targets" require that 20% of energy consumption come from renewables by the year 2020 [10]. In December 2013 President Obama issued a presidential memorandum directing that the United States also meet this 20% renewable penetration target by 2020 [9]. These policies are driven by the fact that renewable energy is environmentally friendly and sustainable. However, most commonly available renewable energy sources are inherently variable and unpredictable. As a result, the benefits of renewable energy come with a risk of decreased grid reliability induced by the uncertainty in predicting renewable power production. To date current grid operation and control techniques do not adequately deal with this uncertainty, leading to the sporadic and undesirable behavior in power markets.

The efficient and reliable integration of renewable resources into the power grid creates substantial operational challenges. Several authors have used the theory of Markov decision processes (MDPs) as an optimization framework to address the cost-efficient operation of energy systems under renewables uncertainty. For instance Bejan et al. [2] discuss a simple
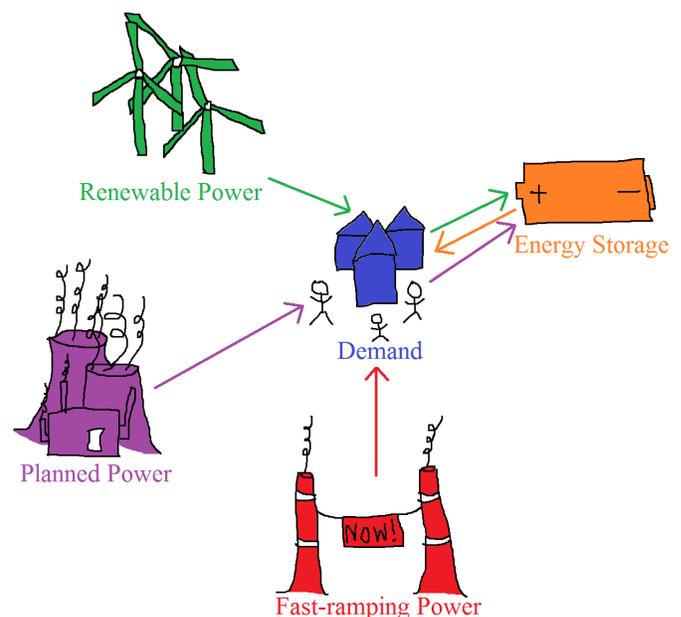

Fig. 1: Our basic energy system

stochastic model to maintain energy storage at a certain percentage of maximum capacity, usually 50%. At every time slot, storage level predictions are used to determine the slow-ramping power production level for the next time slot. Gast et al. [5] consider an MDP model whose state is the storage level, and a simple random variable is used to capture renewable unpredictability. Planning decisions are made at every slot but actualized with a delay. An expected time-average cost is then minimized over an infinite time horizon.

Building upon these earlier contributions, in Section II we propose a novel MDP model which incorporates more realistic features of the power grid and of its operations. We consider an idealized system with (1) perfectly predictable power demand, (2) a renewable power source subject to uncertain forecast, (3) limited energy storage, (4) an unlimited fast-ramping power

source, and (5) a slow-ramping power source which requires (optimal) planning. A finite-horizon stochastic optimization problem is then introduced in Section III as a way to minimize the overall operational cost of such a system. In the MDP formulation used here, the state comprises the current storage level and wind predictions for future slots; the planning is done simultaneously for several slots (which together constitute a *frame*) at the so-called *declaration time* (which marks the beginning of the frame). In Section IV the model is solved numerically using standard approaches (based on backward induction) and available data. To mitigate the usual curse of dimensionality, we propose a state discretization technique. We show that discretizing with few states can already lead to significant operational improvements, while remaining computationally feasible with standard computational resources. In Section V we see through several test cases that it is very beneficial to implement the solution obtained for horizons of 5 to 10 frames, and that no benefit accrues from considering a longer (e.g., infinite) horizon – This should be contrasted with the unit commitment problem which is traditionally formulated over a single planning frame. Section VI summarizes the main conclusions of this work.

II. THE MODEL

A. The Energy System Components

Throughout this paper we work in a discrete time setting: Time is organized in contiguous time intervals or *slots* of equal duration. In practice a time slot would typically have a duration of half an hour. However, without loss of generality, we normalize the common length of these time slots to unity. We identify a time slot with its left boundary, referring to time slot $[t, t+1)$ simply as the slot t for $t \in \{0, 1, 2, \ldots, T\}$ where T is the planning *horizon*. At time t, we have access to historical information about wind power predictions and wind power production from past slots $\theta \in \{-H, -(H-1), \ldots, t-2, t-1\}$. The system is observed at the beginning of every slot and planning is done only at the beginning of specific slots, called *declaration times* (which we define later). We wish to optimize the operation of an energy system with the following five components (see Figure 1):

1) Power Demand: The demand $d(t) \geq 0$ is the amount of power required by the electricity users from the energy system during time slot t. Although this demand becomes known exactly only at the moment it must be served, in all practicality we may assume that it is perfectly predictable. This is not an unreasonable assumption; energy companies are known to have extremely accurate estimates that even take television schedules into account down to the second.[1]

2) Renewable Power Source: To keep the terminology simple, we assume wind to be the only renewable power source. Let $W(t) \geq 0$ denote the renewable power produced from wind turbines during time slot t. For $t \geq 0$ we note

that $W(t)$ is a random variable, reflecting the fact that we cannot predict it exactly. Available to the planner at time t are the realizations of wind power produced in the past, namely $W(\tau)$ for $-H \leq \tau < t$, and the predictions $\widehat{w}(t+i|\tau)$ made at time τ for wind power produced at time $t+i$ for $-H \leq \tau \leq t$ and $i \geq 0$.

3) Energy Storage: Storage facilities, in the form of dams or batteries, are common in energy systems. They mitigate, or buffer, the impact of the volatility of renewable energy on the system as a whole. In our energy system model, let $B(t)$ denote the energy contained in our storage devices at the beginning of time slot t. Every storage device has limited capacity as well as physical constraints on its charge/discharge rates: In addition to the constraint $0 \leq B(t) \leq B_{\max}$, it is also assumed that the device can charge and discharge a maximum of $C_{\max}$ and $D_{\max}$ units of energy per time slot, respectively. Furthermore, $0 < \eta \leq 1$ is the efficiency of the storage facility. A typical value given current technology is $\eta = 0.8$. Energy storage helps to avoid complete waste of residual energy, but still leads to two types of energy loss: the energy that cannot be charged because of the charging constraint or capacity constraint, and the charged energy that is inefficiently stored.

4) Fast-Ramping Power Source: The short-notice power generated $G(t) \geq 0$ during time slot t is also known as fast-ramping power. Unlike the slow-ramping power production discussed below, this power requires virtually no advance planning and may be used as an instant panacea when power shortages are realized. However, this power tends to be very expensive and should thus be used sparingly. In practice, short-notice power may take the form of spinning reserves or Open Cycle Gas Turbines (OCGTs).

5) Slow-Ramping Power Source: The planned power production $p(t|\tau) \geq 0$ is the amount of power produced during time slot t, as planned at some earlier time $0 \leq \tau \leq t$. It may also be called slow-ramping power production or non-dispatchable power production, and must be planned before the status of other variables (such as renewable power at time t) becomes known. In practice it may, for instance, include power produced from nuclear or coal plants. The slow-ramping power production is the single control variable of the energy system operator.

B. The System Dynamics

Next we introduce some definitions and notation concerning system dynamics and operations that will be needed for model design.

1) Time Structure: We refer to the given positive integer f as the *frame length*, so that a single time frame consists of f time slots. The beginning of a frame is called a *declaration time* and is the epoch at which the planner announces the slow-ramping power production level for that frame. For notational convenience, define the set of all declaration times as $DT := \{nf : n = 0, 1, \ldots, \lfloor T/f \rfloor\}$ where nf is called the n-th declaration time. In our model we choose T to be

---

[1]We use a lowercase d to represent demand because of our assumption that demand is known for any time t; it is not random nor does it depend on the policy. We use an uppercase letter to represent any random variable.

a multiple of f. The n-th frame is the frame corresponding to the n-th declaration time and comprises the slots $t = nf, nf + 1, ..., (n+1)f - 1$. Thus, at the n-th declaration time nf the planner announces how much slow-ramping power $p(t|nf)$ needs to be produced during the n-th frame. From now on, we use $t \in \{0, 1, 2, ..., T\}$ to refer to an arbitrary time, and $s \in \mathrm{DT}$ to refer to an arbitrary declaration time.

In many situations, decisions made for production levels are only actualized in the future. Such a delay can however be incorporated in a straightforward manner, e.g., see [5] where production is planned a fixed number of slots ahead. We also ignore plant start-up times and minimum run times; they could also be included easily as additional model constraints. Here these issues will not be discussed any further for the sake of modeling simplicity.

2) Storage Dynamics: Recall that $B(t)$ is a random variable representing the storage level (in energy) at time t. We have constraints

$$0 \le B(t) \le B_{\max}$$

and

$$-D_{\max} \le B(t+1) - B(t) \le \eta C_{\max}.$$

The constants $B_{\max}$, $D_{\max}$, $C_{\max}$, and η represent the maximum storage capacity, the ramping constraints for discharging and charging the storage facility, and the efficiency of the storage facility (assumed active only when charging the storage device), respectively.

The power *mismatch* $\Delta(t)$ at time t is defined as the difference

$$\Delta(t) = d(t) - W(t) - p(t|s),$$

when $t \in [nf, (n+1)f)$ and $s = nf$ is the associated declaration time for some $n \in \{0, 1, \ldots, \lfloor T/f \rfloor\}$. This mismatch represents the additional power required to satisfy the demand $d(t)$ after taking $p(t|s)$ and $W(t)$ into account. Note that $\Delta(t) > 0$ means that more power must be provided through discharging storage or fast-ramping power production in order to meet the demand, while $\Delta(t) < 0$ indicates that there is an excess of power, which should be stored whenever possible.

Formally, the dynamics of storage levels are of the form

$$B(t+1) = \Phi(B(t), \Delta(t))$$

with mapping $\Phi : \mathbb{R}_+ \times \mathbb{R} \to \mathbb{R}_+$ given by

$$\Phi(B, \Delta)$$
$$= \begin{cases} \max\{0, B - \min\{\Delta, D_{\max}\}\} & \Delta \ge 0 \\ \min\{B_{\max}, B + \eta \min\{C_{\max}, -\Delta\}\} & \Delta \le 0. \end{cases}$$

3) Cost Dynamics: As we will see shortly, we are interested in minimizing an expected total cost over the planning horizon. At every time slot t, the one-slot cost $C(t)$ is constructed as a weighted average of the fast-ramping power production $G(t)$ and lost energy $L(t)$ to be defined next.

The fast-ramping power production source will only generate the amount needed to satisfy demand exactly, after using the slow-ramping power, wind power, and the energy in storage. Thus, $G(t)$ is given by

$$G(t) = \begin{cases} \Delta(t) - (B(t) - B(t+1)) & \Delta(t) \ge 0 \\ 0 & \Delta(t) \le 0. \end{cases}$$

Let $L(t) \ge 0$ denote the amount of energy lost at time t by storage facility inefficiencies or energy that is spilled because it cannot be stored. As these are the only two ways energy can be lost, this loss is given by the amount

$$L(t) = \begin{cases} 0 & \Delta(t) \ge 0 \\ -\Delta(t) - (B(t+1) - B(t)) & \Delta(t) \le 0. \end{cases}$$

Using the definition of $B(t+1)$ in $G(t)$ and $L(t)$ we get

$$G(t) = \begin{cases} \Delta(t) - \min\{B(t), \min\{\Delta(t), D_{\max}\}\} & \Delta(t) \ge 0 \\ 0 & \Delta(t) \le 0 \end{cases}$$

and

$$L(t) = \begin{cases} 0 & \Delta(t) \ge 0 \\ -\Delta(t) - \min\{B_{\max} - B(t), \ldots\} & \Delta(t) \le 0 \end{cases}$$

with

$$\ldots = \eta \min\{C_{\max}, -\Delta(t)\}.$$

Finally, for some weight $0 < \lambda < 1$, we introduce the cost $C(t)$ at time t as the quantity given by the convex combination

$$C(t) = \lambda G(t) + (1 - \lambda) L(t)$$

so that

$$C(t) = \begin{cases} \lambda\Big(\Delta(t) - \big(B(t) - B(t+1)\big)\Big) & \Delta(t) \ge 0 \\ (1-\lambda)\Big(-\Delta(t) - \big(B(t+1) - B(t)\big)\Big) & \Delta(t) \le 0 \end{cases}$$

We may write $\lambda = c_G/(c_L + c_G)$, with c_G and c_L as the unit costs of producing fast-ramping power and slow-ramping power, respectively.

4) Forecasts and Controls: At declaration time $s \in \mathrm{DT}$, the planning decision will draw on all knowledge of the current state of the system at time s, namely

$$I_W(s), \{B(t) : t = 0, \ldots, s\}, \{d(t) : t = s, \ldots, T\}$$

where $I_W(s)$ is the information available at time s from an independent weather base about past wind, past wind predictions, and current wind predictions. At declaration time s, we use this information to plan the slow-ramping power production that will occur in the next time frame. In particular, we use the current information to make predictions about the future state of the system.

The prediction $\widehat{w}(t + i|t)$ at time t of wind power produced at time $t + i \ge t$ uses known information about past wind predictions and actual wind power levels before time t. For instance we may use persistent forecast, in which case $\widehat{w}(t+i|t) = W(t-1)$ for all $i \ge 0$. We may also use weather prediction with or without a linear correction, or any other

prediction method. Here the prediction is based on available weather forecast data which is read in from a file, with a linear correction.

At time t, we predict that the storage level $B(t+i)$ will be $\widehat{b}(t+i|t)$, given by

$$\widehat{b}(t+i|t) = \begin{cases} \Phi\left(\widehat{b}(t+i-1|t), \widehat{\Delta}(t+i-1|t)\right) & i > 0 \\ B(t) & i = 0. \end{cases}$$

Recalling that

$$\Delta(t+i) = d(t+i) - W(t+i) - p(t+i|s),$$

by analogy at time t we predict the mismatch at time $t+i$ to be given by

$$\widehat{\Delta}(t+i|t) = d(t+i) - \widehat{w}(t+i|t) - p(t+i|s), \quad i \geq 0$$

These predictions are useful in determining the planned power production level, which we are now in a position to define. The planned power production at declaration time s for power at time $s+i$, for $0 \leq i < f$, is defined as

$$p(s+i|s) = d(s+i) - \widehat{w}(s+i|s) + u(s+i|s)$$

for some production *offset* $u(s+i|s)$ which is our control. Although we follow [2], [5] in taking the production offset as the control variable, the planned production $p(s+i|s)$ is certainly another possible and valid control variable.

III. THE MDP FORMULATION

We model our problem as a Markov decision process. Using the definitions and notation given earlier, we define the states, actions, and transition probabilities as follows:

The *state* at declaration time s is

$$\mathbf{x}(s) := (b(s), \widehat{w}(s|s), \widehat{w}(s+1|s), ..., \widehat{w}(s+f-1|s)),$$

namely the storage level at the declaration time and the wind predictions for each slot in the corresponding frame;

The *action* at declaration time s is

$$\mathbf{u}(s) := (u(s|s), u(s+1|s), ..., u(s+f-1|s)),$$

namely a sequence of production offsets for each slot in the frame which starts at the declaration time;

The *transition probability matrix* $\mathbf{P}$ is generated by iterating through historical wind data and computing the probability of transitioning from a specific state to another specific state, under a specific action; this is done for every possible state-state-action triple.

By standard results from the theory of MDPs [7] it suffices to consider state-dependent policies, i.e., policies π which assign an action to a given state. Such a policy will be implemented at every declaration time, when it determines how much slow-ramping power needs to be produced for the time slot(s) in the corresponding frame, based on our estimations of the storage level under the resulting actions for those time slots.

Recall the definition of the one-slot cost $C(t)$ from the previous section. Our objective at every planning time is to

Data: $s_B, s_W, s_A, \boldsymbol{P}, T$
for *each state* x **do**
  | $V_{x,0}^* = 0$
end
for *each frame* $n = 1$ *to* $\lfloor T/f \rfloor$ **do**
  for *each state* x **do**
    for *each action* u **do**
      | $Q_n^u(x) = \mathbb{E}_x^u[C(x)]$
      | $\quad + \Sigma_y P^u(x,y) \cdot V_{n-1}(y)$
    end
    $u_n^*(x) = \arg\min_u Q_n^u(x)$
    $V_{x,n}^* = Q_n^{u_n^*(x)}(x)$
  end
end
return π_T^*, V_T^*

Algorithm 1: An algorithmic scheme of backward induction.

minimize the expected total cost over a finite time horizon. Our problem at planning time $t = 0$ if initially in state $z = \mathbf{x}(0)$ is to solve for

$$V_{z,T}^* := \min_\pi \left\{ \mathbb{E}_z^\pi \left[\sum_{t=0}^T C(t) \right] \right\}$$

where $\mathbb{E}_z^\pi$ denotes the expectation under policy π if initially in state z.

IV. THE SOLUTION

We numerically solve the MDP problem. For practical purposes, we discretize the storage states, wind predictions, and actions. Within a single time slot, let s_B denote the number of possible discrete storage states, s_W the number of possible discrete wind predictions, and s_A the number of possible discrete actions. We then use historical wind production and prediction data to generate the transition probabilities. These are the probabilities of transitioning from a specific state to a different state given a certain action. This results in a matrix matrix $\mathbf{P}$ of size

$$(s_B)^2 \times (s_W)^{2f} \times (s_A)^f$$

where f is the frame length.

Given this transition probability matrix, we generate expected costs associated with performing a specific action given a specific state. We use this expected cost matrix to perform backward induction (see algorithm 1) over the horizon T, producing the optimal policy π_T^* as the set of optimal actions $\mathbf{u}_T^*(\mathbf{x})$ for each state $\mathbf{x}$, and the optimal values $V_{\mathbf{x},T}^*$.

A. Memory Limitations

MDP problems are well-known to suffer from the *curse of dimensionality*. This model becomes intractable for large values of f, s_B, s_W, and s_A due to memory limitations. A computer with 16GB of RAM, for example, can only solve the model for $f = s_W = s_B = s_A = 4$. Because of the exponential growth of the memory requirements, even using one of today's supercomputers would only allow the exact
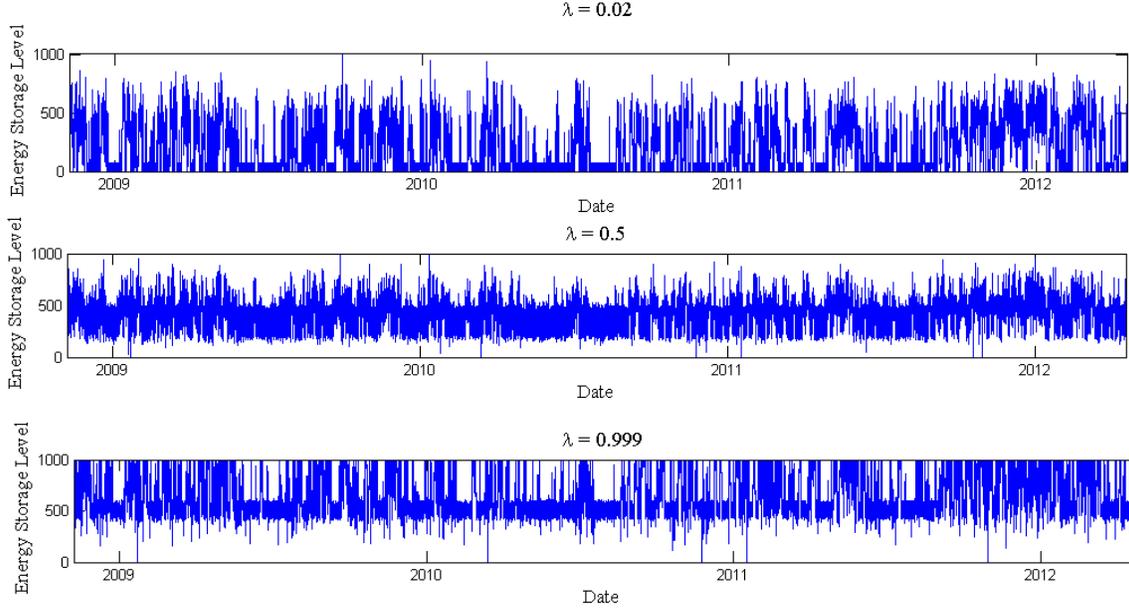
Fig. 2: Stored energy over time for different values of λ

solution of our model for $f = s_W = s_B = s_A = 5$. However, in the next section we show that the precision of the solution obtained from our model might be acceptable for practical purposes.

There exist various methods that address the curse of dimensionality by reducing the "dimension" of the problem and providing an approximate solution to it. For instance, in generating the transition probability matrix from the data, the dimension of the matrix $\mathbf{P}$ can be reduced by discarding any wind forecasts that occur very infrequently.[2] A more in-depth exploration of these methods is beyond the scope of this paper.

V. SIMULATIONS AND RESULTS

We ran simulations to determine the efficacy of the optimal solution obtained by backward induction as described in the previous section. All simulations use UK half-hourly wind data for the period June 2009 - April 2012 available from *elexonportal.co.uk*, and were run on a computer with an Intel Quad Core i7-2600 CPU at 3.40GHz and 16GB of RAM.

Given our general model, the optimal policy is found by using backward induction. We used parameters $w_{\max} = B_{\max} = 1000$, $C_{\max} = D_{\max} = 200$, $a_{\max} = 100$, $\eta = 1$, $f = 2$, and $s_W = s_B = s_A = 4$ to test the runtime and convergence of backward induction over different horizons. The results are shown in Table I, with the costs rounded to the nearest integer. When the number of frames (horizon) is small, there is considerable benefit to increasing the number of frames. However, the impact of increasing the frames to more than 5 is negligible. Also note that the backward

---

[2] In our computations, we discard those wind forecasts that occur fewer than $\frac{0.01 \cdot \lfloor H/f \rfloor}{(s_W)^f}$ times in our historical weather data. Note that $\lfloor H/f \rfloor$ represents the number of historical weather predictions (for an entire frame) to which we have access.

induction solution time is approximately equal to the number of frames we consider. In view of these observations, we chose to run all simulations using backward induction over 10 frames; in other words we have taken $T = 10f$.

Figure 2 illustrates the energy level at the storage device as a function of time for three different cost ratios. In these simulations, we used the same parameters as for Table I. In the upper plot, observe that when fast-ramping power is cheap ($\lambda = 0.02$), the optimal policy maintains a low storage level. This makes sense because the amount of energy lost due to storage device being close to capacity B_{Max} is minimized. In the middle plot, when $\lambda = 0.5$, we see that the storage device is kept around middle capacity, as we would expect. This serves to balance lost energy and unsatisfied demand. Lastly, when $\lambda = 0.999$ and fast-ramping power is expensive, our policy maintains the stored energy relatively near maximum storage capacity. This minimizes the amount of unsatisfied demand that must be met by fast-ramping power.

In Table II, we examine the considerable impact of different values of the frame length f. The same parameters as before were used, and we have rounded all table entries to the nearest integer. Recall that f is the number of time slots for which one must declare power production in advance, when planning at a declaration time. In the scenarios considered here, we have taken $s_W = s_B = s_A = 3, 4, 5, 6$ as indicated in the leftmost column of the figure. The "Time" column refers to the time needed to generate the probability transition matrix from the June 2009 - April 2012 data; this is the most computationally expensive part of the model. The "Total Cost" column contains the total cost incurred by the simulation of the respective optimal solution. For every number of states, we see a positive correlation between f and total cost. The further ahead energy companies must plan, the less effective their plans become due

Frames	Total Cost	Solution Time (s.)
1	6981	0.99
2	5092	1.97
3	4750	2.95
5	4736	4.96
10	4736	9.83
20	4736	19.63
100	4736	100.25
1000	4736	1009.99

TABLE I. Backward induction over varying numbers of frames (horizon).

	f	Time (s.)	Total Cost
3 states	1	125	4557
	2	300	5217
	3	788	9414
4 states	1	290	3844
	2	941	4736
	3	3282	8462
5 states	1	578	3274
	2	2329	4552
	3	—	—
6 states	1	993	3226
	2	4865	4421
	3	—	—

TABLE II. The impact of f and the number of states on the total cost.

to worse predictions, as expected. More interestingly, we note that for fixed f, an increase in the number of states decreases the total cost. Increased granularity is more computationally expensive; the blank entries in the table could not be completed for lack of RAM.

VI. Conclusion and Future Work

In this paper, we examined optimal power production planning in an idealized energy system in which slow-ramping power must be planned ahead periodically at certain "declaration times" throughout the day. A new, general Markov decision process model was formulated for this system, and we used backward induction to find the optimal policy. This approach was tested on historical UK electricity grid data using different parameters, and great reductions in total cost were achieved.

This model is more realistic than existing models at the expense of being more memory expensive. With current technological standards, the model can only be tractably implemented for small frame lengths and coarse discretizations of the states and actions. Nevertheless, even the computationally feasible implementation of our model was proven effective in optimizing the system. This indicates possible benefits for energy markets to reduce the frame length f over which energy companies must project their production levels (see Table II). For instance, instead of the common practice of a single declaration for the whole next day split into half-hourly slots ($f = 48$), several intra-day declarations would allow for a more cost-efficient operation.

More work may be done to reduce the runtime and computational requirements of this model. Known approximation methods may be implemented to mitigate the curse of dimensionality. Sparse matrix techniques may also be applied to our model, further decreasing the computational requirements. The approach presented in this paper constitutes a promising initial step for improving the cost-efficient operation of energy systems under renewables uncertainty; it may also be extended to large electricity grid networks.

References

[1] M. Arnold and G. Andersson. *Model predictive control of energy storage including uncertain forecasts*. Proceedings of the Power Systems Computation Conference (PSCC), Stockholm (Sweden). 2011.

[2] A. Iu. Bejan, R.J. Gibbens, and F.P. Kelly. *Statistical aspects of storage systems modelling in energy networks*. Proceedings of the 46th Annual Conference on Information Sciences and Systems (CISS 2012), Princeton University (NJ), March 2012.

[3] D.S. Callaway. *Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy*. Energy Conversion and Management, vol. 50 (2009), pp. 1389-1400.

[4] E.D. Castronuovo and J.A.P. Lopes. *Optimal operation and hydro storage sizing of a wind-hydro power plant*. International Journal of Electrical Power & Energy Systems, vol. 26 (2004), pp. 771-778.

[5] N. Gast, D.C. Tomozei, and J.Y. Le Boudec. *Optimal energy storage policies with renewable forecast uncertainties*. Proceedings of Greenmetrics 2012, Imperial College, London (UK), June 2012.

[6] V.S. Pappala, I. Erlich, K. Rohrig, and J. Dobschinski. *A stochastic model for the optimal operation of a wind-thermal power system*. IEEE Transactions on Power Systems, vol. 24 (2009), pp. 940-950.

[7] M.L. Puterman *Markov Decision Processes: Discrete Dynamic Stochastic Programming*. 1994.

[8] A. Tuohy and M. O'Malley. *Impact of pumped storage on power systems with increasing wind penetration*. Proceedings of the IEEE Power & Energy Society General Meeting, pp. 1-8. Calgary (Canada), 2009.

[9] "FACT SHEET: Presidential Memorandum on Federal Leadership on Energy Management." Retrieved December 6, 2013 from *http://www.whitehouse.gov/the-press-office/2013/12/05/fact-sheet-presidential-memorandum-federal-leadership-energy-management*.

[10] "The EU Climate and Energy Package". Retrieved December 6, 2013 from *http://ec.europa.eu/clima/policies/package/index_en.htm*.