

Meeting Inelastic Demand in Systems with Storage and Renewable Sources

Natarajan Gautam

Dept. of Industrial & Systems Engr.,
Texas A&M University, USA
Email: gautam@tamu.edu

Yunjian Xu

Engineering Systems and Design,
Singapore Univ. of Technology and Design,
Email: yunjian_xu@sutd.edu.sg

Jeremy T. Bradley

Department of Computing,
Imperial College, London, UK
Email: jb@imperial.ac.uk

Abstract—We consider a system where inelastic demand for electric power is met from three sources: the grid, in-house renewables such as wind turbines or solar panels, and an in-house energy storage device. In our setting, power demand, renewable power supply, and cost for grid power are all time-varying and stochastic. Further, there are limits and efficiency issues for charging and discharging the energy storage device. Under such a scenario, at all times across an infinite horizon, we need to determine how to split power demand among the various sources. For that, we formulate an optimization framework that minimizes the long-run average operational cost subject to satisfying demand and meeting capacity constraints. To determine the optimal actions, we construct a dynamic program with parameters estimated from a training data set (that uses real demand, supply and cost data). As the dynamic program is computationally intensive for large-scale problems, we explore an algorithm based on approximate dynamic programming, and apply it to a test data set. We compare its performance against other heuristics that may or may not use the training data. In addition, we ascertain the value of storage as well as the value of installing a renewable source.

I. INTRODUCTION

In this research, we consider a *consumer* of electricity with inelastic power demand, i.e. demand must be satisfied instantaneously and cannot be either postponed or cut back. *Demand* for power is time-varying and stochastic. Part of the demand can be met by a renewable energy source (such as photovoltaic (PV) solar panels or wind turbines) that is situated locally and owned by the consumer. Note that renewable power *supply* is time-varying and stochastic. Remaining demand (if any), beyond what the renewable source can supply, is satisfied either from the *grid* or by an in-house *Energy Storage Device* (ESD) or both. It is also possible to charge the ESD, thus the ESD is both a source and a sink of power. Like power demand and renewable supply, *price* for power from the grid is also time-varying and stochastic.

In the last few years this problem has received a lot of attention in the literature. The literature can be broadly divided into two categories: (i) articles that assume that all parameters are deterministic time-varying quantities and make strategic decisions such as whether to invest in renewable sources, and/or ESDs; what fraction of demand is met by renewable supply; whether it is possible to go net zero, etc., and (ii) articles that assume that demand, supply and costs are either IID or follow stationary Markov chains; these result in Markov

decision processes (MDPs) or Lyapunov-based techniques to model the storage process.

In particular, [1] considers case (i) with a data center as the consumer. Besides the modeling nuances we consider, [1] also models the use of diesel generators and external renewable sources (besides on-site renewables). They evaluate strategic decisions such as whether to use renewables, from where, whether to invest in storage, whether to use diesel generators, etc. Their key finding is that the results are dependent on the renewable source's capacity factor.

For case (ii), there have been recent studies on the operation of consumer-owned ESDs. Threshold structures are established for the optimal control policy in a variety of MDP settings without renewable generation [2], [3]. Leveraging on techniques from queuing theory, a few recent papers propose a variety of on-line algorithms that are shown to be asymptotically optimal, as the storage capacity increases to infinity [4], [5]. Closer to the present paper, the authors of [5] consider a similar storage operation problem faced by a consumer who has additional options to sell power back to the grid (but through the ESD) and to shift demand across time. Unlike MDP, which is computationally complex and requires substantial statistical information of the system dynamics, their method requires no statistical knowledge and uses an extremely light linear program. The algorithm is expected to perform well when the storage capacity is significantly larger than the maximum charging/discharging rates.

While our article considers a single consumer, [6] takes the perspective of a renewable source (such as a wind farm) with large storage device. They consider renewable supply to be controlled and sold/stored, and also negative prices. Their objective is to develop a fast algorithm to determine how much to generate and what fraction of that must be sold versus stored so that revenue is maximized. At each discrete time instant, they develop three thresholds on the ESD contents to determine when to: generate, buy and store; generate and store; do nothing; or sell.

It is our belief that both the deterministic variability and the stochastic variability are important to determine optimal operational actions. The contribution of this paper is twofold: (i) we construct an MDP framework that incorporates the randomness in consumer demand, renewable generation, and electricity prices; (ii) using a set of real training data on




Fig. 1. Schematic representation of scenario and notation

electricity prices, solar generation, and consumer demand, we numerically compare approaches ranging from Markovian models, to hybrid methods based on statistics, optimization and heuristics, to those that require no historical data. The insight we obtain from these numerical experiments will guide the design of more efficient heuristics. Interestingly, all the aforementioned articles in the literature consider wind energy while our focus for the numerical study is on solar PV.

The storage operation problem is intimately related to inventory control problems with random production cost and uncertain demand [7], [8]. We note, however, our model is significantly different from the setting in the inventory control literature. In our model, there is no “inventory” holding cost that is proportional to the storage level; instead, the major losses resulting from storage operation are due to energy injection and withdrawal (e.g., ESD charging and discharging). In addition there is a finite maximum rate at which the ESD can be charged and discharged.

We describe our problem in Section II. In Section III we develop a probabilistic model and suggest approaches to solve it. We consider other heuristic algorithms in Section IV and demonstrate their quality in Section V by obtaining parameters using a real training data set and testing them.

II. PROBLEM DESCRIPTION

Before describing our model, we state a few *assumptions*. Recall that demand is inelastic. There are no operational costs for the consumer other than price of electricity from grid. Price is exogenous and not affected by the consumer’s decisions. There are inefficiencies in charging and discharging the ESD. We assume no leakages in the ESD, since the storage efficiency of many different types of modern batteries (e.g., Lead acid, Li-ion and Vanadium redox batteries) is close to 100% [9]. The ESD has a finite energy storage capacity. Power cannot be sold to the grid (we will relax this assumption in an extension of this work).

Next we describe some of the *notation* used in this work (pictorially described in Fig. 1). We consider a discrete-time model where time periods are indexed by $t = 0, 1, \dots$. The *stochastic* uncontrollable variables are: D_t , the *demand* for energy in period t (in kWh); S_t , the *energy supply* from renewable source in period t (in kWh); C_t , the *cost* in period t for a unit of energy from grid (in \$ per kWh). Let U_t be the amount of energy in the ESD at the beginning of period t (in kWh). We assume that D_t for all t is smaller than the

grid capacity. However there are capacities and constraints in the ESD charging and discharging processes. A maximum of K kWh of energy can be stored in the ESD at any time. The ESD can be discharged and charged at a maximum rate of c_{dis} and c_{char} (in kW) respectively. Also, the ESD discharging and charging efficiencies are $\eta_{dis} \leq 1$ and $\eta_{char} \leq 1$ respectively (which we explain next). If ρ kW of power is used to charge the ESD in period t , then $U_{t+1} - U_t$ (when $U_{t+1} \leq K$ and $\rho \leq c_{char}$) is ρ/η_{char} hours per unit time. Likewise if ρ kW of power is needed from the ESD, then $U_{t+1} - U_t$ (when $U_{t+1} \geq 0$ and $\rho \leq c_{dis}$) is $-\rho/\eta_{dis}$ hours per unit time. Next we describe the *decision variables* under the control of the consumer. Let X_t , Y_t and Z_t be the energy drawn (in kWh) from the grid, renewable source and ESD respectively at period t . While $X_t \geq 0$ and $Y_t \geq 0$ for all t , Z_t can be positive or negative (latter denoting power used to charge ESD).

Having described the notation, we are in a position to state the problem we consider. During every period t , given demand D_t , renewable supply S_t , cost C_t and ESD charge level U_t , we need to *determine* the supply from grid X_t , the draw from renewable source Y_t , and contribution from the ESD Z_t so that the long-run average cost per time period is minimized subject to satisfying demand, staying within ESD capacities and other constraints such as dynamics and non-negativity. This sequential decision making problem can be formulated mathematically as follows

$$\begin{aligned}
 & \text{Minimize } \lim_{\tau \rightarrow \infty} \frac{1}{\tau} \sum_{t=0}^{\tau} C_t X_t \\
 & \text{Subject to the following } \forall t \in \{0, 1, 2, \dots\} \\
 & \quad X_t + Y_t + Z_t \geq D_t \\
 & \quad 0 \leq Y_t \leq S_t \\
 & \quad Z_t \leq c_{dis} \\
 & \quad -\min\{Z_t, 0\} \leq c_{char} \\
 & \quad \psi(U_{t+1} - U_t) = -\max\{Z_t, 0\}/\eta_{dis} + \eta_{char} \min\{Z_t, 0\} \\
 & \quad 0 \leq U_t \leq K \\
 & \quad X_t \geq 0
 \end{aligned}$$

where ψ is a constant for time-unit conversion, i.e. number of time units per hour (viz. when U_t is in kWh and Z_t in kW, if length of period t is 1 second then $\psi = 3600$).

Remark 1: The above formulation has a trivial solution if C_t stays constant over t . In particular, if $D_t > S_t$ for all t then $Z_t = 0$, $Y_t = S_t$ and $X_t = D_t - S_t$ is the optimal solution with no need for storage. However, if C_t stays constant over t but D_t is not always greater than S_t , then the greedy solution of charging or discharging at maximum feasible capacities when D_t is less or more respectively than S_t is optimal. Thus for the rest of this article we consider the non-trivial case of C_t varying with time and the possibility that for some time periods S_t can be greater than D_t .

Remark 2: We can reduce the above problem to a 1-dimensional control in X_t or Z_t by realizing that we can let $Y_t = S_t$ and $Z_t = D_t - X_t - Y_t$. However, for ease of presentation

we will use all variables, not just X_t or Z_t .

Remark 3: Due to the linear nature of the objective function and constraints, the value function at every period t is piecewise-linear in U_t [3]. As a result, given the system state at period t , namely D_t , S_t , C_t , and U_t , the optimal policy is such that Z_t is one of the following: (i) zero, (ii) the minimum of c_{dis} , ψU_t and $D_t - S_t$, or (iii) negative of certain threshold that depends on future demand and cost.

III. PROBABILISTIC MODEL WITH CYCLES

Analyzing the data described in [1], the NREL labs <http://www.nrel.gov/midc/> and [10] respectively, it is evident that demand, solar PV supply and cost are time-varying and stochastic. However, it is also not unreasonable to assume that there are daily or weekly seasonality effects. In other words, there is a deterministic variability as well as stochastic variability. To model such a phenomenon we consider what we call *probabilistic model with cycles*.

Definition 1: An uncontrolled process $\{V_t, t \geq 0\}$ is *cyclic* with cycle length T if V_t is stochastically identical to $V_{t+\ell N}$ for all $\ell \in \{1, 2, \dots\}$, where N is the number of periods in a cycle of length T hours, i.e. $N = \psi T$.

Based on the above definition, we assume that $\{D_t, t \geq 0\}$, $\{S_t, t \geq 0\}$ and $\{C_t, t \geq 0\}$ are *cyclic* with cycle length T (T typically is the equivalent of 1 day or 1 week). Further, we write down for all $t \in \{0, 1, \dots\}$, with $n = (t \bmod N)$,

$$\begin{aligned} D_t &= d_n Z_t^d + \delta_n \\ S_t &= s_n Z_t^s \\ C_t &= c_n Z_t^c + \gamma_n, \end{aligned}$$

where $\{d_1, d_2, \dots, d_N\}$, $\{\delta_1, \delta_2, \dots, \delta_N\}$, $\{s_1, s_2, \dots, s_N\}$, $\{c_1, c_2, \dots, c_N\}$ and $\{\gamma_1, \gamma_2, \dots, \gamma_N\}$, are sets of deterministic constants while $\{Z_t^d, t \geq 0\}$, $\{Z_t^s, t \geq 0\}$, and $\{Z_t^c, t \geq 0\}$, are stationary and independent *discrete time Markov chains* on discrete state spaces \mathcal{S}^d , \mathcal{S}^s , and \mathcal{S}^c and transition probability matrices \mathcal{P}^d , \mathcal{P}^s , and \mathcal{P}^c respectively.

One can think of $\{s_1, s_2, \dots, s_N\}$ as the power supplied by PV panels on a perfectly sunny day while \mathcal{S}^s as a discrete set of values between 0 and 1. The *demand* and *cost* terms do not have such a nice interpretation and one would have to model them carefully based on data. In Section V, we will use training data to estimate d_n , δ_n , s_n , c_n and γ_n for all $n \in \{1, \dots, N\}$ as well as \mathcal{P}^d , \mathcal{P}^s , and \mathcal{P}^c . However, for the rest of this section we take a probabilistic approach assuming all the aforementioned parameters are known and formulate the system as an MDP.

We denote the **system state** at time $t \in \{0, 1, 2, \dots\}$ as a 5-tuple

$$\mathbf{x}_t = \{\{t/N\} + 1, Z_t^d, Z_t^s, Z_t^c, U_t\}$$

(where $\{t/N\}$ denotes $(t \bmod N)$) with state space \mathcal{S} given by the cartesian product

$$\{1, 2, \dots, N\} \times \mathcal{S}^d \times \mathcal{S}^s \times \mathcal{S}^c \times \mathcal{S}^u$$

where \mathcal{S}^u is the discrete set of values between 0 and K that U_t can take. It is worth noting that the time dependency and

correlations of consumer demand and renewable generation are incorporated by including in the system state a periodic Markov chain that describes time evolution.

Let the **action** at time t denote the amount of power to be supplied from the grid, i.e. X_t , with action space \mathcal{A} corresponding to the set of all possible non-negative values that are less than or equal to

$$\psi \min\left\{\frac{K - U_{\{t/N\}}}{\eta_{char}}, c_{char}\right\} + \max\{D_{\{t/N\}} - S_{\{t/N\}}, 0\},$$

for all $t \in \{1, 2, \dots, N\}$. As noted in Remark 2, the energy drawn from renewable source and ESD at period t is determined by X_t , i.e., $Y_t = S_t$ and $Z_t = D_t - X_t - Y_t$.

For every action $X_t \in \mathcal{A}$, $\mathbf{x} \in \mathcal{S}$, and $\mathbf{y} \in \mathcal{S}$, we can compute the **transition probability** $P_{\mathbf{xy}}(X_t)$ using appropriate Kronecker products of \mathcal{P}^d , \mathcal{P}^s , \mathcal{P}^c and other matrices of zeros and ones (which are not explained due to space constraints). The **stage cost** at time t is the product of the corresponding power cost in state i times $X_t \in \mathcal{A}$, $C_t X_t$.

By incorporating the time element into the state of the MDP, we have indeed formulated an MDP with a “stationary” policy (the quotes are because the policy is stationary from one cycle of N values to the next cycle, but not within a cycle). For a given stationary policy which maps every point in the state space \mathcal{S} to a point in the action space \mathcal{A} , the long-run average cost per unit time corresponds to the objective function of our optimization problem defined in Section II (and also results in a feasible solution).

Note that the above MDP has a finite state-space and a finite action-space. There are many methods to obtain the optimal action $a \in \mathcal{A}$ in state i for all $i \in \mathcal{S}$. We consider a linear program (LP) based method described in [11]. The LP can be solved (especially by packages such as MATLAB) when the state and action spaces are small, $\eta_{char} = \eta_{dis} = 1$, and D_n and S_n belong to a small discrete set for all $n \in \{1, \dots, N\}$. MATLAB runs out of memory when the dimension becomes large and approximating by rounding off to integer values sometimes results in solutions that are too far from being optimal. A common procedure adopted when one faces such a curse of dimensionality is the use of Approximate Dynamic Programming (ADP) (see [12]). While there are several ADP algorithms to choose from, we select a simple one based on a single stage look ahead. In this version of **ADP** which is a 1-step optimization, given the current state, we determine the best action so that the expected cost for this state and the next state is minimized.

IV. OTHER HEURISTICS

Before describing two heuristics that we develop to approximately solve the MDP, we revisit Huang, Walrand and Ramchandran [5] (details explained in Section I). We call their algorithm **HWR**. It is a remarkable online algorithm that does not use any historical information and makes decisions based only on current state information (such as D_t , S_t , C_t and U_t). As the authors of HWR argue, it is an ideal algorithm under

such situations when there is tremendous variability, and there may be too many parameters to fit while modeling as an MDP.

With that understanding, we are in a position to describe two heuristics that use the cyclic structure in Section III. Under such a setting, we first obtain a fluid model of the system by taking expectation of D_τ , S_τ and C_τ for all $\tau \in \{1, 2, \dots, N\}$. Note that it is relatively straightforward to estimate $\mathbb{E}[D_\tau]$, $\mathbb{E}[S_\tau]$ and $\mathbb{E}[C_\tau]$ as the sample mean from historical observations. Then we can obtain the resulting policy as the fluid model (FL) with the variables \bar{X}_τ , \bar{Y}_τ , \bar{Z}_τ and \bar{U}_τ .

$$\text{FL: Minimize } \sum_{\tau=1}^N \mathbb{E}[C_\tau] \bar{X}_\tau$$

Subject to the following $\forall \tau \in \{1, \dots, N\}$

$$\bar{X}_\tau + \bar{Y}_\tau + \bar{Z}_\tau \geq \mathbb{E}[D_\tau]$$

$$0 \leq \bar{Y}_\tau \leq \mathbb{E}[S_\tau]$$

$$\bar{Z}_\tau \leq c_{dis}$$

$$-\min\{\bar{Z}_\tau, 0\} \leq c_{char}$$

$$\frac{N}{24}(\bar{U}_{\tau+1} - \bar{U}_\tau) = -\max\{\bar{Z}_\tau, 0\}/\eta_{dis} + \eta_{char} \min\{\bar{Z}_\tau, 0\}$$

$$0 \leq \bar{U}_\tau \leq K$$

$$\bar{X}_\tau \geq 0,$$

where $\bar{U}_{N+1} = \bar{U}_1$.

The two heuristics that we are about to explain essentially use the above fluid model's solution. The idea for the heuristics is based on Remark 3. The state at time t is the 5-tuple $(\{t/N\}, D_t, S_t, C_t, U_t)$. However, it is not easy to determine in each state which case to choose among (i), (ii) or (iii) of Remark 3. Thus as an approximation we consider guidelines provided by the fluid model and propose two heuristics: TBA (threshold-based approximation) and NOA (naive opportunistic algorithm).

Heuristic TBA: Given the state at time t namely, $(\{t/N\}, D_t, S_t, C_t, U_t)$, we determine Z_t so that at time $t+1$, U_{t+1} is as close to $\bar{U}_{\{t/N\}+1}$ by appropriately charging or discharging. The goal is to reach threshold level $\bar{U}_{\{t/N\}+1}$ in the next time. Thus TBA is as follows (with $n = \{t/N\}$):

$$\text{if } U_t < \bar{U}_{n+1}, Z_t = -\min\left\{\frac{N}{24} \frac{(\bar{U}_{n+1} - U_t)}{\eta_{char}}, c_{char}\right\}$$

else if $U_t = \bar{U}_{n+1}$, $Z_t = 0$,

else $Z_t = \min\left\{\frac{N}{24}(U_t - \bar{U}_{n+1})\eta_{dis}, c_{dis}\right\}$. In all the above cases, $X_t = \max\{D_t - S_t - Z_t, 0\}$.

Heuristic NOA: Given the state at time t namely, $(\{t/N\}, D_t, S_t, C_t, U_t)$, we adopt a naive (but intuitive) strategy – if C_t is cheap, charge the ESD as much as possible; and if D_t is much higher than S_t , discharge as much as possible; otherwise do what the fluid model suggests. For that we use $\mathbb{E}[C_N]$ as the grand average cost (computed over entire cycle N), and $\text{Var}[C_N]$ the corresponding grand variance; also ϕ_c and ϕ are parameters to be tuned. It leads to the following NOA (with $n = \{t/N\}$):

if $C_t < \mathbb{E}[C_N] - \phi_c \sqrt{\text{Var}[C_N]}$, then $Z_t =$

$$-\min\left\{N \frac{(K - U_t)}{24\eta_{char}}, c_{char}\right\} \text{ otherwise,}$$

if $D_t - S_t > \mathbb{E}[D_n] - \mathbb{E}[S_n] + \phi \sqrt{\text{Var}[D_n] + \text{Var}[S_n]}$,

$$Z_t = \min\left\{D_t - S_t - \bar{X}_n, \frac{N U_t \eta_{dis}}{24}, c_{dis}\right\}$$

else

(i.e. $D_t - S_t < \mathbb{E}[D_n] - \mathbb{E}[S_n] + \phi \sqrt{\text{Var}[D_n] + \text{Var}[S_n]}$)

$Z_t = \min\left\{\max(S_t - D_t, \bar{Z}_n), c_{char}, \frac{N(K - U_t)}{24\eta_{char}}\right\}$. In all the above cases, $X_t = \max\{D_t - S_t - Z_t, 0\}$.

V. NUMERICAL EXPERIMENTATION AND RESULTS

We obtained 26 days of demand, supply and cost data in a single month that we split into two groups: 16 days for training our models and 10 days for testing them. Our main purpose was to get a representative sample that adequately captures the deterministic and stochastic variability over time. In that spirit we collected demand data from households (see Acknowledgments section for credits), solar PV supply data from NREL (<http://www.nrel.gov/midc/>) and cost data at 2 granularities: 1-hour pricing (<https://www.nationalgridus.com>) and 5-minute pricing (<http://iso-ne.org/>). Next we used the training data to estimate/fit parameters in the MDP model described in Section III. To estimate d_t and δ_t for any $t \in [1, N]$, we use $D(1, t)$, $D(2, t)$, \dots , $D(16, t)$, the 16 realized demands and compute $\delta_t = \min_i[D(i, t)]$ and $d_t = \max_i[D(i, t)] - \delta_t$. Likewise for c_t and γ_t . In case of supply s_t , the minimum value is zero. Then for the DTMCs $\{Z_t^d, t \geq 0\}$, $\{Z_t^s, t \geq 0\}$, and $\{Z_t^c, t \geq 0\}$, we first arbitrarily select the number of states M . The state space is a set of discrete values $0, \frac{1}{M-1}, \frac{2}{M-1}, \dots, 1$. Then we estimate the elements of \mathcal{P}^d , \mathcal{P}^s , and \mathcal{P}^c as the respective frequency of transitioning based on the 16 days' data. For all the numerical experiments we use $c_{char} = c_{dis}$ and $\eta = \eta_{char} = \eta_{dis}$.

A. Results of MDP-based Analysis

Recall from Section III that due to curse of dimensionality we are unable to build very large sized probability matrices. Using the version of MATLAB installed, the following is the largest size problem that can be solved before memory limits are reached: storage efficiency $\eta = 1$; 1-hour intervals; and number of states in demand, storage and cost Markov chain is (3,3,2). For the analysis we considered the case where the hours of ‘‘average’’ demand in storage, i.e. $K/\mathbb{E}[D_n]$ is 1.64 while the Hours to fully charge/discharge, i.e. K/c_{char} is 2. We use the ratio of average PV supply to average demand as 0.59 to avoid non-trivial solution and tractability.

Once the training data is used to model the system, we perform a simulation (sampling from fitted data). Naturally, the MDP solution would be the best and we benchmark the four heuristics ADP, HWR, TBA and NOA. The results are described in Fig. 2 where we compare the four heuristic algorithms with the y-axis denoting $(b - a)/a$ where a is the MDP-based average cost and b is the corresponding heuristic's average cost. While the left side display corresponds to the above conditions and models fitted with training data, the right side figure is based on 10 randomly generated examples to test. From the figure it is clear that ADP is reasonable.




Fig. 2. Performance of heuristics: fraction higher than MDP solution

B. Benchmarking Heuristics Against ADP

Here we compare the four heuristics ADP, HWR, TBA and NOA but without MDP. Motivated by the excellent performance of ADP, we compare the other 3 algorithms against ADP in the next set of experiments. However, unlike the previous set, here we consider the test model based on real data (see the discussion at the beginning of Section V for details). In addition we use storage efficiency $\eta = 0.85$, and also both 1-hour/5-min intervals (corresponding to $N = 24$, & $N = 288$). While we will consider variations, we will mainly consider the baseline of: Hours of “average” demand in storage, i.e. $K/\mathbb{E}[D_n]$ as 2.17; Hours to fully charge/discharge, i.e. K/c_{char} as 8; Ratio of average PV supply to average demand as 0.468. We first estimate parameters using training data. We tried various alternatives for size of the state space. We chose number of states in demand, storage and cost Markov chain to be (4,4,4). Incidentally, when we increased the number of states to (10,10,10), the results remain the same.

For the testing we used 10 days of demand, supply and cost data (from the same month as training) assuming $U_0 = K/2$. For NOA, we selected tolerance parameters $\phi_c = \phi = 0.25$ by testing several options. Interestingly the 0.25 value is robust and the solutions do not change with much higher or lower values. Since the ADP algorithm always outperforms other heuristics, in Fig. 3, 4, and 5, we compare the three heuristic algorithms against ADP with the y-axis denoting $(b - a)/a$ where a is the ADP-based average cost and b is the corresponding heuristic’s average cost. Although HWR in general performs reasonably well, we observe that TBA achieves a lower cost than HWR in all the 5-minute cases.




Fig. 3. Heuristics’ performance over varying storage capacity (via $K/\mathbb{E}[D_n]$) keeping c_{char} constant

In Fig. 3, we fix the maximum charging rate c_{char} and vary




Fig. 4. Heuristics’ performance over varying storage capacity (via $K/\mathbb{E}[D_n]$) keeping $K/c_{char} = 8$




Fig. 5. Heuristics’ performance over varying hours to completely charge/discharge storage (via K/c_{char}) keeping K fixed

the storage capacity K . Note that the ratio K/c_{char} is 4, 8, 16, and 32 hours for the four cases, respectively. We observe from Fig. 3 that HWR performs reasonably well in all cases, and yields a 2% – 10% more cost than ADP. In Fig. 4, we fix the ratio $K/c_{char} = 8$ and vary the storage capacity K . This is a more practical setting since the maximum charging rate usually grows (nearly) proportionally to the storage capacity. The performance gap between ADP and the three heuristics increases with the battery capacity.

Finally, in Fig. 5 we fix the storage capacity K and vary the capacity to charging rate ratio K/c_{char} . The parameter setting in our simulation is motivated by the development of fast-charging batteries, for example, the lithium-ion titanate batteries are capable of recharging to 95% of full capacity within approximately ten minutes [13]. We observe from Fig. 5 that the performance of HWR heavily depends on the ratio K/c_{char} : it achieves almost the same performance as ADP when this ratio is larger than 8 (i.e., it takes more than 8 hours to fully charge the storage); however, for fast-response storage devices with $K/c_{char} \leq 1$, ADP outperforms HWR.

C. The Value of Storage and PV

Although there is no operational cost other than buying power from the grid, there are costs to install a solar PV system and/or an ESD. A natural question to ask is whether the PV and/or ESD installation was worth it. For that we consider two parameters: value of storage and value of PV and storage. We use the same test data as the previous sub-section and policy based on ADP. In Fig. 6 and 7, the y-axis denotes $(a - b)/a$ where b is the average cost using both PV and storage, while

a in the left bars correspond to the (optimal) use of only PV¹, and a in the right bars correspond to the use of neither PV nor storage.

The left bars present the percentage cost savings due the operation of storage, and can be therefore viewed as an illustrator on the value of storage. Similarly, the right bars illustrates the value of storage and PV. We observe from Fig. 6 that the value of storage increases sharply with the storage capacity K , when the maximum charging rate c_{char} grows in proportional to K . Fig. 7 shows that both the value of storage and the value of PV increases with average solar PV generation.




Fig. 6. Value of storage/PV over varying storage capacity (via $K/\mathbb{E}[D_n]$) keeping $K/c_{char} = 8$




Fig. 7. Value of storage/PV over varying ratio of average solar supply to average demand

VI. CONCLUSION

Although deceptively easy to state, the problem of determining energy mix from the grid, renewable source and storage device is a fairly complex one to solve. We explored four heuristics ADP, HWR, TBA and NOA. The following were our findings. ADP outperforms the other three heuristics in all cases, but the algorithm is relatively low in fidelity (because of poor fitting). Except when charging/discharging rates are high, in general all algorithms perform well for 1-hour data. NOA performs quite poorly in some of the 5-minute data cases (but NOA does well when others fail). Tuning tolerance and coefficient parameters had virtually no effect on NOA.

HWR is an easily implementable algorithm that needs no training. Its performance to a large extent depends on the number of hours to fully charge storage (i.e. K/c_{char}). It

achieves almost the same average cost as ADP when K/c_{char} is large (e.g. > 8). On the other hand, ADP significantly outperforms HWR for the case with $K/c_{char} \leq 1$. The insight we obtain from numerical comparisons has led to on-going efforts to design new ADP-based heuristics to better explore the flexibility provided by energy storage.

Under the simple one-step look-ahead (ADP) algorithm, value of storage is not too high. Demand and PV supply are correlated making a case for PV value. Value of storage would improve greatly if the storage size increases along with speed of charging and discharging. As solar penetration goes higher, storage has more value.

ACKNOWLEDGMENT

This material is based on work done while the first author was visiting Imperial College and SUTD, and partially supported by the AFOSR under Contract No. FA9550-13-1-0008 and the MIT-SUTD International Design Center (IDC) Grant IDG21400103. Thanks to Gareth Jones for collecting household power consumption data from <http://www.doc.ic.ac.uk/~dk3810/data/>

REFERENCES

- [1] C. Ren, D. Wang, B. Urgaonkar, and A. Sivasubramanian, "Carbon-aware energy capacity planning for datacenters," in *Modeling, Analysis & Simulation of Computer and Telecommunication Systems (MAS-COTS), 2012 IEEE 20th International Symposium on*. IEEE, 2012, pp. 391–400.
- [2] P. van de Ven, N. Hegde, L. Massoulié, and T. Salonidis, "Optimal control of end-user energy storage," 2013. [Online]. Available: <http://arxiv.org/pdf/1203.1891.pdf>
- [3] Y. Xu and L. Tong, "On the operation and value of storage in consumer demand response," *Submitted to IEEE Annual Conference on Decision and Control (CDC)*, 2014.
- [4] I. Koutsopoulos, V. Hatzis, and L. Tassioulas, "Optimal energy storage control policies for the smart power grid," *Proc. of IEEE International Conference on Smart Grid Communications (SmartGridComm)*, 2011.
- [5] L. Huang, J. Walrand, and K. Ramchandran, "Optimal demand response with energy storage management," in *Smart Grid Communications (SmartGridComm), 2012 IEEE Third International Conference on*. IEEE, 2012, pp. 61–66.
- [6] Y. Zhou, A. Scheller-Wolf, N. Secomandi, and S. Smith, "Managing wind-based electricity generation in the presence of storage and transmission capacity," 2013.
- [7] A. Federgruen and P. Zipkin, "An inventory model with limited production capacity and uncertain demands i: The average-cost criterion," *Mathematics of Operations Research*, vol. 11, no. 2, pp. 193–207, 1986.
- [8] S. P. Sethi and F. Cheng, "Optimality of (s, s) policies in inventory models with markovian demand," *Operations Research*, vol. 45, no. 6, pp. 931–939, 1997.
- [9] K. C. Divya and J. Østergaard, "Battery energy storage technology for power systems - an overview," *Electric Power Systems Research*, vol. 7, no. 4, pp. 511–520, 2009.
- [10] Z. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. Wang, M. Marwah, and C. Hyser, "Renewable and cooling aware workload management for sustainable data centers," in *ACM SIGMETRICS Performance Evaluation Review*, vol. 40, no. 1. ACM, 2012, pp. 175–186.
- [11] S. M. Ross, *Introduction to stochastic dynamic programming: Probability and mathematical*. Academic Press, Inc., 1983.
- [12] W. B. Powell, *Approximate Dynamic Programming: Solving the curses of dimensionality*. John Wiley & Sons, 2007, vol. 703.
- [13] M. Etezadi-Amoli, K. Choma, and J. Stefani, "Rapid-charge electric-vehicle stations," *IEEE Trans. Power Delivery*, vol. 25, no. 3, pp. 1883–1887, 2010.

¹We note that without storage, the optimal operation of PV is trivial: simply use as much solar generation as possible to fulfill the current demand.