Computational Management Science manuscript No.
(will be inserted by the editor)

A Dynamic Programming Model of Energy Storage and Transformer Deployments to Relieve Distribution Constraints

Xiaomin Xi · Ramteen Sioshansi

Received: 14 November, 2013 / Revised: 19 May, 2014

Abstract We introduce a stochastic dynamic programming (SDP) model that co-optimizes multiple uses of distributed energy storage, including energy and ancillary service sales, backup capacity, and transformer loading relief, while accounting for market and system uncertainty. We propose an approximation technique to efficiently solve the SDP. We also use a case study with high residential loads to demonstrate that a deployment consisting of both storage and transformer upgrades decreases costs and increases value relative to a transformer-only deployment.

Keywords Energy storage · stochastic dynamic program · electricity distribution

Mathematics Subject Classification (2000) 90C15 · 90C39 · 90B05

1 Introduction

Recent developments in electric power systems have increased interest in energy storage. Graves et al (1999) note the development of markets, which signal the cost and value of many of the services that storage can provide. The increasing use of electric and electronic devices by consumers can also yield extreme demand peaks. Of particular concern are grid-chargeable plug-in electric vehicles (PEVs). Collins and Mader (1983) and Mohseni and Stevie (2009) analyze possible PEV adoption in two regions of the United States finding that significant clustering of PEVs, which can yield extremely high distribution-level loads, is possible. Conversely, Kintner-Meyer et al (2007) and Sioshansi and Denholm (2010) show that there is sufficient transmission

Integrated Systems Engineering Department, The Ohio State University, Columbus, Ohio, United States
Phone: +1-614-292-3932
E-mail: xi.12@osu.edu, sioshansi.1@osu.edu

and generation capacity to serve charging loads with relatively high PEV penetration levels. Thus, at least initially, the negative impacts of PEVs may be at the distribution level.

The distribution system is normally built to accommodate the anticipated peak demand. This can be inefficient, however, since the system may only achieve this peak during a handful of hours each year. An alternative is to site storage on the constrained side of the distribution system. By charging storage when distribution is unconstrained and discharging when loads are higher, the distribution system can be downsized. Moreover, such storage can provide additional value to the utility, system operator (SO), or customers beyond the distribution benefits. For instance, by operating storage in a dynamic islanding mode it can provide backup energy to customers if there is a service outage. Similarly, storage can be used to provide ancillary services (AS) or to arbitrage diurnal day-ahead or real-time energy price differences. AS are excess generating capacity that a utility or SO reserves to provide a buffer for real-time deviations between actual and forecasted energy demand or supply. Nourai (2007) discusses a 1 MW distributed sodium-sulfur battery used by American Electric Power (AEP) to relieve a distribution-level transformer in West Virginia.

Using storage for multiple applications presents operational challenges, however. This is because using storage for one service may prevent or reduce its ability to provide another. For instance, discharging storage to relieve a distribution constraint leaves less energy available for backup energy. In this paper we adapt the stochastic dynamic programming (SDP) model developed by Xi et al (2014) to study the use of energy storage to relieve distribution constraints and to provide energy and AS sales and backup energy. More specifically, the model used here is similar to that in the work of Xi et al (2014), except we include a more detailed representation of the distribution transformer's and storage's operating constraints and aging characteristics. Our model accounts for market and system uncertainty in deriving an optimal storage use policy. We use an approximation technique, developed by Nascimento and Powell (2009), to solve the SDP. Using a case study based on several residential homes and PEVs connected to an overloaded distribution-level transformer, we demonstrate that supplementing the transformer with storage decreases costs and increases value relative to a transformer-only deployment. The remainder of this paper is organized as follows: Section 2 presents the SDP model, Section 3 presents our proposed solution algorithm, Section 4 provides the details of our case study, Section 5 summarizes our results, and Section 6 concludes.

2 Stochastic Dynamic Programming Model

We model distribution-sited storage at hourly timesteps. Our model assumes that storage provides four services: arbitraging hourly energy price differences, selling AS, providing backup energy to serve distribution-level loads in the

event of a system outage, and relieving transformer loading. Although storage can provide different AS products, only regulation services are modeled since Tomić and Kempton (2007) show that they are the most valuable. Regulation is an AS product that is used to maintain the system's frequency and voltage within acceptable bounds by exactly matching real-time energy demand and supply. Regulation is further subdivided into regulation-up and -down services. A regulation up provider must increase its output from a baseline level if called in real-time. A regulation down provider must instead decrease its output. We assume that the market includes a single regulation product. Thus, an entity that sells regulation capacity may be called by the SO to provide both regulation-up and -down services in real-time. Regulation services are assumed to receive a one-part capacity payment only. This payment is determined by the amount of regulation capacity offered and a uniform market-clearing price for regulation capacity. If all real-time calls for regulation energy are satisfied during the contract period, the payment is given by the product of the contracted capacity and the price. Otherwise, the payment is prorated based on the amount of requested regulation energy actually provided. These assumptions are consistent with the definition and pricing of regulation used by PJM Interconnection, the largest SO market in the United States.

In each hour, the storage operator decides how much energy to charge and discharge (which can be sold to the market or used to serve the distribution-level load) and the amount of regulation capacity to sell. These decisions are made based on the state of the market (*i.e.*, energy and regulation prices and anticipated regulation calls), state of the system (*i.e.*, whether there is an outage or not), and an exogenous distribution-level load. If there is no system outage, storage can be charged and discharged from the grid and can provide regulation service, otherwise it can only be discharged to serve distribution-level loads. We further assume that the storage is sufficiently small compared to the rest of the market that storage decisions do not impact the market or system. Our model formulation follows the conventions and notation used by Powell (2007) and Xi et al (2014). Random variables with a subscript t are unknown (stochastic) before hour t and become known (deterministic) at hour t. The overall model structure is similar to that used by Xi et al (2014), with the exception of our handling of the distribution transformer constraint here.

2.1 Model Parameters

$\bar{P}^b$: maximum power capacity of storage [kW]
$\bar{R}^{max}$: maximum energy level of storage [kWh]
$\bar{R}^{min}$: minimum energy level of storage [kWh]
η^c: charging efficiency of storage
η^d: discharging efficiency of storage
$\bar{P}^{tr}$: rated power capacity of transformer [kW]
$V^{tr}(v)$: penalty for operating transformer above rated capacity, as a function of loading [$]

V^L: penalty for unserved building load [$/kWh]
γ: hourly discount factor

Our model assumes that the storage has a minimum storage level, $\bar{R}^{min}$. This accounts for technologies, such as lithium-ion batteries, which suffer extreme cycle-life degradation if the state of charge falls too low. The unitless ratios, η^c and η^d, reflect efficiency losses from charging and discharging storage, respectively.

$\bar{P}^{tr}$ is the transformer's rated power capacity. The transformer can be operated above this capacity, however, which accelerates transformer aging and imposes a cost. This accelerated aging is typically estimated using hot-spot or top-oil temperature models. Susa et al (2005) and Gong et al (2011) present two examples of such models. These models estimate the effect of operating the transformer above its rated capacity on its overall lifetime. Combining this aging effect with an assumed transformer replacement cost gives a cost for operating the transformer above its rated capacity, which we denote $V^{tr}(v)$. This function, which we assume to be convex, accrues on an hourly basis and represents the cost incurred in each hour during which the transformer is operated above its rated capacity. Allowing the transformer to be operated above its rated capacity is an added feature of our model, compared to that developed by Xi et al (2014). V^L is the cost penalty for curtailing distribution-level loads, which are caused by system outages and distribution constraints. Stored energy can be used, however, to reduce such curtailments.

2.2 Decision Variables

e_t^d: energy discharged for sales in hour t [kWh]
e_t^c: energy charged into storage in hour t [kWh]
e_t^l: energy discharged from storage in hour t to serve distribution-level load [kWh]
l_t: distribution-level load met in hour t [kWh]
k_t: regulation capacity sold in hour t [kW-h]
v_t: amount transformer is overloaded in hour $t-1$ [kW]

We also define $A_t = \left(e_t^d, e_t^c, e_t^l, l_t, k_t, v_t\right)$ as a vector of hour-t decision variables.

2.3 State Variables

x_t: total energy in storage at the beginning of hour t [kWh]
p_t^e: market price of energy in hour t [$/kWh]
p_t^r: market price of regulation in hour t [$/kW-h]
D_t: distribution-level energy demand in hour t [kWh]
I_t: binary variable indicating if there is a system outage (equals 1) in hour t
δ_t^u: dispatch-to-contract ratio of regulation-up in hour $t-1$

δ_t^d: dispatch-to-contract ratio of regulation-down in hour $t-1$

Following the work of Kempton and Tomić (2005) and Xi et al (2014), we use dispatch-to-contract ratios, δ_t^u and δ_t^d, to model the relationship between regulation capacity sales and real-time regulation deployments. These ratios are defined as the actual amount of regulation-up and -down energy called in real-time, divided by the amount of regulation capacity procured by the SO. Thus, by definition, the terms $\delta_{t+1}^u \cdot k_t$ and $\delta_{t+1}^d \cdot k_t$ are the amount of regulation-up and -down energy called by the SO in hour t, respectively. These ratios can be estimated using historical system data. For instance, Kempton and Tomić (2005) use several years of historical data to estimate an average dispatch-to-contract ratio for regulation in the California ISO market of 0.08. This indicates that each kW of regulation capacity sold in the market for one hour results, on average, in 0.08 kWh of regulation energy being called in real-time.

We assume that the energy and regulation prices, p_t^e and p_t^r, are all non-negative. Although regulation prices are typically non-negative, negative energy prices are occasionally observed in energy markets. Negative energy prices typically occur because of unpriced non-convexities stemming from unit commitment or subsidies provided to renewable generators. Negative energy prices tend to occur more often in real-time, as opposed to day-ahead, energy markets. We assume storage use is scheduled based on day-ahead prices, when this non-negativity assumption is more likely to hold.

We also define $S_t = \left(x_t, p_t^e, p_t^r, D_t, I_t, \delta_t^u, \delta_t^d\right)$ as a vector of hour-t state variables.

2.4 Exogenous Variables

The state variables, p_t^e, p_t^r, D_t, I_t, δ_t^u, and δ_t^d, are assumed to evolve exogenously of the decision variables. We define $\hat{p}_t^e$, $\hat{p}_t^r$, $\hat{D}_t$, $\hat{I}_t$, $\hat{\delta}_t^u$, and $\hat{\delta}_t^d$ as exogenous random variables that represent the change in the value of these variables between hour $t-1$ and hour t. These random variables may be dependent on one another, which would imply that the exogenous state variables are probabilistically dependent. We define $\omega_t = \left(p_t^e, p_t^r, D_t, I_t, \delta_t^u, \delta_t^d\right)$ as a vector of hour-t exogenous state variables, thus we can also define $S_t = (x_t, \omega_t)$.

2.5 State-Transition Function

The exogenous state variables are assumed to evolve randomly according to the following transition functions:

$$p_{t+1}^e = \hat{p}_{t+1}^e + p_t^e,$$

$$p_{t+1}^r = \hat{p}_{t+1}^r + p_t^r,$$

$$D_{t+1} = \hat{D}_{t+1} + D_t,$$

$$I_{t+1} = \hat{I}_{t+1} + I_t,$$
$$\delta^u_{t+1} = \hat{\delta}^u_{t+1} + \delta^u_t,$$

and:

$$\delta^d_{t+1} = \hat{\delta}^d_{t+1} + \delta^d_t.$$

We define the amount of regulation-up that is unserved in hour $t-1$ as:

$$n^u_t = \max\{0, \delta^u_t \cdot k_{t-1} - \eta^d \cdot (x_{t-1} - \bar{R}^{min}) + e^d_{t-1} + e^l_{t-1} - e^c_{t-1}\}. \quad (1)$$

We similarly define unserved regulation-down as:

$$n^d_t = \max\{0, \delta^d_t \cdot k_{t-1} - (\bar{R}^{max} - x_{t-1})/\eta^c - e^d_{t-1} - e^l_{t-1} + e^c_{t-1}\}. \quad (2)$$

Equations (1) and (2) define unserved regulation as the difference between the amount of regulation called and the maximum amount of regulation that can be feasibly provided without violating storage energy or power constraints. Based on these definitions, the amount of energy in storage evolves according to:

$$x_{t+1} = x_t + \eta^c \cdot (e^c_t + \delta^d_{t+1} \cdot k_t - n^d_{t+1}) - (e^d_t + e^l_t + \delta^u_{t+1} \cdot k_t - n^u_{t+1})/\eta^d.$$

We also define the proportion of regulation energy called in hour $t-1$ that is unserved as:

$$n^r_t = \frac{n^u_t + n^d_t}{(\delta^u_t + \delta^d_t) \cdot k_{t-1}}. \quad (3)$$

2.6 Constraints

The total amount of energy charged and discharged in each hour must be within the power capacity of the storage device:

$$0 \leq e^c_t + k_t \leq \bar{P}^b, \quad (4)$$

and:

$$0 \leq e^d_t + e^l_t + k_t \leq \bar{P}^b. \quad (5)$$

The energy level of the storage device is constrained to be within its upper and lower bounds:

$$\bar{R}^{min} \leq x_t \leq \bar{R}^{max}. \quad (6)$$

Net power flows are only allowed into or out of the distribution system if there is not an outage:

$$-\bar{P}^{tr} \cdot (1 - I_t) - v_{t+1} \leq l_t - e^l_t + e^c_t - e^d_t - k_t, \quad (7)$$

and

$$l_t - e^l_t + e^c_t - e^d_t + k_t \leq \bar{P}^{tr} \cdot (1 - I_t) + v_{t+1}. \quad (8)$$

Constraints (7) and (8) also define v_{t+1}, the amount that the transformer is operating above its rated capacity. Charging storage and selling energy and regulation are not possible if there is an outage:

$$e_t^d, e_t^c, k_t, v_{t+1} = 0, \text{ if } I_t = 1. \tag{9}$$

Constraints (7) through (9) force the distribution-level load to be either served using stored energy or curtailed in any hour with a system outage. We further constrain the distribution-level load served to be no greater than the actual demand:

$$l_t \leq D_t. \tag{10}$$

We also assume that the amount of regulation capacity sold cannot exceed available nameplate transformer capacity:

$$k_t \leq \max\{0, \bar{P}^{tr} - D_t\}. \tag{11}$$

All of the decision variables are non-negative:

$$e_t^d, e_t^c, e_t^l, l_t, k_t, v_t \geq 0. \tag{12}$$

We further define $\mathcal{A}_s$ as the set of decision vectors, a, that are feasible in (4) through (12) when the system is in state s.

2.7 Objective Function

The hour-t objective function contribution is given by:

$$C_t(S_t, a_t) = p_t^e \cdot (e_t^d - e_t^c) - V^L \cdot (D_t - l_t) + p_{t-1}^r \cdot k_{t-1} \cdot (1 - n_t^r) - V^{tr}(v_t).$$

The first term, $p_t^e \cdot (e_t^d - e_t^c)$, represents revenues from net energy sales, excluding regulation energy provided. The second term, $-V^L \cdot (D_t - l_t)$, represents the penalty for unserved distribution-level loads. The third term, $p_{t-1}^r \cdot k_{t-1} \cdot (1 - n_t^r)$, are the regulation capacity payments, which are prorated based on how much regulation energy is provided in real-time. The fourth term, $-V^{tr}(v_t)$, represents the penalty for aging the transformer.

We define a policy, $A_t^\pi(S_t)$, as a mapping between an hour-t state variable, S_t, and an hour-t feasible decision, $a \in \mathcal{A}_{S_t}$. We further define $\mathit{\Pi}$ as the set of all feasible policies. For each $\pi \in \mathit{\Pi}$, the total expected discounted profit from hour t is defined as:

$$G_t^\pi(S_t) = \mathbb{E}\left[\sum_{\tau=t}^{T+1} \gamma^{\tau-t} C_\tau(S_\tau, A_\tau^\pi(S_\tau)) \middle| S_t\right],$$

where T is the optimization horizon. The objective is then to find an optimal policy, π^*, which satisfies:

$$G_t^{\pi^*}(S_t) = \sup_{\pi \in \mathit{\Pi}} G_t^\pi(S_t), \ \forall\ t = 1, 2, \ldots, T+1.$$

2.8 System Schematic

Figure 1 is a schematic of the overall system, showing the interaction among the state and decision variables. The l_t kW of distribution-level demand served is satisfied using e_t^l kW from storage and the remaining $l_t - e_t^l$ kW is satisfied using energy from the grid. The quantities $e_t^c + \delta_{t+1}^d \cdot k_t - n_{t+1}^d$ and $e_t^d + \delta_{t+1}^u \cdot k_t - n_{t+1}^u$ represent gross energy charged and discharged into storage for grid services (*i.e.*, direct energy sales and regulation services). The quantity $l_t - e_t^l + e_t^c - e_t^d + (\delta_{t+1}^d - \delta_{t+1}^u) \cdot k_t + n_{t+1}^u - n_{t+1}^d$ represents net energy drawn from the grid through the transformer. This quantity may be negative, for instance if the distribution-level load is low and storage discharges more net energy than the building draws from the grid.

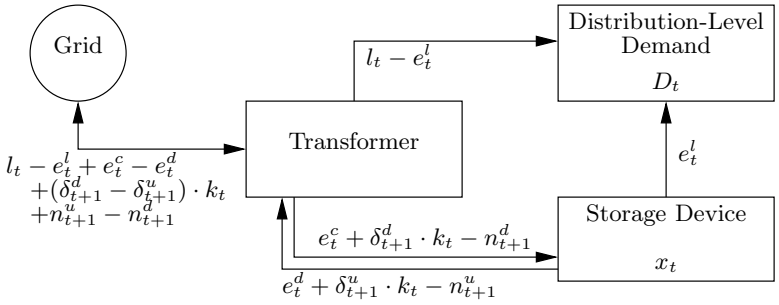

Fig. 1: System schematic

3 Solution Technique

One approach to find a near-optimal policy is to discretize all of the continuous decision and state variables and apply the dynamic programming algorithm to solve the problem. This is impractical given the large dimension of the decision and state variables. Hence we propose adapting an approximate stochastic dynamic programming (ASDP) method developed by Nascimento and Powell (2009). We also use sampling techniques to derive statistical upper and lower bounds for the true optimum of the SDP, demonstrating the quality of the solutions found by the algorithm.

We define:

$$x_t^{post} = f^{post}(x_t, a_t) = x_t + \eta^c \cdot e_t^c - (e_t^d + e_t^l)/\eta^d,$$

as the post-decision hour-t storage level. This represents the amount of energy in storage immediately after the hour-t decisions are made, but before the amount of energy needed for regulation is known (*i.e.*, it accounts for the

amount of energy charged and discharged for arbitrage and serving the distribution loads). Note that we can also refer to x_t as the pre-decision hour-t storage level, since it is the storage level before any hour-t decisions are made. We further define $S_t^{post} = (x_t^{post}, \omega_t)$ as the post-decision hour-t state variable and:

$$F_t^*(x_t, \omega_t) = G_t^{\pi^*}(S_t),$$

and:

$$F_t^{post}(x_t^{post}, \omega_t) = \mathbb{E}\left[F_{t+1}^*(x_{t+1}, \omega_{t+1}) \middle| \omega_t\right],$$

to be the optimal value function around the pre- and post-decision hour-t storage level, respectively. By definition $F_t^*(x_t, \omega_t)$ satisfies the following Bellman-type equation:

$$F_t^*(x_t, \omega_t) = \max_{a_t \in A_{S_t}} \left\{C_t(S_t, a_t) + \gamma F_t^{post}(x_t^{post}, \omega_t)\right\}. \tag{13}$$

The ASDP algorithm requires $F_t^*(x_t, \omega_t)$ and $F_t^{post}(x_t^{post}, \omega_t)$ to be concave in x_t and x_t^{post}, respectively. We prove concavity in the following lemma.

Lemma 1 *If $V^{tr}(v_t)$ is convex in v_t and the energy and regulation prices, p_t^e and p_t^r, are all non-negative, then $F_t^*(x_t, \omega_t)$ and $F_t^{post}(x_t^{post}, \omega_t)$ are concave in x_t and x_t^{post}, respectively.*

Proof This is shown using an inductive argument. In period, T, the boundary condition and (3) give:

$$F_T^{post}(x_T^{post}, \omega_T) = \mathbb{E}\left[p_T^r \cdot \left(k_T - \frac{n_{T+1}^u + n_{T+1}^d}{\delta_{T+1}^u + \delta_{T+1}^d}\right) - V^{tr}(v_{T+1}) \middle| \omega_T\right], \tag{14}$$

where n_{T+1}^u and n_{T+1}^d are defined by (1) and (2). The terms n_{T+1}^u, n_{T+1}^d, and v_{T+1} are the only elements of (14) dependent on x_T^{post}. The term $-(n_{T+1}^u + n_{T+1}^d)$ is a piecewise-linear concave function of x_T^{post}. Due to the convexity assumption, $-V^{tr}(v_{T+1})$ is concave in x_T^{post}. Note, also, that system outages do not affect concavity, since an hour-T system outage simply forces all of n_{T+1}^u, n_{T+1}^d, and v_{T+1} to equal zero. Thus, $F_T^{post}(x_T^{post}, \omega_T)$ is concave in x_T^{post} as the expectation gives a convex combination of concave functions.

Furthermore, (13) shows that $F_T^*(x_T, \omega_T)$ is a positive-weighted sum of $C_T(S_T, a_T)$, which is linear in x_T, and $F_T^{post}(x_T^{post}, \omega_T)$, which is concave in x_T. Since constraints (4) through (12) define a convex feasible set for x_T, $F_T^*(x_T, \omega_T)$ is concave in x_T since it is the supremal convolution of a concave function, which Rockafellar (1970) shows to be concave.

For $t < T$ we have:

$$\begin{aligned} F_t^{post}(x_t^{post}, \omega_t) &= \mathbb{E}\left[F_{t+1}^*(x_{t+1}, \omega_{t+1}) \middle| \omega_t\right] \\ &= \text{Prob}\left\{I_{t+1} = 0 | \omega_t\right\} \cdot \mathbb{E}\left[F_{t+1}^*(x_{t+1}, \omega_{t+1}) \middle| \omega_t\right] \\ &\quad + \text{Prob}\left\{I_{t+1} = 1 | \omega_t\right\} \cdot \mathbb{E}\left[F_{t+1}^*(x_{t+1}, \omega_{t+1}) \middle| \omega_t\right], \end{aligned} \tag{15}$$

where equation (15) explicitly shows the relationship between I_{t+1} and the expectation involved in computing $F_t^{post}(x_t^{post}, \omega_t)$. Using an inductive argument, $F_{t+1}^*(x_{t+1}, \omega_{t+1})$ is concave in x_{t+1} and, thus, in x_t^{post} as well. Moreover, equation (15) defines $F_t^{post}(x_t^{post}, \omega_t)$ as a convex combination of two concave functions, showing the concavity of $F_t^{post}(x_t^{post}, \omega_t)$. The concavity of $F_t^*(x_t, \omega_t)$ follows by applying the same inductive argument to (13). □

Rather than working directly with the F_t^{post} functions the ASDP algorithm uses approximations, which are denoted $\hat{F}_t^{post}$. The algorithm works in two phases: Monte Carlo simulation is used in the first to iteratively update the $\hat{F}_t^{post}$ functions, improving the approximations; the approximations are then used in the second phase to determine a near-optimal policy.

3.1 Approximate Bellman Equation

$\hat{F}_t^{post}$ is assumed to be a concave piecewise-linear function of x_t^{post}. It is estimated by iteratively updating $\hat{F}_t^{post}$ at a discrete set of values. We denote the values at which the function is updated $\tilde{x}_{t,i}^{post}, i = 1, 2, \ldots, M_t$, and assume that they are rank ordered (*i.e.*, $\tilde{x}_{t,1}^{post} < \cdots < \tilde{x}_{t,M_t}^{post}$). $\hat{F}_t^{post}$ is estimated for a discrete set of exogenous state variables, which are denoted $\tilde{\omega}_t$. We define $\tilde{\Omega}_t$ as the discrete set of exogenous state variables at which $\hat{F}_t^{post}$ is estimated. The decision variables are also restricted to a discrete set of values in the first phase of the algorithm. We use $\tilde{a}_t$ to denote such a discrete decision vector and define $\tilde{\mathcal{A}}_{\tilde{S}_t}$ as the set of $\tilde{a}_t$ that satisfy the discretization and are feasible in constraints (4) through (12) when the system is in the discrete state, $\tilde{S}_t$. Further details of the discretization used are given in Section 4.5.

We denote the slopes of $\hat{F}_t^{post}(x_t^{post}, \tilde{\omega}_t)$ as:

$$\hat{m}_t(\tilde{\omega}_t) = \left(\hat{m}_t(\tilde{x}_{t,2}^{post}, \tilde{\omega}_t), \cdots, \hat{m}_t(\tilde{x}_{t,M_t}^{post}, \tilde{\omega}_t)\right), \tag{16}$$

where:

$$\hat{m}_t(\tilde{x}_{t,i}^{post}, \tilde{\omega}_t) = \frac{\hat{F}_t^{post}(\tilde{x}_{t,i}^{post}, \tilde{\omega}) - \hat{F}_t^{post}(\tilde{x}_{t,i-1}^{post}, \tilde{\omega})}{\tilde{x}_{t,i}^{post} - \tilde{x}_{t,i-1}^{post}}. \tag{17}$$

Figure 2 shows an example of a $F_t^{post}(x, \omega_t)$ function and its piecewise-linear approximation, in which the discrete exogenous state variable, $\tilde{\omega}_t$, is used and $M_t = 5$.

Using these slopes, equation (13) can be approximated in the discretized exogenous variable space as:

$$\hat{F}_t^*(x_t, \tilde{\omega}_t, \tilde{m}_t(\tilde{\omega}_t)) = \max_{\tilde{a}_t \in \mathcal{A}_{\tilde{S}_t}, y_t} \left\{ C_t(x_t, \tilde{\omega}_t, \tilde{a}_t) + \gamma \sum_{i=2}^{M_t} \hat{m}_t(\tilde{x}_{t,i}^{post}, \tilde{\omega}_t) \cdot y_{t,i} \right\}, \tag{18}$$

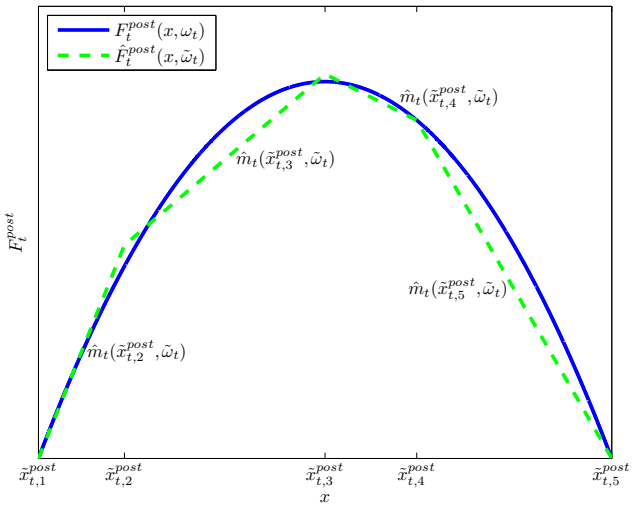

Fig. 2: Optimal objective function and piecewise-linear approximation

where:

$$\sum_{i=2}^{M_t} y_{t,i} = f^{post}(x_t, \tilde{a}_t), \tag{19}$$

and,

$$0 \leq y_{t,i} \leq \tilde{x}_{t,i}^{post} - \tilde{x}_{t,i-1}^{post}, \forall\ i = 2, \ldots, M_t. \tag{20}$$

Note that $\tilde{a}_t$ and $\tilde{\omega}_t$ are restricted to the discretization in (18). Since our original SDP involves exogenous state variables from a continuous space, we use the discretized exogenous state variable that is closest to the actual exogenous state variable. We could, instead, conduct multi-dimensional interpolation on x_t^{post} and ω_t, which could provide more robust results at higher computational costs. Alternatively, we could simulate directly from the discretized distribution, rather than rounding continuous variables to discretized values. Doing so would only provide value function approximations at the discretized points, however, whereas we aim to approximate the value function at points from the continuous distribution. Our results in Section 5 show that our interpolation and discretization scheme provides reasonably near-optimal policies with relatively small optimality gaps. Note also that if the discretizations, $\tilde{a}_t$, $\tilde{\omega}_t$, and $\tilde{x}_t^{post}$, are defined appropriately, then $x_t^{post} = f^{post}(x_t, \tilde{a}_t)$ will necessarily take on one of the M_t discrete values assumed.

3.2 ASDP Algorithm

Our ASDP approach is summarized by the pseudocode in Algorithms 1 and 2. The first phase of the algorithm, in which the $\hat{F}_t^{post}$ functions are estimated, is

summarized in Algorithm 1. We use j throughout this algorithm to denote the iteration number and $\hat{m}_t^{(j)}(\tilde{\omega}_t)$ denotes the slopes of the $\hat{F}_t^{post}$ functions after the jth iteration. This algorithm begins with an initialization (Step 1). This includes discretizing the action, exogenous state, and post-decision variables, choosing an initial value for the slopes of $\hat{F}_t^{post}$, and fixing the starting state of charge of storage, x_1. The starting values, $\hat{m}_t^{(0)}(\tilde{\omega}_t)$, are not restricted, except that they must be non-decreasing (*i.e.*, the $\hat{F}_t^{post}$ functions must be concave). Nascimento and Powell (2009) suggest using empirical results, consulting experts, or simply setting the slopes equal to zero. In our case study we initialize the slopes by generating a single sample path of exogenous state variables, $\{\omega_t\}_{t=1}^{T+1}$ and solving a deterministic optimization problem in which the sample path is fully known. The solution of this problem is used to estimate the slopes and the projection procedure in Steps 14–16 of Algorithm 1 is used to ensure concavity.

The remaining steps of the algorithm are repeated in each iteration. First a path of the exogenous state variables is randomly sampled from the continuous space (Step 3) and this sample path is rounded to the nearest discretized sample path (Step 4). The algorithm then cycles through hours 1 to $T-1$ of the optimization horizon (the algorithm does not cycle through hour T, since the exact computation of F_T^{post} is trivial). In each hour it first determines an optimal action variable from the discretized space, using the piecewise-linear approximation of F_t^{post} from the previous iteration (Step 6). The resulting post-decision hour-t storage level is computed (Step 7), and without loss of generality we assume that this is the ξth discretized value of $\tilde{x}_t^{post}$ (in rank order). The unserved regulation (Step 8) and change between the post-decision hour-t and pre-decision hour-$(t+1)$ storage level (Step 9), which we denote $\Delta\tilde{x}_{t+1}$, are computed. Two random samples of F_t^{post} are then computed (Step 10). These random samples, which we denote $\tilde{\theta}_{t+1,\xi}$ and $\tilde{\theta}_{t+1,\xi+1}$,[1] are computed using the approximation of F_{t+1}^* from the previous iteration. The random samples are then used to update the slopes (Step 12), based on a learning rate parameter, α_j. Note that we only update the slopes of the piecewise-linear function associated with the actual sampled exogenous state variable, $\tilde{\omega}_t$, and the breakpoint, $\tilde{x}_{t,\xi}^{post}$. Finally, we use a projection operation (Step 15) to ensure that the final updated slopes are non-increasing. Steps 3 through 19 of the algorithm are repeated J times, which yields a final set of piecewise-linear approximations of the F_t^{post} functions, which are denoted $\tilde{m}_t^{(J)}(\tilde{\omega}_t)$.

The second phase of the ASDP technique, which uses the final approximation to compute a near-optimal policy, is summarized in Algorithm 2. This algorithm begins by fixing the starting storage level (Step 1). It then iterates through each hour of the optimization horizon, first observing the exogenous state variable and then rounding this to the nearest discretized exogenous state variable (Step 3). If $t > 1$, unserved hour-$(t-1)$ regulation energy (Step 5) and

[1] In keeping with our convention that variables with the subscript t are stochastic before hour t and become known at hour t, we put a $t+1$ subscript on these random samples. This is because the random samples are computed based on a realization of $\tilde{\omega}_{t+1}$.

Algorithm 1 Phase 1 of ASDP Algorithm: Estimation of $\hat{F}_t^{post}$

1: Initialize:
   - Discretize a_t, ω_t, x_t^{post}
   - Initialize $\hat{m}_t^{(0)}(\tilde{\omega}_t)$
   - Fix x_1
2: **for** $j = 1$ to J **do**
3: Randomly generate a sample path, $\{\omega_t\}_{t=1}^{T+1}$, from the continuous distribution
4: Round the continuous sample path to the nearest discretized sample path, $\{\tilde{\omega}_t\}_{t=1}^{T+1}$
5: **for** $t = 1$ to $T-1$ **do**
6: $(\tilde{a}_t, y_t) \in \arg\max \left\{ \hat{F}_t^*(x_t, \tilde{\omega}_t, \tilde{m}_t^{(j-1)}(\tilde{\omega}_t)) \middle| \tilde{a}_t \in \mathcal{A}_{\tilde{S}_t}, (19), (20) \right\}$
7: Without loss of generality, suppose $f^{post}(x_t, \tilde{a}_t)$ is the ξth discretized value of $\tilde{x}_t^{post}$ in rank ordering (*i.e.*, $f^{post}(x_t, \tilde{a}_t) = \tilde{x}_{t,\xi}^{post}$)
8: $n_{t+1}^u \leftarrow \max\{0, \tilde{\delta}_{t+1}^u \cdot \tilde{k}_t - \eta^d \cdot (x_t - \bar{R}^{min}) + \tilde{e}_t^d + \tilde{e}_t^l - \tilde{e}_t^c\}$ and $n_{t+1}^d \leftarrow \max\{0, \tilde{\delta}_{t+1}^d \cdot \tilde{k}_t - (\bar{R}^{max} - x_t)/\eta^c - \tilde{e}_t^d - \tilde{e}_t^l + \tilde{e}_t^c\}$ {compute unserved regulation}
9: $\Delta\tilde{x}_{t+1} \leftarrow \eta^c \cdot (\tilde{\delta}_{t+1}^d \cdot \tilde{k}_t - n_{t+1}^d) - (\tilde{\delta}_{t+1}^u \cdot \tilde{k}_t - n_{t+1}^u)/\eta^d$ {compute change in storage level from regulation calls}
10: $\tilde{\theta}_{t+1,\xi} \leftarrow \left[\hat{F}_{t+1}^*(\tilde{x}_{t,\xi}^{post} + \Delta\tilde{x}_{t+1}, \tilde{\omega}_{t+1}, \tilde{m}_{t+1}^{(j-1)}(\tilde{\omega}_{t+1})) - \hat{F}_{t+1}^*(\tilde{x}_{t,\xi-1}^{post} + \Delta\tilde{x}_{t+1}, \tilde{\omega}_{t+1}, \tilde{m}_{t+1}^{(j-1)}(\tilde{\omega}_{t+1}))\right] / \left(\tilde{x}_{t,\xi}^{post} - \tilde{x}_{t,\xi-1}^{post}\right)$ and $\tilde{\theta}_{t+1,\xi+1} \leftarrow \left[\hat{F}_{t+1}^*(\tilde{x}_{t,\xi+1}^{post} + \Delta\tilde{x}_{t+1}, \tilde{\omega}_{t+1}, \tilde{m}_{t+1}^{(j-1)}(\tilde{\omega}_{t+1})) - \hat{F}_{t+1}^*(\tilde{x}_{t,\xi}^{post} + \Delta\tilde{x}_{t+1}, \tilde{\omega}_{t+1}, \tilde{m}_{t+1}^{(j-1)}(\tilde{\omega}_{t+1}))\right] / \left(\tilde{x}_{t,\xi+1}^{post} - \tilde{x}_{t,\xi}^{post}\right)$
11: **for** $i = 1, \ldots, M_t$ and $z \in \tilde{\Omega}_t$ **do**
12: $\rho_t(\tilde{x}_{t,i}^{post}, z) \leftarrow \begin{cases} (1-\alpha_j) \cdot \hat{m}_t^{(j-1)}(\tilde{x}_{t,i}^{post}, z) + \alpha_j \cdot \tilde{\theta}_{t+1,i}, & \text{if } z = \tilde{\omega}_t \\ & \text{and } i \in \{\xi, \xi+1\}; \\ \hat{m}_t^{(j-1)}(\tilde{x}_{t,i}^{post}, z), & \text{otherwise} \end{cases}$
13: **end for**
14: **for** $i = 1, \ldots, M_t$ and $z \in \tilde{\Omega}_t$ **do**
15: $\hat{m}_t^{(j)}(\tilde{x}_{t,i}^{post}, z) \leftarrow \begin{cases} \rho_t(\tilde{x}_{t,\xi}^{post}, \tilde{\omega}_t), & \text{if } z = \tilde{\omega}_t, \tilde{x}_{t,i}^{post} < \tilde{x}_{t,\xi}^{post} \\ & \text{and } \rho_t(\tilde{x}_{t,i}^{post}, z) \le \rho_t(\tilde{x}_{t,\xi}^{post}, \tilde{\omega}_t) \\ \rho_t(\tilde{x}_{t,\xi+1}^{post}, \tilde{\omega}_t), & \text{if } z = \tilde{\omega}_t, \tilde{x}_{t,i}^{post} > \tilde{x}_{t,\xi+1}^{post} \\ & \text{and } \rho_t(\tilde{x}_{t,i}^{post}, z) \ge \rho_t(\tilde{x}_{t,\xi+1}^{post}, \tilde{\omega}_t) \\ \rho_t(\tilde{x}_{t,i}^{post}, \tilde{\omega}_t), & \text{otherwise} \end{cases}$
16: **end for**
17: $x_{t+1} \leftarrow f^{post}(x_t, \tilde{a}_t) + \Delta\tilde{x}_{t+1}$
18: **end for**
19: **end for**

the starting hour-t storage level (Step 6) are computed. A near-optimal hour-t decision is then determined, based on the final piecewise-linear approximation of the F_t^{post} function (Step 8). Note that the rounded exogenous state variable, $\tilde{\omega}_t$, is used to determine what set of slopes to use in this maximization. Finally, the hour-t profit contribution is calculated (Step 9).

Algorithm 2 Phase 2 of ASDP Algorithm: Obtain a Near-Optimal Policy

1: Fix x_1
2: **for** $t = 1$ to T **do**
3: Observe ω_t from continuous distribution and round it to the nearest discrete $\tilde{\omega}_t$
4: **if** $t > 1$ **then**
5: $n_t^u \leftarrow \max\{0, \delta_t^u \cdot k_{t-1} - \eta^d \cdot (x_{t-1} - \bar{R}^{min}) + e_{t-1}^d + e_{t-1}^l - e_{t-1}^c\}$, $n_t^d \leftarrow \max\{0, \delta_t^d \cdot k_{t-1} - (\bar{R}^{max} - x_{t-1})/\eta^c - e_{t-1}^d - e_{t-1}^l + e_{t-1}^c\}$, and $n_t^r \leftarrow \frac{n_t^u + n_t^d}{(\delta_t^u + \delta_t^d) \cdot k_{t-1}}$ {compute unserved regulation}
6: $x_t \leftarrow x_{t-1} + \eta^c \cdot (e_{t-1}^c + \delta_t^d \cdot k_{t-1} - n_t^d) - (e_{t-1}^d + e_{t-1}^l + \delta_t^u \cdot k_{t-1} - n_t^u)/\eta^d$ {update beginning hour-t storage level}
7: **end if**
8: $(a_t, y_t) \in \arg\max \left\{ \hat{F}_t^*(x_t, \omega_t, \tilde{m}_t^{(J)}(\tilde{\omega}_t)) \middle| a_t \in \mathcal{A}_{S_t} \right\}$ {find a near-optimal hour-t decision}
9: $C_t \leftarrow p_t^e \cdot (e_t^d - e_t^c) - V^L \cdot (D_t - l_t) + p_{t-1}^r \cdot k_{t-1} \cdot (1 - n_t^r) - V^{tr}(v_t)$ {compute hour-t profit contribution}
10: **end for**

3.3 Statistical Optimality Bounds

As noted before, the ASDP algorithm provides near-optimal policies for our original SDP. To demonstrate the quality of these policies, we also generate statistical upper and lower bounds on the true optimal value of the SDP. Shapiro (2003) shows that implementable and feasible policies provide valid statistical lower bounds, which we use. A policy is implementable and feasible if it satisfies the constraints of the original SDP and is non-anticipative, meaning that hour-t decisions depend solely on information available at hour-t. Our ASDP algorithm provides such feasible and implementable policies. This is because the hour-t decision that is determined in Step 8 of Algorithm 2 depends solely on the hour-t state variable. Moreover, the only exogenous variables that are unknown when making hour-t decisions that affect feasibility are the dispatch-to-contract ratios. Since regulation energy calls can go unserved with an associated profit penalty, the quantities, n_{t+1}^u, n_{t+1}^d, and n_{t+1}^r, are defined in Step 5 of Algorithm 2 to ensure that any hour-t decision is feasible. Thus, we compute our lower bound by randomly generating K sample paths, $\{\omega_t\}_{t=1}^T$, of the exogenous state variables and determining near-optimal policies using Algorithm 2. Let $G_0^{IF}(\omega^k)$ denote the resulting objective function value over the T-hour horizon with exogenous state variable sample path, ω^k. We then compute our lower bound as:

$$B_L = \frac{1}{K} \sum_{k=1}^{K} G_0^{IF}(\omega^k).$$

We compute statistical upper bounds using a sample path averaging technique. This is done by randomly generating K' sample paths, $\{\omega_t\}_{t=1}^T$, of the exogenous state variables. For each sample path we solve a deterministic optimization problem in which the full sequence of exogenous random variable values is known in hour 1 and profits are maximized over the T-hour hori-

zon. Let $G_0^{DET}(\omega^k)$ denote the resulting profit with exogenous state variable sample path, ω^k. We compute our upper bound as:

$$B_U = \frac{1}{K'} \sum_{k=1}^{K'} G_0^{DET}(\omega^k).$$

We compute 100 lower and upper bounds, using 1000 randomly generated sample paths for each. This provides sample means and standard errors of the bounds, which can be used to determine confidence intervals.

4 Case Study

We demonstrate our model and solution algorithm using a case study, in which a battery is co-located with a residential transformer. We allow the battery to provide all of the services modeled by our SDP. Our case study assumes a residential subdivision consisting of six houses connected to the grid through a 25 kVA transformer. We further assume that two of the homes have one hybrid PEV each and another two have one pure battery PEV each. The added PEV charging loads overload the transformer in some hours. We examine how a battery would be used in such a setting and also examine long-term economics of using a battery in this manner. Specifically, we allow the utility to address these distribution overloads (in addition to providing other services) using a mix of transformer upgrades and by installing a battery.

The region studied is in the PJM Interconnection system and we model the system based on PJM market rules. We further use PJM market and system data from 2009 to study battery use and economics over a typical year. This typical year is assumed to repeat in perpetuity for our long-term economic analysis. We optimize battery use over the year 24 hours at a time using a rolling 48-hour optimization horizon. This is done by first solving the ASDP for hours one through 48, to derive an optimal policy for days one and two. We then generate a sample path of random variable realizations for the first 24 hours and use the optimized policy to determine day-one decisions, objective function contributions, and the beginning state of the system on day two, S_{25}. We then roll forward and solve another ASDP for hours 25 through 71, using S_{25} as the starting state of the system. We proceed in this fashion to model battery use over the year studied.

Sioshansi et al (2009) argue that this use of an additional 24-hour 'look-ahead' period ensures that energy that has carryover value on a subsequent day is kept in storage. It is important to note that determining a proper length of the look-ahead period can depend on the battery's power and energy capacities. A battery that can store and discharge energy for many consecutive hours may require a longer look-ahead period to optimally carry energy over to subsequent periods. Our case study assumes a battery with a high power capacity relative to its energy capacity, implying that it can only charge and discharge for a limited number of hours. Thus, the 24-hour look-ahead period

should be sufficient. Since we optimize over a 48-hour period, we assume a discount factor $\gamma = 1$ in the ASDP. We now detail other assumptions underlying the case study.

4.1 Distribution-Level Load

Residential demand consists of PEV and non-PEV loads. We model the non-PEV loads based on two datasets provided by AEP for residential customers in its central-Ohio service territory in 2009. One data set provides summary statistics (*i.e.*, minimum, average, and coincident peak load) for a typical 25 kVA transformer connecting six homes to the grid in 2009. The second data set is an average hourly residential home load profile for 2009. We scale the hourly average loads, using generalized method of moments, to generate a load profile with summary statistics that match the values reported by AEP for a 25 kVA transformer connected to six homes.

PEV charging load data are generated based on the mass simulation model developed by Gong et al (2011). Their technique generates vehicle driving profiles using a Markov chain method and converts the driving profiles into battery energy usage. It then determines when vehicles are charged and the charging load assuming that the PEVs begin charging no earlier than 7 pm (but possibly later, depending on return time to the home) using a 4.4 kW level-2 charger.

Figure 3 shows three load duration curves corresponding to these data sets. The first corresponds to the 'unscaled' average hourly residential loads provided by AEP. The second are the loads scaled to match the summary statistics given by AEP. AEP reports that the minimum, average, and coincident peak loads for a typical 25 kVA transformer connected to six homes are 0.20 kW, 10.24 kW, and 25.97 kW, respectively. Our scaled load profile has values of 0.22 kW, 10.16 kW, and 26.00 kW for these summary statistics. The third load duration curve corresponds to the sum of the PEV and scaled non-PEV loads. Figure 3 further shows that the three load profiles exceed the 25 kVA transformer capacity in 11, 63, and 398 hours of the year, respectively.

Adding the PEV charging loads changes the distribution-level charging profile in two important ways. The first is that it increases the average and peak load on the distribution transformer to 11.37 kW and 41.47 kW, respectively. This effect is clearly seen in Figure 3. Secondly, the PEV charging changes the hours in which the distribution-level load peaks. Without PEVs, the distribution-level load peaks midday. Under our modeling assumption, the PEV charging load peaks in hours 19–21, when most PEVs begin their charging cycles. When the PEV charging loads are added, 99 of the 100 highest-load hours of the year occur between hours 19 and 21.

We assume in our case study that the PEV loads are deterministic and known, based on outputs of the mass simulation model (these are represented by the difference between the third and second load duration curves shown in Figure 3). The non-PEV loads are assumed to follow serially independent log-

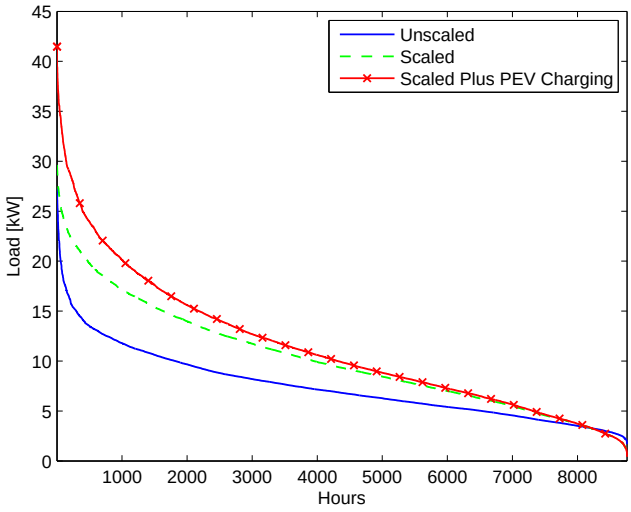

Fig. 3: Load duration curves for scaled and unscaled distribution-level loads

normal distributions. The mean of the distribution in each hour is given by the scaled non-PEV load modeled using the AEP data (these are represented by the second load duration curve shown in Figure 3) and we assume a variance of 0.15. We assume a penalty of $V^L = \$3.72$/kWh on any unserved residential load, which is the value reported by Kariuki and Allan (1996).

4.2 Battery Characteristics

The batteries' energy capacity depends on two factors: the amount of energy that has previously been cycled through the batteries and ambient temperature. We model the effects of both factors based on the results of battery cycling experiments conducted at The Ohio State University's Center for Automotive Research.[2] We assume that the batteries have a nominal capacity, $\bar{R}^{nom}$, which is reduced by 1.048×10^{-5} kWh per Ah of net energy exchanged (*i.e.*, charged or discharged) at 220 V. This is equivalent to 4.762×10^{-5} kWh of nominal energy capacity lost per kWh of net energy exchange.

Given the extremely low rate of battery capacity deterioration, our SDP does not directly account for this capacity loss in the optimization. Rather, after the operation of the battery over a one-week period is determined, the nominal battery capacity is updated before proceeding to model operations over the following week. Since the battery capacity is updated *a posteriori*, our SDP does not fully tradeoff the implicit cost of battery cycle life degradation

[2] Since this testing is still preliminary, results are not yet publicly available. The experiments cycle batteries repeatedly and under different temperature conditions to determine their capacities and aging characteristics.

against the value of the services that the battery provides. One could attempt to capture this tradeoff by adding a battery cycling cost to the objective function. This may be difficult to implement in practice, however, since the opportunity cost of battery degradation depends on how the battery would be optimally operated in the future.

We further assume that once the nominal battery capacity reaches some threshold level, the battery is no longer usable. This can either represent an actual battery failure, or the fact that once the nominal capacity is sufficiently low, it is more efficient to remove the battery from service and recycle it. We consider cases in which this failure occurs once the nominal capacity is less than between 50% and 75% of the starting nominal capacity. Figure 4 shows the deterioration of the battery capacity as a function of the number times it is cycled under the two different threshold assumptions. The figure defines a battery cycle as completely discharging and then completely recharging the battery. The figure also suggests that by examining the effect of the failure threshold on battery operations and economics, we can also indirectly examine the effect of different battery deterioration rates. This is because varying the failure threshold affects the number of cycles that the battery can operate through.

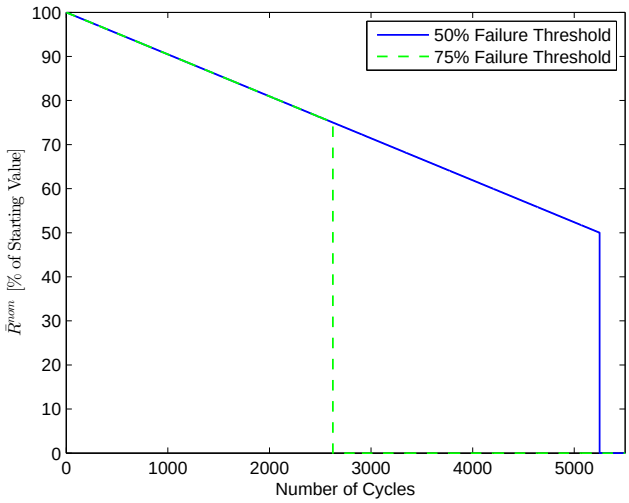

Fig. 4: Nominal capacity deterioration as a function of battery cycling

We capture temperature effects by modeling the batteries' minimum and maximum energy capacities, $\bar{R}^{min}$ and $\bar{R}^{max}$, as being a fraction of $\bar{R}^{nom}$, where this fraction is temperature-dependent. Figure 5 shows weekly-average temperature data for the city of Columbus, OH from the year 2009, which are used to model temperature effects (daily-average temperatures are used in the modeling, but Figure 5 shows the weekly-averaged data). It also shows the

modeled minimum and maximum capacity for each week (as a percentage of $\bar{R}^{nom}$). Figure 5 further indicates that the batteries' charging and discharging efficiencies are temperature-dependent.

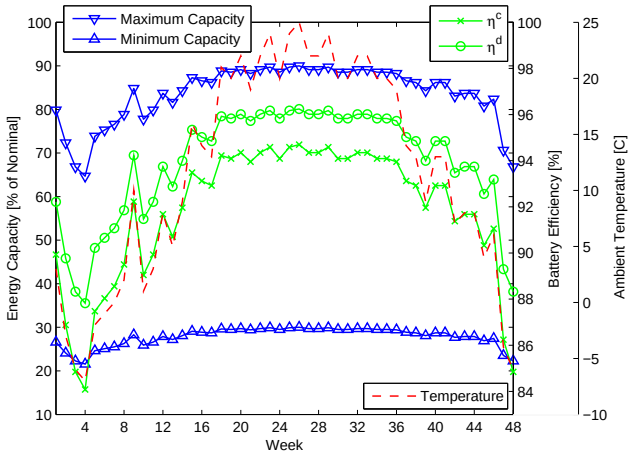

Fig. 5: Weekly average temperatures and resulting battery capacities and efficiencies

We assume that future battery capacities and efficiencies are known with perfect foresight and deterministic within the 48-hour optimization horizon. We further assume that the batteries have a maximum charging and discharging power capacity, $\bar{P}^b$, which is equal to 8.5 kW per kWh of nominal energy capacity, based on tested characteristics of lithium ion batteries.

4.3 Transformer Characteristics

A transformer can be operated above its rated capacity, however this accelerates transformer aging. Susa et al (2005) and Gong et al (2011) develop a transformer aging model based on hot-spot and top-oil temperatures, which we use to estimate the effect of operating a transformer above its rated capacity. This is done by increasing transformer loading during a single hour, while holding the remaining loads constant, and modeling the expected decrease in transformer life. We then use an assumed lump-sum replacement cost of $750 for a 25 kVA transformer, based on typical values reported by AEP, to translate this expected transformer life loss into a cost. Since transformers are a relatively mature technology, we do not expect or model substantive future cost reductions. The resulting cost of operating the 25 kVA transformer above

its rated capacity is given by the following convex piecewise-linear function:

$$V^{tr}(v_t) = \begin{cases} 0.68v_t, & \text{if } v_t \in [0, 0.1\bar{P}^{tr}), \\ 1.76(v_t - 0.1\bar{P}^{tr}) + 0.68 \cdot 0.1\bar{P}^{tr}, & \text{if } v_t \in [0.1\bar{P}^{tr}, 0.2\bar{P}^{tr}), \\ 5.14(v_t - 0.2\bar{P}^{tr}) + 2.44 \cdot 0.1\bar{P}^{tr}, & \text{if } v_t \in [0.2\bar{P}^{tr}, 0.3\bar{P}^{tr}), \\ 16.30(v_t - 0.3\bar{P}^{tr}) + 7.58 \cdot 0.1\bar{P}^{tr}, & \text{if } v_t \in [0.3\bar{P}^{tr}, 0.4\bar{P}^{tr}), \\ 50.21(v_t - 0.4\bar{P}^{tr}) + 23.88 \cdot 0.1\bar{P}^{tr}, & \text{if } v_t \in [0.4\bar{P}^{tr}, 0.5\bar{P}^{tr}), \\ 158.40(v_t - 0.5\bar{P}^{tr}) + 74.09 \cdot 0.1\bar{P}^{tr}, & \text{if } v_t \in [0.5\bar{P}^{tr}, 0.6\bar{P}^{tr}), \\ 520.70(v_t - 0.6\bar{P}^{tr}) + 232.49 \cdot 0.1\bar{P}^{tr}, & \text{if } v_t \in [0.6\bar{P}^{tr}, +\infty). \end{cases} \tag{21}$$

4.4 Exogenous Random Variables

Although our model and solution algorithm do not require any specific correlation among the exogenous random variables (other than the Markovian property), we assume in our case study that $\hat{p}_t^e$, $\hat{p}_t^r$, $\hat{D}_t$, $\hat{I}_t$, $\hat{\delta}_t^u$, and $\hat{\delta}_t^d$ are all mutually independent. We further assume that $\hat{p}_t^e$, $\hat{p}_t^r$, $\hat{D}_t$, $\hat{\delta}_t^u$, and $\hat{\delta}_t^d$ are serially independent over time.

Historical energy and regulation capacity prices show very little correlation. In 2009 these prices had a correlation of about -0.15. This is because high energy prices signal less generating capacity being available and higher-cost generation having to be used to serve the load. High regulation prices signal a lack of fast-responding generation. Indeed, energy prices tend to peak midday when electricity demand is the highest, whereas regulation prices can peak overnight when loads are low and only base-load generators with slow ramping rates are online. These differences in the diurnal price patterns also explain the slightly negative correlation between the two sets of prices. Knittel and Roberts (2005) and Shrestha and Songbo (2010) test historical energy price data from the California ISO and Singapore electricity markets, respectively, and show that these best fit a log-normal distribution. Our own examination of PJM system marginal prices in 2009 reveals the same. Thus, we assume that the energy prices have a log-normal distribution. Although the distribution of regulation prices has not been examined, our analysis of 2009 PJM data indicates that these prices also fit a log-normal distribution, which we assume.

To capture seasonal and diurnal price patterns, we allow the price distributions to differ for each of the 8760 hours of the year. Specifically, the location parameter of the price distribution in each hour is set equal to the actual historical energy or regulation price. The scale parameters are set equal to the maximum likelihood estimator given by all of the prices observed for the same hour of the day during the course of the month. That is, a different scale parameter is estimated for each of 24 hours in the month of January, as is done for the 24 hours of each of other months of the year.

The distributions of the dispatch-to-contract ratios of regulation up and down are estimated using historical PJM data from 2009. These data specify the amount of regulation capacity reserved in each hour and the amount of

gross regulation up and down energy deployed in real-time. The data do not show any diurnal or seasonal patterns in the ratios. Thus, we assume that the distributions are time-invariant. Hypothesis testing suggests that a Gaussian distribution best fits the historical data, which we assume. Maximum-likelihood estimators of the mean and standard deviation are used.

A number of approaches are used to model power system reliability, with Markov-based models being common. DeSieno and Stine (1965) show that power system failures approximately follow a Markov process. We model system outages using a two-state Markov chain in which the system can either be in an outage or non-outage state. Transitions between these states depend solely on the present state of the system, and the transition probabilities are time-invariant. We use system reliability data reported by FirstEnergy for its Ohio service territory in 2009 to estimate the transition probabilities. FirstEnergy reports an average of 1.24 outages per customer during the year, which lasted an average of 2 hours. Based on these values, we assume a 0.000142 probability that the system has an outage at hour $t+1$ if it is in a non-outage state at hour t, and a 0.5 probability that it recovers from an outage in each hour.

The low probability with which it occurs raises an issue in capturing the effect of system outages on battery usage. If outages are not observed in the sample paths generated in Step 3 of Algorithm 1, our estimates of $\hat{F}_t^*$ may underestimate the benefit of keeping energy stored for future outages. To overcome this issue, we modify Algorithm 1. Specifically, after randomly generating a sample path and updating the estimated slopes of $\hat{F}_t^*$ in Steps 5–18, we repeat Steps 5–18 by fixing the outage variable equal to 1 in each hour. In this way, the $\hat{F}_t^*$ functions are explicitly estimated taking into account the cost savings of having stored energy available to help mitigate an outage. Our optimized policies using Algorithm 1 (without this modification) and this variant in which outages are explicitly included in the function updates are quite similar to each other. This suggests that although there is a high cost of load curtailment, the low probability with which outages occur and the value of other services that the battery can be used for overshadows the benefits of keeping additional energy stored. These results are further discussed in Section 5.

4.5 Model Discretization

The pre- and post-decision storage level state variables are discretized to take on values:

$$\tilde{x}_t, \tilde{x}_t^{post} \in \{\bar{R}^{min}, \bar{R}^{min} + 0.5, \bar{R}^{min} + 1.0, \cdots, \bar{R}^{max}\}.$$

This implies that the charging and discharging variables, $\tilde{e}_t^c$, $\tilde{e}_t^d$, and $\tilde{e}_t^l$, can only take on a discrete number of values, corresponding to possible transitions between $\tilde{x}_t$ and $\tilde{x}_t^{post}$. We discretize the k_t variables using 1 kW interval widths between their upper and lower bounds, which are defined by equations (4)

and (5). We discretize the v_t variables using the breakpoints shown in equation (21). The distributions of the $\hat{p}_t^e$, $\hat{\delta}_t^u$, $\hat{D}_t$, and $\hat{\delta}_t^d$ random variables are discretized into five possible outcomes and the distributions of the $\hat{p}_t^r$ random variables into four possible values, using bracket medians. The distribution of $\hat{I}_t$ needs no discretization, as it can only take on two values.

These assumptions yield a discrete dynamic program, given by (18), which can be solved using the dynamic programming algorithm. Moreover, we can exploit the structure of our problem to further reduce the feasible action space, over which we must search for an optimal solution at Step 6 of Algorithm 2. We first note that due to roundtrip efficiency losses, it is suboptimal to simultaneously charge and discharge energy. Thus, for all t, e_t^c cannot be non-zero if at least one of e_t^d and e_t^l is and *vice versa*. Hence, the total number of combinations that e_t^c, e_t^d, and e_t^l can take in a given hour is M^2, where:

$$M = \left|\{\bar{R}^{min}, \bar{R}^{min} + 0.5, \bar{R}^{min} + 1.0, \cdots, \bar{R}^{max}\}\right|.$$

Furthermore, due to the high penalty on unserved building loads, the value of $\tilde{l}_t$ can be determined by (7) through (10) and the values of I_t, $\tilde{e}_t^c$, $\tilde{e}_t^d$, $\tilde{e}_t^l$, $\tilde{k}_t$, and $\tilde{v}_t$. Specifically, if $I_t = 1$, then $\tilde{l}_t = \min\{\tilde{D}_t, \tilde{e}_t^l\}$. Otherwise, if $I_t = 0$ then $\tilde{l}_t = \min\{\tilde{D}_t, \bar{P}^{tr} + \tilde{v}_t + \tilde{e}_t^l\}$. This structure implies that a maximum of $M^2 \cdot |\tilde{v}_t| \cdot \bar{P}^{tr}$ combinations of action variables are feasible and could be optimal in (18).

4.6 ASDP Algorithm Implementation

Nascimento and Powell (2009) prove that the piecewise-linear approximations of the F_t^{post} functions, which are estimated in Algorithm 1, converge to the true cost-to-go function as the number of iterations approaches infinity. We conduct 500 iterations of this algorithm, since for most of the case studies run, the approximation is within a 1% relative error after about 100 iterations. The values of α_j used in Step 11 of Algorithm 1 are taken to be 1/2 for the first 50 iterations, 1/3 for the next 50, 1/4 for the following 50, *etc.*

5 Results

We first model and study the operation of an exogenously fixed battery and transformer deployment over a typical one-year period. We then use the results of this typical year to conduct a more thorough long-term economic analysis, in which the utility has the option of installing batteries and transformer upgrades to maximize life-cycle value of the deployment.

5.1 Short-Term Operational Modeling Results

Figure 6 shows the operation of a battery with a nominal capacity of $\bar{R}^{nom} =$ 15 kWh over a one-day period. The figure shows that much of the value of the

battery comes from regulation, and that it provides much more regulation than arbitrage. This is because regulation is primarily a capacity service resulting in relatively little energy charging or discharging. This means that this service tends to incur little cost and comparably high revenues. Although the dispatch-to-contract ratio in a particular hour can be high (*e.g.*, we find cases of up to 0.35 in the historical PJM data) the regulation-up and -down signals tend to cancel-out in the long-run. Our simulation has high ratios of up to 0.36, but the average ratio over the year is much lower at 0.10, which is consistent with the historical PJM data. Thus, on average, providing regulation results in small net charging of the battery. When the energy price is sufficiently high compared to the price of regulation, such as in hour 15, the battery provides arbitrage energy as well.

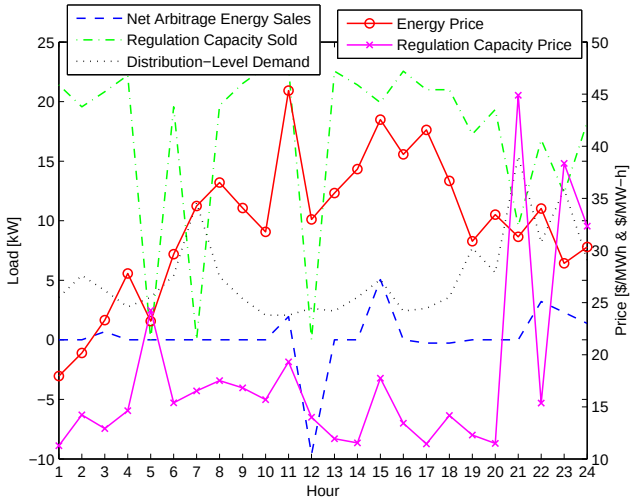

Fig. 6: Battery operation over a one-day period

Figure 6 also illustrates that determining an optimal policy is more complex than following diurnal price patterns. For instance, the battery is charged in hour 12, despite the energy price being relatively high at that time. This is because a high dispatch-to-contract ratio in hour 10 results in 3.5 kWh of energy discharge. This, coupled with energy sales due to the high energy price in hour 11, nearly depletes the battery's state of charge. The battery is recharged in hour 12 so it can provide more regulation capacity in subsequent hours while reducing the likelihood of not receiving 100% of the regulation payment due to an unserved energy call. Table 1 reports the mean upper- and lower-bounds on the optimal SDP objective function value and standard errors for the bounds over the typical year studied. The optimality gap is around 1.97%, showing that the ASDP algorithm provides a relatively good near-optimal solution.

Table 1: Upper and lower bounds on optimal SDP objective function value

Bound	Value [$]	Standard Error [$]
B_L	2172.7	18.74
B_U	2215.4	22.39

Another benefit of the battery, which is illustrated in Figure 7, is that it improves the net distribution-level load profile. This is because the battery allows stored energy to be used during extremely high-load periods while also increasing transformer loading during low-load periods when the transformer is not being fully utilized. Adding the 15 kWh battery reduces the peak net distribution-load from 41.47 kWh (without the battery) to 35.45 kWh. Moreover, adding the battery reduces the net distribution-level load by an average of 12.3 kW during the 10 highest-load hours of the year. The battery increases the average (over the course of the year) distribution-level load from 11.37 kW to 11.41 kW, which reflects the effect of energy lost in the storage process. The battery also reduces net distribution-level demand variability slightly. Without the battery the hourly loads have a standard deviation of about 6.76 as opposed to 6.23 with the battery.

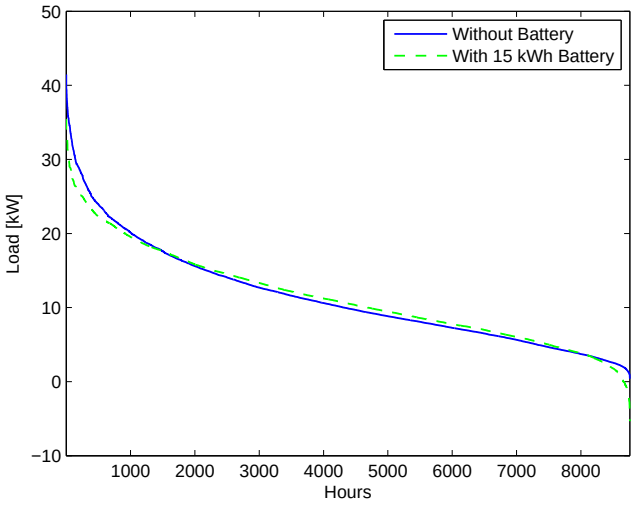

Fig. 7: Net distribution-level load duration curve without and with 15 kWh battery

5.2 Long-Term Distribution Infrastructure Design

Optimizing distribution infrastructure design involves an economic tradeoff between the upfront transformer upgrade and battery installation capital costs and the associated stream of revenues and averted costs. These are combined with an assumed annual discount rate to compute the net present value (NPV) of different infrastructure designs and to determine an NPV-maximizing deployment. We compute these NPVs over a 20-year period, which is a standard design life of a distribution-level transformer. Before presenting the results of this analysis, we detail the cost and revenue assumptions underlying it.

5.2.1 Cost Assumptions

Our cost assumptions are based on typical values reported to us by AEP. The transformer is assumed to be available in 25 kVA, 37.5 kVA, 50 kVA, and 75 kVA sizes only, with lump-sum per-unit capital costs of \$750, \$1000, \$1300, and \$1750, respectively. Although the transformer has a 20-year design life, it may fail prematurely depending on loading characteristics. If this occurs, we assume that a replacement transformer of equal capacity is immediately installed, with the same capital cost. As noted before, since transformers are a relatively mature technology, we do not model substantive future cost reductions. Since distributed battery storage devices are not widely available, the cost of such systems is uncertain. Thus, we consider two bounding cases in which such batteries cost \$50/kWh and \$200/kWh. We assume that all battery and transformer installations incur annual maintenance costs equal to 5% of their upfront capital costs. We further assume a 7% annual discount rate.

5.2.2 Revenue Assumptions

Deployment revenues are estimated using the SDP model. This is done by running the SDP in the manner outlined in Section 4 and repeating the process using 1000 sample paths of the exogenous random variables to estimate expected annual revenues. Figure 8 shows annual expected battery operating revenues, as a function of the starting nominal battery capacity at the beginning of the year and the transformer capacity. The revenues shown account for battery cycle-life losses during the course of the year, since the starting nominal battery capacity is updated after each week of simulated operations. As expected, it shows that the battery value is increasing in its capacity, although this can be greatly limited by transformer capacity as well. This reflects the transformer potentially limiting the ability of the battery to provide market services, especially regulation.

While using the battery accrues the revenues shown in Figure 8, it also causes cycle-life losses to the battery. Figure 9 shows these cycle-life losses, as a function of transformer and starting nominal battery capacities. As expected,

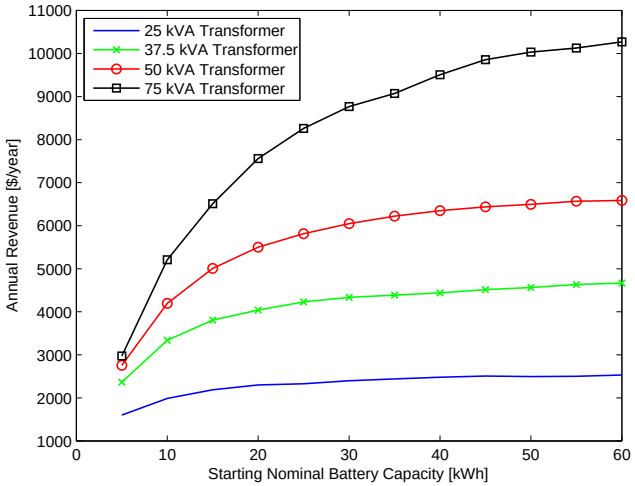

Fig. 8: Expected annual battery operating revenues

a larger battery or transformer gives greater losses, since such deployments allow greater operational flexibility and more energy exchange to occur.

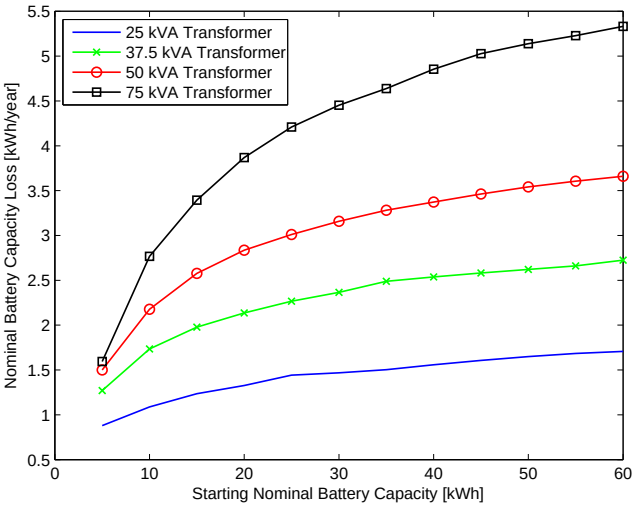

Fig. 9: Expected annual battery capacity loss

To compute the expected NPV of revenues earned, we define $\upsilon(\bar{P}^{tr}, \bar{R}^{nom})$ as expected revenues earned by the deployment during the modeled typical year, as a function of the transformer and starting (at the beginning of the year) nominal battery capacities. We also define $\lambda(\bar{P}^{tr}, \bar{R}^{nom})$ as expected

nominal battery capacity loss after the modeled typical year, as a function of the transformer and starting nominal battery capacities. Both of these functions are found from Figures 8 and 9, using linear interpolation. We impose the battery failure condition (discussed in Section 4.2) when defining $\upsilon(\bar{P}^{tr}, \bar{R}^{nom})$. This is done by fixing $\upsilon(\bar{P}^{tr}, \bar{R}^{nom}) = \upsilon(\bar{P}^{tr}, 0)$ for any $\bar{R}^{nom}$ below the cutoff value, at which the battery is assumed to fail. We also define γ^{ann} as the annual discount rate.

5.2.3 Discounted Deployment Profits

The NPV of revenues earned by a deployment consisting of a $\bar{P}^{tr}$ kVA transformer and a battery with a starting (at the beginning of year 0) nominal capacity of $\bar{R}_0^{nom}$ kWh is given by:

$$\upsilon^{NPV}(\bar{P}^{tr}, \bar{R}_0^{nom}) = \sum_{y=0}^{19} \gamma_{ann}^{y} \cdot \upsilon(\bar{P}^{tr}, \bar{R}_y^{nom}),$$

where the starting nominal battery capacity in each year is updated according to:

$$\bar{R}_{y+1}^{nom} = \bar{R}_y^{nom} - \lambda(\bar{P}^{tr}, \bar{R}_y^{nom}), \ \forall\ y = 1, 2, \ldots, 18.$$

The net discounted profit of the deployment is then given by:

$$\upsilon^{NPV}(\bar{P}^{tr}, \bar{R}_0^{nom}) - K^{tr}(\bar{P}^{tr}) - K^{ba}(\bar{R}^{nom}) - \sum_{y=1}^{r} \gamma_{ann}^{y \cdot L(\bar{P}^{tr})} K^{tr}(\bar{P}^{tr}), \quad (22)$$

where:

$$r = \left\lfloor \frac{20}{L(\bar{P}^{tr})} \right\rfloor,$$

is the number of times that the transformer must be replaced during the 20-year optimization period, $K^{tr}(\bar{P}^{tr})$ and $K^{ba}(\bar{R}^{nom})$ are the transformer and battery capital costs, respectively, as a function of their starting nominal capacities, and $L(\bar{P}^{tr})$ is the expected transformer failure time in years.

Figure 10 shows the estimated NPV, as defined by (22), of different deployment configurations over the 20-year study horizon. The figures assume different underlying battery costs and failure characteristics. Contrasting the figures shows that while battery capital cost plays a role in determining the net return and economic viability of a distributed battery storage system, its failure characteristic is a much more important parameter. This is indicative of the fact that with the assumed 7% annual discount rate, the ability of a battery to earn revenues over a prolonged period of time is important in making such an investment viable. These results, further, suggest that the principal focus of battery manufacturers targeting distributed storage solutions should be on maximizing battery cycle-life, with manufacturing cost a less important factor.

All of the configurations without a battery result in a negative NPV. This is because such deployments incur the cost of installing a transformer, without

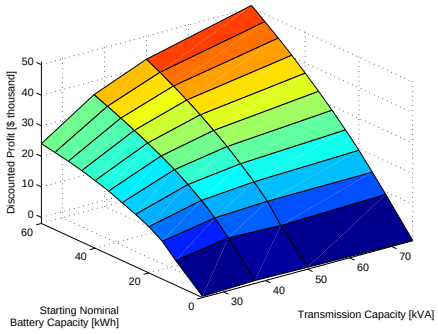

(a) $50/kWh battery cost and battery failure at 50% of starting nominal capacity

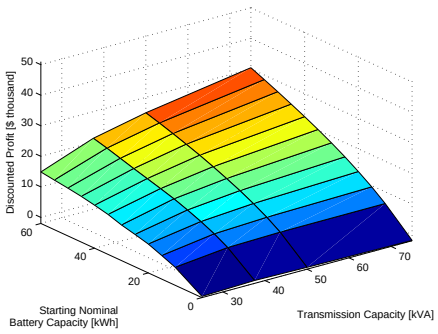

(b) $50/kWh battery cost and battery failure at 75% of starting nominal capacity

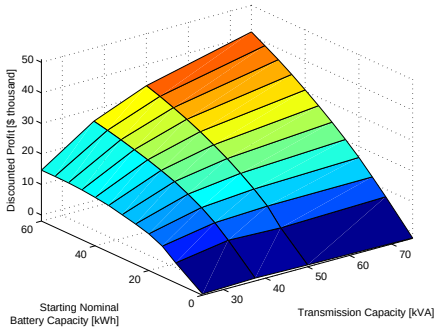

(c) $200/kWh battery cost and battery failure at 50% of starting nominal capacity

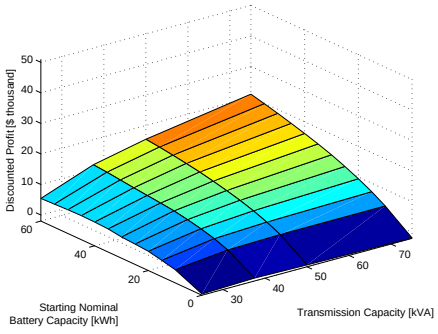

(d) $200/kWh battery cost and battery failure at 75% of starting nominal capacity

Fig. 10: Expected NPV of deployment profits with different battery costs and failure points

earning any revenues through providing market services with the battery. The highest-NPV configuration without a battery includes a 37.5 kVA transformer. Installing a 25 kVA transformer requires two transformer replacements during the 20-year planning horizon. A 37.5 kVA transformer, on the other hand, does not require any replacement within the 20-year horizon. If a 25 kVA transformer is installed, a battery with a 10 kWh starting nominal capacity is needed to maintain a 20-year transformer life.[3] None of the other deployments examined require transformer replacements within the 20-year planning horizon considered.

[3] We only model battery capacities with starting energy capacities in 5 kWh increments. A deployment with a 25 kVA transformer and a battery with a starting nominal capacity of 5 kWh results in an expected transformer life of 18.4 years. A 10 kWh battery results in a 21.3 year transformer life.

6 Conclusions

This paper presents an SDP model that co-optimizes the use of a battery for multiple applications. This includes energy arbitrage, the provision of AS, backup energy for distribution-level loads, and distribution infrastructure relief. Our SDP accounts for both market and system uncertainty, and is fairly flexible in how such uncertainties are modeled. We also adapt an approximation technique developed by Nascimento and Powell (2009) to solve the SDP efficiently. We use a case study, based on central Ohio in the year 2009, to demonstrate the model. We show the inherent complexity in co-optimizing the use of a storage device for multiple applications. We also use the results of our short-run operational model to study long-term distribution infrastructure design. Our results show that the aging characteristic of the battery and the point at which it fails are critically important parameters in determining the long-run financial viability of a distributed energy storage system. Although we focus on small-scale distributed storage, the model is also applicable to utility-scale storage.

Acknowledgements The authors thank A. Sorooshian, the editors, and two anonymous reviewers for helpful discussions and suggestions and Y. Guezennec for providing useful battery aging and temperature performance data. Financial support for this work was provided by the SMART@CAR consortium. This work was also supported in part by an allocation of computing time from the Ohio Supercomputer Center. Any opinions and conclusions expressed in this paper are solely those of the authors.

References

Collins MM, Mader GH (1983) The Timing of EV Recharging and Its Effect on Utilities. IEEE Transactions on Vehicular Technology VT-32:90–97

DeSieno CF, Stine LL (1965) A Probability Method for Determining the Reliability of Electric Power Systems. IEEE Transactions on Reliability R-14:30–35

Gong Q, Midlam-Mohler S, Marano V, Rizzoni G (2011) PEV Charging Impact on Residential Distribution Transformer Life. In: 2011 IEEE Energytech, Institute of Electrical and Electronics Engineers, Cleveland, OH, United States

Graves F, Jenkin T, Murphy D (1999) Opportunities for Electricity Storage in Deregulating Markets. The Electricity Journal 12:46–56

Kariuki KK, Allan RN (1996) Evaluation of reliability worth and value of lost load. IEE Proceedings—Generation, Transmission, and Distribution 143:171–180

Kempton W, Tomić J (2005) Vehicle-to-grid power fundamentals: Calculating capacity and net revenue. Journal of Power Sources 144:268–279

Kintner-Meyer MCW, Schneider KP, Pratt RG (2007) Impacts Assessment of Plug-in Hybrid Vehicles on Electric Utilities and Regional US Power Grids: Part 1: Technical Analysis. Online Journal of EUEC 1(4)

Knittel CR, Roberts MR (2005) An empirical examination of restructured electricity prices. Energy Economics 27:791–817

Mohseni P, Stevie RG (2009) Electric vehicles: Holy grail or Fool's gold. In: Power & Energy Society General Meeting, 2009, Institute of Electrical and Electronics Engineers, Calgary, AB, pp 1–5

Nascimento JM, Powell WB (2009) An Optimal Approximate Dynamic Programming Algorithm for the Energy Dispatch Problem with Grid-Level Storage, working paper

Nourai A (2007) Installation of the First Distributed Energy Storage System (DESS) at American Electric Power (AEP). Tech. Rep. SAND2007-3580, Sandia National Laboratories

Powell WB (2007) Approximate Dynamic Programming: Solving the Curses of Dimensionality. Wiley-Interscience, Hoboken, New Jersey

Rockafellar RT (1970) Convex Analysis. Princeton University Press, Princeton, New Jersey

Shapiro A (2003) Inference of statistical bounds for multistage stochastic programming problems. Mathematical Methods of Operations Research 58:57–68

Shrestha GB, Songbo Q (2010) Statistical Characterization of Electricity Price in Competitive Power Markets. In: 2010 IEEE 11th International Conference on Probabilistic Methods Applied to Power Systems (PMAPS), Institute of Electrical and Electronics Engineers, Singapore

Sioshansi R, Denholm P (2010) The value of plug-in hybrid electric vehicles as grid resources. The Energy Journal 31:1–23

Sioshansi R, Denholm P, Jenkin T, Weiss J (2009) Estimating the Value of Electricity Storage in PJM: Arbitrage and Some Welfare Effects. Energy Economics 31:269–277

Susa D, Lehtonen M, Nordman H (2005) Dynamic Thermal Modelling of Power Transformers. IEEE Transactions on Power Delivery 20:197–204

Tomić J, Kempton W (2007) Using fleets of electric-drive vehicles for grid support. Journal of Power Sources 168:459–468

Xi X, Sioshansi R, Marano V (2014) A Stochastic Dynamic Programming Model for Co-optimization of Distributed Energy Storage. Energy Systems Forthcoming