A Partially Observable Markov Decision Process Approach to Residential Home Energy Management

Timothy M. Hansen, *Member, IEEE*, Edwin K. P. Chong, *Fellow, IEEE*, Siddharth Suryanarayanan, *Senior Member, IEEE*, Anthony A. Maciejewski, *Fellow, IEEE*, and Howard Jay Siegel, *Fellow, IEEE*

***Abstract*—Real-time pricing (RTP) is a utility-offered dynamic pricing program to incentivize customers to make changes in their energy usage. A home energy management system (HEMS) automates the energy usage in a smart home in response to utility pricing signals. We present three new HEMS techniques—one myopic approach and two non-myopic partially observable Markov decision process (POMDP) approaches—for minimizing the household electricity bill in such a RTP market. In a simulation study, we compare the performance of the new HEMS methods with a mathematical lower bound and the status quo. We show that the non-myopic POMDP approach can provide a 10%–30% saving over the status quo.**

***Index Terms*—Cyber-physical systems, demand response, home energy management systems, partially observable Markov decision process, real-time pricing.**

I. Introduction

THE UNITED States Energy Information Administration predicts a 21% increase in residential electricity use from a 2012 reference case to the year 2040 [1]. Studies show that small and targeted reductions in peak demand can have large impacts on wholesale electricity prices [2]. Given that residential customers can account for over half of the system peak demand in summertime, such as in markets like the Electric Reliability Council of Texas (ERCOT) [3], residential demand response (DR) programs are attractive solutions for relieving the stress on the system and market.

Dynamic pricing programs are one way to accomplish DR. These utility-offered programs, such as time-of-use (TOU) and real-time pricing (RTP), fluctuate the price of electricity throughout the day in accordance with system load levels to elicit a change in the consumption of electricity [4]. Residential customers can take advantage of these time-varying rates by changing electricity use to reduce their electricity bill. An automated method for changing electricity usage in response to time-varying price is a residential home energy management system (HEMS), a form of demand-side management (DSM). The challenges of an effective HEMS are (a) the uncertainty in the time-varying price of electricity, and (b) that as a customer, the benefit received from changing energy usage must exceed the inconvenience caused. To overcome these challenges and to maximize the benefit of dynamic pricing, we design a HEMS using a non-myopic sequential decision technique known as a partially observable Markov decision process (POMDP). The POMDP HEMS determines energy use at each point in time to minimize the electricity bill under the uncertainty of the time-varying price and customer comfort constraints.

HEMS and DSM are active research areas [5]–[17]. HEM has been approached through dynamic programming [8], stochastic optimization [9], [14], two-horizon algorithms [16], MILP [9], [10], convex programming with integer relaxation [12], and heuristic optimization [15], [17]. The HEMS optimize for cost [7], [9], [10], [12], [15], [16], user preference [7], [11], power caps [11], [14], and peak-to-average ratio [15]. Our POMDP HEMS is a stochastic control process for making non-myopic decisions when the underlying state is uncertain. The resulting delayed gratification approach to HEMS, combined with a continuously-updating RTP prediction method, results in significant cost savings in a RTP market. Our work also presents a new method for accurately modeling appliance usage within a household. The considered time-period of one-month for quantifying results is also significantly longer than previous studies.

Markov decision processes (MDPs) have been used in other problems relating to power systems. In [18], an MDP approach for commercial building energy management is presented. Multiple energy systems (e.g., wind, photovoltaics (PV), batteries) are jointly scheduled to match the load requirements of a commercial building to minimize cost. An MDP approach is used in [19] to schedule the use of residential pool pumps. A HEM unit "plays a game" with a central energy management unit that provides a dynamic price to meet its energy needs within a desired budget in [20]. Our work differs from previous power system MDP papers in that we use the *partially observable* framework and we make different decisions within our POMDP framework only at the customer level in an existing RTP market. The POMDP HEMS presented here manages the use of flexible, non-interruptible smart appliances according to the definition in [9].

Manuscript received July 31, 2015; revised January 22, 2016 and April 28, 2016; accepted June 10, 2016. Date of publication June 21, 2016; date of current version February 16, 2018. Paper no. TSG-00890-2015.

T. M. Hansen is with South Dakota State University, Brookings, SD 57007 USA (e-mail: timothy.hansen@sdstate.edu).

E. K. P. Chong, S. Suryanarayanan, A. A. Maciejewski, and H. J. Siegel are with Colorado State University, Fort Collins, CO 80523 USA (e-mail: edwin.chong@colostate.edu; sid@colostate.edu; aam@colostate.edu; hj@colostate.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TSG.2016.2582701

Our primary contributions in this work are:

(a) the design of a non-myopic residential HEMS using a sequential decision technique (i.e., POMDP) that optimizes energy usage over a long time horizon to minimize cost;
(b) the creation of a new appliance energy usage pattern based on queueing theory that models residential household usage for smart home simulations; and
(c) the comparison of the POMDP HEMS against three methods, including a new myopic algorithm, using month-long simulation studies.

The system model and problem statement are described in Section II. Section III introduces the design of the optimization methods. In Sections IV and V, the setup and results of the simulation study are presented. Concluding remarks are given in Section VI.

II. System Model

A. Overview

This work relates to changing the energy usage within a single residential household in response to dynamic pricing of electricity. The house exists in a RTP market, where the price for electricity varies every hour in response to demand. If many homes change their load profile using HEMS the RTP may change, however in this work we consider the HEMS as a price-taker [21] — an entity whose transactions are unable to affect the market price. Flexible, non-interruptible smart appliances [9] are dynamically *arriving* (i.e., the residential customer is ready to use the appliance — not that the appliance is physically arriving) to be scheduled by the HEMS. The goal of the HEMS is to minimize the total cost of electricity. We assume that the enabling communication and control technologies are available and installed [22]. We are aware that cyber-security (e.g., [23]) is a concern moving forward in the implementation of smart homes — and, more generally, the Internet-of-Things (IoT) — but it is out of the immediate scope of this work.

B. Appliance Model

For each appliance i that arrives at time $t_{i-\text{arr}}$, we are given the power rating in kW (p_i), the duration in hours (d_i), and the *start-time deadline* ($t_{i-\text{dead}}$). To consider customer comfort, the customer provides the start-time deadline for each appliance, indicating the latest time the appliance can be started. This allows the customer to be more (later start-time deadline) or less (sooner start-time deadline) flexible depending on the specific appliance and the customer comfort. The goal of the HEMS is to find the start-time of each appliance ($t_{i-\text{start}}$).

We consider a generic appliance model with constant power draw. An example of the difference between an actual appliance output and our generic model is shown in Fig. 1. In Fig. 1(a), the output of the generic constant-power draw model is compared to a clothes washer with time-varying power consumption from [24] at one-minute resolution. The generic model has a load equal to the average load of the actual appliance (i.e., they consume the same amount of energy, just at different rates throughout the appliance duration). To validate this approximation, we compare the energy consumption of the two models through the appliance duration and the error between the two models in Figs. 1(b) and 1(c), respectively. Although the two models rarely consume the same power at the same time, the total energy consumed is the same, the maximum absolute error is under 6% (maximum absolute error is 15 Wh of 254 Wh consumed by the device), and at four points in time the two models consumed the same amount of energy (the zero crossings in Fig. 1(c)).

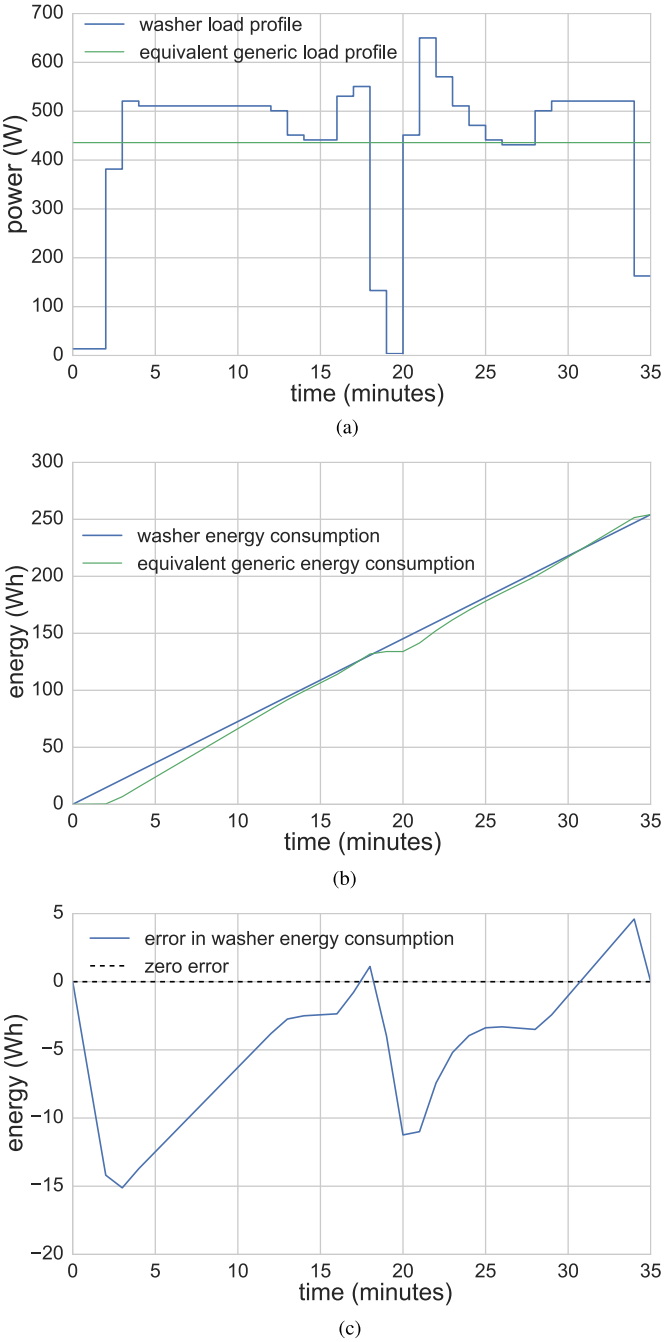

Fig. 1. (a) The load profile of the washer from [24] (blue line) versus the equivalent constant average load profile at a one-minute resolution; (b) the energy consumption of the real (blue line) versus equivalent model (green line) through time; and (c) the error between the two methods through the appliance duration (blue line) compared to zero error (black dashed line).

Therefore, for simulation purposes, the generic constant load model is considered adequate.

For simulation, we must determine $t_{i-\text{arr}}$ for each appliance that models residential energy usage. We introduce a novel probabilistic model for the usage pattern of a residential household based on queueing theory, specifically an $M_t/G/\infty$ queue. We posit that a household can be modeled as an infinite computing server with appliance usage analogous to task arrivals — where the capacity of a server is equivalent to the number of tasks it can execute, and the infinite capacity of the household indicates there is no limit on the number of concurrently running appliances. In the $M_t/G/\infty$ queue, applications arrive non-homogeneously with *Markovian* probability (i.e., time-varying Poisson process), *generally* distributed execution times, and *infinite* capacity [25]. In the smart home realm, the run-time of an appliance is analogous to the execution time of a task, and the resident wanting to turn on an appliance is analogous to the arrival of a compute task. The capacity of the household is *infinite* in the sense that if a resident wants to turn on one of their appliances, they are not delayed (or "queued") by the capacity of the household.

Let D be a random variable describing the duration of the set of appliances and $m(t)$ be the average number of appliances running at time t. The appliances arrive into the system according to a Poisson process with the time-varying rate $\lambda(t)$. Using the linear-with-time-shift (LIN-S) approximation of $\lambda(t)$ from [25], the time-varying rate of the Poisson process describing the appliance arrivals at time t is given as

$$\lambda(t - \mathbb{E}[D]) = {}^{m(t)}/_{\mathbb{E}[D]}. \tag{1}$$

Because the arrivals are for simulation, we do not have to assume causality. If we substitute the current time as $t \implies t + \mathbb{E}[D]$, then we have the required rate equation as

$$\lambda(t) = {}^{m(t+\mathbb{E}[D])}/_{\mathbb{E}[D]}. \tag{2}$$

To obtain a realistic number of simultaneously running appliances, we must determine $m(t)$. Let L be a random variable describing the power rating of the set of appliances and $l(t)$ be the desired aggregate household load at time t. The average number of running appliances is then

$$m(t) = {}^{l(t)}/_{\mathbb{E}[L]}. \tag{3}$$

Substituting (3) into (2), we determine the time-varying arrival rate of the appliances according to the Poisson parameter

$$\lambda(t) = {}^{l(t+\mathbb{E}[D])}/_{\mathbb{E}[L]\mathbb{E}[D]}. \tag{4}$$

An example generation of appliance arrivals is presented in Fig. 2 showing the output of the $M_t/G/\infty$ queue with the real load curves of the set of appliances from [24] (blue line), using the same appliances represented by the generic average power consumption model (green line), and the desired load profile (red dashed line) of the house for three days in June 2008. The generated usage pattern is close to the desired usage pattern, and most important, the differences between the real appliance load curves and average load models are similar. This is because the errors, both positive and negative, average out over time between the two models, and thus, we believe, the constant load model is appropriate for these simulation studies. The discrepancies between the generated curves from the $M_t/G/\infty$ queue model and the desired load curve can be explained because (a) the arriving appliances are a random process and inherently have some form of stochasticity, and (b) the LIN-S is an approximation of the average number of running appliances and is not exact. The generated pattern is averaged over 500 samples to minimize the impact of (a).

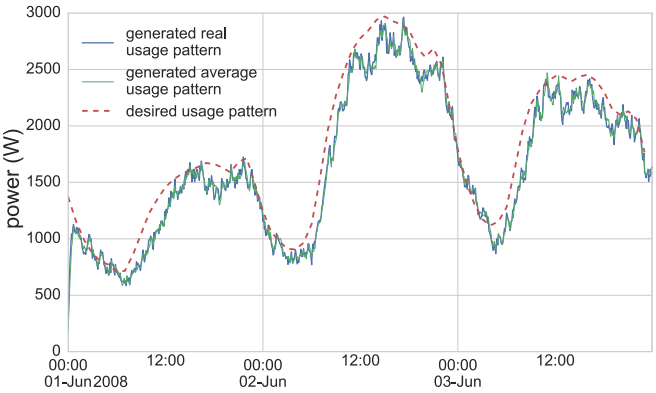

Fig. 2. An example usage pattern generated by the $M_t/G/\infty$ queue model using the real appliance set from [24] with the real load curve (blue line) and equivalent generic model (green line). The dashed red line represents the desired load curve ($l(t)$). The solid blue and green lines are the generated usage patterns averaged over 500 samples.

C. Real-Time Pricing Model

The RTP market is modeled after the ComEd Residential RTP (RRTP) program [26]. The price of electricity changes every hour in response to the PJM real-time hourly market price. At approximately 4:30 p.m., the day-ahead clearing price for the next day's hourly prices are provided to the customer — henceforth called the *forecast price*. At the start of each hour, the actual price for that hour is provided.

An example of the RTP market is given in Fig. 3. At 11:00 a.m. in Fig. 3, the price of electricity is known until noon. A forecast price is known for each hour until midnight. At 4:30 p.m., the exact price of electricity is known until 5:00 p.m., but an additional 24-hours of forecast information is provided. To the left side of the red-dotted line in Fig. 3, the actual and forecast prices are much different. At time $(0, 1)$ and $(2, 3)$, the price of electricity is actually *negative*. This occurs when the demand for electricity is low compared to the available supply and it is cheaper for large generators to pay for the consumption of electricity than it is to shut the generator down and re-start it at a later time. The non-myopic HEMS can take advantage of this phenomenon to provide superior cost savings on the electricity bill.

At the current time t, let $c(t)$ be the cost of electricity (in cents/kWh) and $c_f(t, \tau)$ be the forecast price of electricity at time τ given the current forecast. The maximum forecast time, $\tau_{\max}$, corresponds to the latest forecast provided by the utility (according to the operation illustrated in Fig. 3).

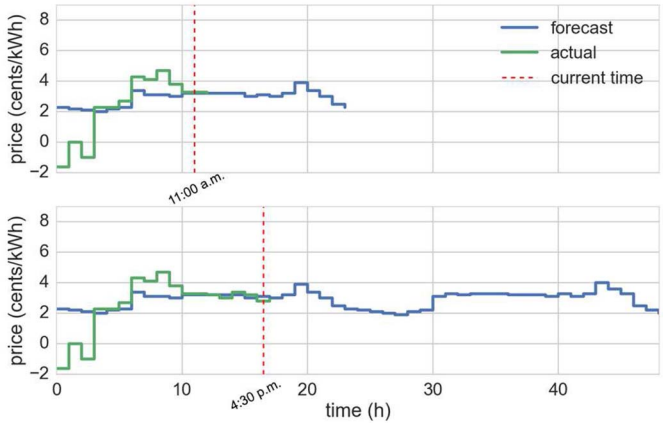

Fig. 3. An example of the ComEd Residential RTP market [26].

To validate the equivalent generic appliance load model from the previous subsection, we quantify the maximum error in electricity cost of using an appliance in the ComEd RRTP market compared to the real appliance load curves from [24]. The maximum error in energy consumption between the two models from the set of appliances is 0.44 kW. In addition, the maximum absolute difference in price of electricity between two consecutive hours in the ComEd RRTP market is 41.1 cents/kWh. If that appliance runs at a time that the price of electricity changes at exactly this maximum error during the largest difference between two consecutive pricing periods, the worst-case error in terms of cost between the two appliance load models is \$0.18. However, the difference between two consecutive prices in the ComEd RRTP market can be positive or negative, with the average difference close to zero (-1.8×10^{-5} cents/kWh), so the error in cost between models will cancel through time. Therefore, because of the canceling of errors in both load (as in Fig. 2) and price, we conclude that the equivalent generic appliance load model is adequate for long-term simulation studies.

D. Problem Statement

In the residential house, we have a vector of dynamically arriving appliances to be used by the customer at times $\hat{t}_{\text{arr}}$. We only know information about appliances with an arrival time before the current time t. We have a corresponding vector with information about the start-time deadlines ($\hat{t}_{\text{dead}}$). Let the i^{th} element in the vectors correspond to the i^{th} appliance. The goal is to find the vector of start-times, $\hat{t}_{\text{start}}$, to minimize the total cost of using the appliances.

Let $C(\hat{t}_{\text{start}})$ be the cost to run appliances at the scheduled start-times, $\hat{t}_{\text{start}}$; and $|\cdot|$ be a unary operator that returns the length or size of a vector or set, respectively.

$$C\left(\hat{t}_{\text{start}}\right) = \sum_{i=1}^{|\hat{t}_{\text{start}}|} \int_{\hat{t}_{\text{start}}[i]}^{\hat{t}_{\text{start}}[i]+d_i} p_i c(t) dt \tag{5}$$

We model the customer inconvenience cost as being insignificant within the scheduling window of the appliances. However, an additional inconvenience cost (as defined by the customer) could be added within the summation term in (5), such as in [27]. The formal problem statement is

$$\min_{\hat{t}_{\text{start}}} C\left(\hat{t}_{\text{start}}\right) \quad \text{s.t.} \quad \hat{t}_{\text{arr}} \leq \hat{t}_{\text{start}} \leq \hat{t}_{\text{dead}}. \tag{6}$$

III. Optimization Methods

A. Overview

To evaluate the delayed gratification approach, we have designed a myopic comparison algorithm, denoted minimum forecast cost (MFC). These new methods are evaluated against the status quo (appliances are used without regard to price) and the mathematical lower bound on cost.

B. Immediate

The status quo of energy usage within the household is to maximize personal comfort, often with little or no regard to price. In our simulation, this is analogous to using an appliance as soon as it arrives. The status quo optimization method, denoted *immediate* (imm), is then $\hat{t}_{\text{start}} = \hat{t}_{\text{arr}}$.

C. Minimum Forecast Cost

The obvious optimization approach is to run the appliances at the minimum forecast cost, as provided by the local utility. When an appliance i arrives at time $t_{i-\text{arr}}$, MFC schedules the appliance to run at the cheapest time over its duration, d_i, that satisfies $t_{i-\text{start}} \leq t_{i-\text{dead}}$. Let $C_i(t)$ be the exact cost of running appliance i at time t, given by (7):

$$C_i(t) = \int_t^{t+d_i} p_i c(t) dt \tag{7}$$

The *forecast* cost of running an appliance i at time t, if the current time is t_{cur}, is then

$$C_{i-\text{fore}}(t) = \int_t^{t+d_i} p_i c_f(t_{\text{cur}}, t) dt. \tag{8}$$

The goal of the myopic MFC algorithm is, for each appliance i, to set $t_{i-\text{sched}}$ to the time that minimizes (8). Formally, for each appliance i, MFC solves (9):

$$t_{i-\text{sched}} = \underset{t}{\text{argmin}}\, C_{i-\text{fore}}(t) \tag{9}$$

$$\text{s.t. } t_{\text{cur}} \leq t \leq \min\{\tau_{\max} - d_i, t_{i-\text{dead}}\}. \tag{10}$$

D. Partially Observable Markov Decision Process

1) Overview: We adopt the POMDP optimization approach from [28], a non-myopic receding horizon control method that balances the trade-off between immediate knowledge and future performance (in this case, cost). By approaching the price of electricity as a stochastic input using historical information, we can determine the conditional probability (*posterior probability distribution*) of future electricity prices based on the current price and utility forecast error.

To balance the trade-off between immediate and future decisions, we use Q-value approximation in the form of Bellman's equation [29]. For each appliance currently ready to run, the customer can take an action, a_i, from the set of possible actions A (i.e., $a_i \in A$). Let $\hat{a}$ be the vector consisting of individual

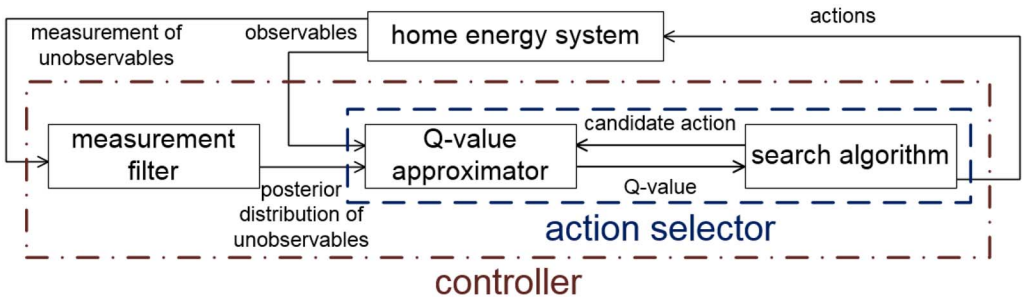

Fig. 4. The general POMDP framework. The underlying belief state is split between *observables* and *unobservables*. The addition of unobservable portions of the state require the formulation of the POMDP, as opposed to an MDP. The unobservables are measured, and the conditional probability of their true state is determined with a measurement filter. The combination of the observable state and conditional probability of the unobservable state are used to select actions.

appliance actions (henceforth known as *the action*) to be determined by the HEMS, x be the current state, x' be the next state (after taking action $\hat{a}$), $R(x, \hat{a})$ be the immediate reward for taking action $\hat{a}$ in state x, and $V^*(x)$ be the optimal cumulative reward value over the time horizon given an initial state x. We want to find the optimal action policy, $\pi^*(x)$, that maps states to actions to maximize the Q-value, $Q(x, \hat{a})$, given by

$$Q(x, \hat{a}) = R(x, \hat{a}) + \mathbb{E}\left[V^*(x')|x, \hat{a}\right]. \tag{11}$$

The action $\pi^*(x)$ is based on Bellman's principle [29], and given by

$$\pi^*(x) = \underset{\hat{a}}{\operatorname{argmax}}\, Q(x, \hat{a}). \tag{12}$$

The HEMS will take actions $\hat{a} = \pi^*(x)$ at each state x. The POMDP framework in Fig. 4 is adapted from [28].

In formulating the HEMS problem as a POMDP, we assume that there are two types of output from our home energy system: (a) the observables and (b) measurements of the unobservables. The measurements of the unobservables are filtered to determine the posterior distribution of the unobservables that, along with the observables, determine the *belief state*. At time t, let y_t represent the underlying state of the POMDP HEMS, $\hat{\Psi}_t$ be a vector of random variables describing the future prices of electricity, ϵ_t be the error between the utility forecast and current price of electricity, and H_t be the set of appliances ready to start. The underlying state is then $y_t = (c(t), \hat{\Psi}_t, \epsilon_t, H_t)$, where $\hat{\Psi}_t$ is unobservable. The measurement we have for future price, $\hat{\Psi}_t$, is the utility forecast price, $c_f(t, \tau)$, where $|\hat{\Psi}_t[\tau]| = \tau_{\max}$. Given measurements $c_f(t, \tau)$, we can determine $P\left(\hat{\Psi}_t | c_f(t, \tau)\right)$ using a filtering method. Here, we use particle filtering [30] — a sequential Monte Carlo sampling method that maintains a set of representative samples (particles) — to calculate the posterior distribution.

The action set available for each appliance is to either run now or wait to run at a later time, i.e., $A = \{\text{run}, \text{wait}\}$. At each decision event at time t, we must determine $\hat{a}$, where $|\hat{a}| = |H_t|$ and each a_i corresponds to appliance $H_t[i]$, to maximize $Q(x, \hat{a})$. This decision is performed in the action selector of Fig. 4. The action selector is comprised of the Q-value approximator and a search algorithm. To approximate the Q-value, we use Monte Carlo sampling in conjunction with the particle filter and a rollout-based search algorithm (explained in detail later).

2) Particle Filter: Two separate particle filters are used within the POMDP framework to create two distinct POMDP HEMS algorithms. The first, which we denote *POMDP-Gauss*, keeps track of the mean and standard deviation of a Gaussian estimate of the error between the forecast and actual RTP at each hour of the day. The second, which we denote *POMDP-GARCH*, keeps track of the parameters in an autoregressive (AR) process with a generalized autoregressive conditional heteroskedastic (GARCH) error process. Let the set of particles in the filter be denoted K.

For the POMDP-Gauss particle filter, we can construct an estimate of the error from previous errors for each hour of the day. For each hour of the day h, we have an associated error $\delta_h \sim \mathcal{N}(\mu_h, \sigma_h)$, where $\mathcal{N}(\mu_h, \sigma_h)$ is a normal distribution with mean μ_h and standard deviation σ_h. To determine the initial mean and standard deviation, we calculate the sample mean and standard deviation from the previous N_d errors. Let us assume the simulation start time t^0 starts at hour 0, and $h = t \mod 24$. At the current time t, given a measured observation $c_f(t, \tau)$, then each particle $\rho_i \in K$ is defined by

$$\rho_i = c_f(t, \tau) + \delta_h. \tag{13}$$

Assuming each particle has the same weight, the particle samples are used to estimate the conditional probability of the price $c(t + \tau)$ given $c_f(t, \tau)$.

The second filter, POMDP-GARCH, an AR process with GARCH error is used to model the error in RTP. The GARCH model has been used previously to estimate high variability in power markets [31]. In an AR process, the current output depends on a linear combination of prior values, plus some error term. At time t, let $c_{\text{ar}}(t)$ be the cost output of the AR process, k be the AR constant, $c_{\text{ar}}(t - i)$ be the i^{th} previous output, γ_i be the coefficient corresponding to $c_{\text{ar}}(t - i)$, m be the number of modeled coefficients, and $\epsilon_{t-\text{ar}}$ be the error. The AR process is given as

$$c_{\text{ar}}(t) = k + \sum_{i=1}^{m} (\gamma_i c_{\text{ar}}(t - i)) + \epsilon_{t-\text{ar}}. \tag{14}$$

The error in (14) is modeled as a GARCH process. The GARCH process is a specialized AR process that has a conditional variance based on prior inputs. At time t, let σ_t be the standard deviation and $z_t \sim \mathcal{N}(0, 1)$. The GARCH error process is given as $\epsilon_{t-\text{ar}} = \sigma_t z_t$.

The term σ_t is itself a linear combination of prior inputs, hence the conditional heteroskedasticity of the GARCH process. At time t, let χ be the GARCH constant, σ^2_{t-i} be the i^{th} previous variance, ϕ_i be the coefficient corresponding to σ^2_{t-i}, P be the number of GARCH terms (i.e., prior variances), ϵ^2_{t-j} be the j^{th} previous square-error, q_j be the coefficient corresponding to ϵ^2_{t-j}, and Q be the number of ARCH terms (i.e., prior square-errors). A GARCH(P, Q) process is fully described by $\epsilon_{t-\mathrm{ar}}$ and

$$\sigma_t^2 = \chi + \sum_{i=1}^{P} \phi_i \sigma^2_{t-i} + \sum_{j=1}^{Q} q_i \epsilon^2_{t-j}. \tag{15}$$

The AR+GARCH process to be used in the POMDP-GARCH is obtained by substituting $\epsilon_{t-\mathrm{ar}}$ into (14), resulting in

$$c_{\mathrm{ar}}(t) = k + \sum_{i=1}^{m} (\gamma_i c_{\mathrm{ar}}(t-i)) + \sigma_t z_t. \tag{16}$$

At time t, the each particle $\rho_i \in K$ of the particle filter corresponding to POMDP-GARCH has its own list of the prior m AR outputs, P error variances, and Q square-errors. It uses this to determine the conditional probability of the next $\tau_{\max}$ prices given $c_f(t, \tau)$.

3) Action Selector: The action selector, shown as the red dotted box in Fig. 4, consists of an estimate of the Q-value of taking an action and an optimization algorithm that maximizes this value over the possible actions. The Q-function maps an action at a belief state to the Q-value. The HEMS problem is relatively unique in that an accurate cost of taking a given action is known and that the actions may be evaluated independently for their cost, reducing the search space at each belief state. Recall that (11) consists of the summation of the *immediate* reward for taking an action, $R(x, \hat{a})$, and the *expected future* reward, $\mathbb{E}[V^*(x')|x, \hat{a}]$ (also known as the expected reward-to-go). These map directly to our action set, $A = \{\text{run}, \text{wait}\}$.

Because we can evaluate each appliance action individually, we reduce the maximum number of evaluations from $2^{|H|}$ to $2|H|$. At state x, we have an estimate of the price of electricity through $\tau_{\max}$, denoted $c'(t), t_{\mathrm{cur}} \leq t \leq \tau_{\max}$. The cost of running appliance i at time t using $c'(t)$ is

$$C'_i(t) = \int_t^{t+d_i} p_i c'(t) dt. \tag{17}$$

To map this state x to a Q-value given an action a_i corresponding to appliance i, we use the following

$$Q_i(x, a_i) = \begin{cases} -C'_i(t_{\mathrm{cur}}) & a_i = \text{run} \\ \max_t -C'_i(t) \ t > t_{\mathrm{cur}} & a_i = \text{wait}. \end{cases} \tag{18}$$

The costs of running the appliances are negated because the HEMS is a cost minimization problem. The total Q-value determined from the Q-function is

$$Q(x, \hat{a}) = \sum_{i=1}^{|H|} Q_i(x, \hat{a}[i]). \tag{19}$$

Note that the Q-value in (19) captures the effect of the current action $\hat{a}$ on the future of the system through (17).

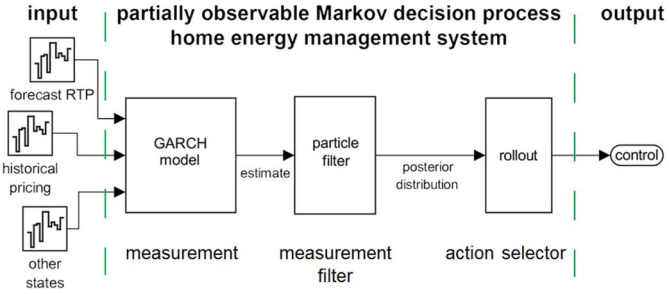

Fig. 5. The POMDP-GARCH HEMS. The measurements of the unobservable state are provided by the utility-forecast cost. The GARCH process combined with the particle filter provides an estimate of the actual RTP. The output of the particle filter, along with the appliances ready-to-run, are used to determine which action to take for each appliance. The GARCH block can be replaced with the Gaussian-noise estimate to obtain the POMDP-Gauss HEMS.[1]

Recall from Bellman's principle that the future effect represented by the expected future reward depends on the optimal policy. Because the optimal policy is unknown, the usual approach is to approximate the Q-value. To approximate the Q-value, we use a method called *policy rollout* [32]. We use a policy approximation because exact methods, such as policy iteration and backward induction, are impractical due to the size of the state space. In policy rollout, we replace the optimal future reward V^* in (11) with the future reward associated with a base policy $V_{\pi\mathrm{base}}$. This gives us the rollout-approximate Q-value

$$Q_{\pi\mathrm{base}}(x, \hat{a}) = R(x, \hat{a}) + \mathbb{E}\left[V_{\pi\mathrm{base}}(x')|x, \hat{a}\right] \tag{20}$$

and the corresponding rollout policy

$$\pi(x) = \underset{\hat{a}}{\operatorname{argmax}}\, Q_{\pi\mathrm{base}}(x, \hat{a}). \tag{21}$$

To compute the output of this policy, we use Monte Carlo sampling to estimate the expected future reward, $\mathbb{E}[V_{\pi\mathrm{base}}(x')|x, \hat{a}]$, which works well with the output of the particle filter, where each particle represents a single sample of the unobservable state. For each appliance i, we determine the action a_i at time t_{cur} as

$$a_i = \underset{A}{\operatorname{argmax}}\, Q_i(x, a_i) \tag{22}$$

where $C'_{i,j}(t)$ is the sample cost of appliance i in particle j

$$Q_i(x, a_i) = \begin{cases} |K|^{-1} \sum_{j=1}^{|K|} -C'_{i,j}(t_{\mathrm{cur}}), & a_i = \text{run} \\ |K|^{-1} \sum_{j=1}^{|K|} \max_t -C'_{i,j}(t), & a_i = \text{wait}, \end{cases}$$
$$t_{\mathrm{cur}} \leq t \leq \min\{\tau_{\max} - d_i, t_{i-\mathrm{dead}}\}. \tag{23}$$

The complete POMDP HEMS controller is given in Fig. 5. Although Monte Carlo sampling is used, the action at each time t can be calculated in real-time.

[1]Note that the GARCH forecast model can be replaced with any other forecast model in the POMDP HEMS that outputs a measurement of the unobservable states, in this case the future price of electricity. The inputs to the measurement process can be any information available to the entity (e.g., historical RTP pricing and the forecast RTP in our GARCH model).

E. Theoretical Lower Bound

We have designed a theoretical lower bound (LB) using *a posteriori* information. The LB determines the maximum gap between our methods and the actual optimal solution, and is obtained by scheduling each appliance i at the absolute minimum cost time given by

$$t_{i-\text{sched}} = \underset{t}{\operatorname{argmin}}\, C_i(t) \tag{24}$$

$$\text{s.t. } t_{\text{cur}} \leq t \leq t_{i-\text{dead}}. \tag{25}$$

Note the difference in the scheduling constraint between (10) and (25). MFC can only make decisions up to the forecast information provided by the utility whereas LB assumes perfect information of the RTP up to each appliance start-time deadline, providing a mathematical lower bound.

IV. Simulation Setup

A. Overview

This section describes the models, input data, and simulation parameters used. The methods introduced up to this point can be used on any inputs for a HEMS, including dynamic pricing markets. To generalize the above methods for time-varying loads, p_i can be replaced by $p_i(t)$ in (5), (7), (8), and (17).

B. Appliance Types

For simulation purposes, we model abstract appliance types. Let N_a be the number of appliance types. Each appliance *type* j has a power rating in kW (p_j) and a random variable modeling the duration in hours (ϕ_j). The reason that the duration is stochastic is that, although appliance run-times may be similar between uses, the run-time of an appliance usually has some variance. To generate p_j and ϕ_j for each appliance type, we use a Gaussian distribution and the coefficient-of-variation-based (CVB [17], [33]) method, respectively. Let $p_j \sim \mathcal{N}(\mu_p, \sigma_p)$, where μ_p and σ_p are the mean and standard deviation, in kW, of the distribution of appliance power ratings. Let $\phi_j \sim \mathcal{G}(\mu_{j-t}, \theta_t)$ where $\mathcal{G}(\mu_{j-t}, \theta_t)$ is a Gamma distribution with mean μ_{j-t} and coefficient-of-variation θ_t (where the relationship to the shape is θ_t^{-2} and the scale is $\mu_{j-t}\theta_t^2$). The mean, μ_{j-t}, is itself generated by a Gamma distribution, $\mathcal{G}(\mu_d, \theta_d)$. The parameters N_a, μ_p, σ_p, μ_d, θ_d, and θ_t are determined empirically, and given in the Appendix. In an implementation of the HEMS for a specific household, the actual power rating and distribution of duration times for the set of appliances can be obtained through analytical, historical, or experimental techniques [34]–[36].

C. Household Usage Pattern

Recall from (4) in Section II-B that the $M_t/G/\infty$ queue requires the desired load curve for the household, $l(t)$, as an input. To obtain a time-varying household load curve that accurately models the daily, weekly, and seasonal change in energy usage, we scale the ComEd system load to match the load of a single household. Let $l_{\text{sys}}(t)$ be the ComEd system load at time t, obtained from [37]. Let $f(l)$ be a function that scales $l_{\text{sys}}(t)$ to $l(t)$, where $l(t) = f(l_{\text{sys}}(t))$. The function $f(l)$ is comprised of two parts: (a) normalizing the system load to [0, 1], and (b) upscaling the normalized load to the desired $l(t)$. Let $S_{\text{scale}}(l)$ be the normalizing function for a given load l, given by

$$S_{\text{scale}}(l) = \frac{l - \min_t l_{\text{sys}}(t)}{\max_t l_{\text{sys}}(t) - \min_t l_{\text{sys}}(t)}. \tag{26}$$

Let $o_{\min}$ and $o_{\max}$ be the minimum and maximum household usage, in kW, respectively. We can then define the scaling function as $f(l) = o_{\min} + S_{\text{scale}}(l)(o_{\max} - o_{\min})$. Defining the previous equation in terms of household load through time, we get $l(t) = o_{\min} + S_{\text{scale}}(l_{\text{sys}}(t))(o_{\max} - o_{\min})$.

We still need to obtain specific appliance arrivals from the appliance types and the $M_t/G/\infty$ queue. At a given time t, the next appliance arrival, i, is sampled from an exponential distribution with rate $\lambda(t)$ to obtain t_{next}. This gives appliance i an arrival time of $t_{i-\text{start}} = t + t_{\text{next}}$. The appliance type j corresponding to the arrived appliance i is selected randomly from the N_a types with equal probability. This sets $p_i = p_j$ and d_i to a sample of ϕ_j. To set the start-time deadline, $t_{i-\text{dead}}$, we randomly sample from the set of $t_{i-\text{start}} + \{1\text{h}, 2\text{h}, 4\text{h}, 8\text{h}\}$. The varying length of the deadlines represents the flexibility the customer may have with each appliance.

TABLE I
Simulation Scenario Characteristics

Name	Month	RMS Error (cents)	Min. Price (cents/kWh)	Avg. Price (cents/kWh)	Max. Price (cents/kWh)
A	Oct. 2009	0.72	-1.6	3.00	8.1
B	Jan. 2011	1.46	-7.5	3.93	20.7
C	June 2008	4.22	-21.1	5.94	48.7

D. Simulation Scenarios

For our simulation study, we chose scenarios that represent one month of electricity usage to obtain a monthly electricity bill. We used the actual and forecast pricing data from ComEd [26] between 2007 and 2013. Our GARCH model was trained on the errors in 2007.[2] For the POMDP-Gauss model, we set $N_d = 60$. Both particle filters used 1000 particles to estimate the posterior distribution of unobservables ($|K| = 1000$).

We chose three scenarios, each representing one month in the ComEd system. The scenarios (A, B, and C) correspond to the minimum, median, and maximum months in the input data in terms of root-mean-square (RMS) error between the forecast and actual price and are summarized in Table I. Each scenario was run for 50 trials. For a given scenario, the actual and forecast RTP (from the ComEd data) and the scaled household load stay the same. Between trials, (a) the power and duration values for appliance types and (b) appliance arrivals differ.

[2]The coefficients for the AR+GARCH process can be found at the project website at http://www.engr.colostate.edu/sgra.

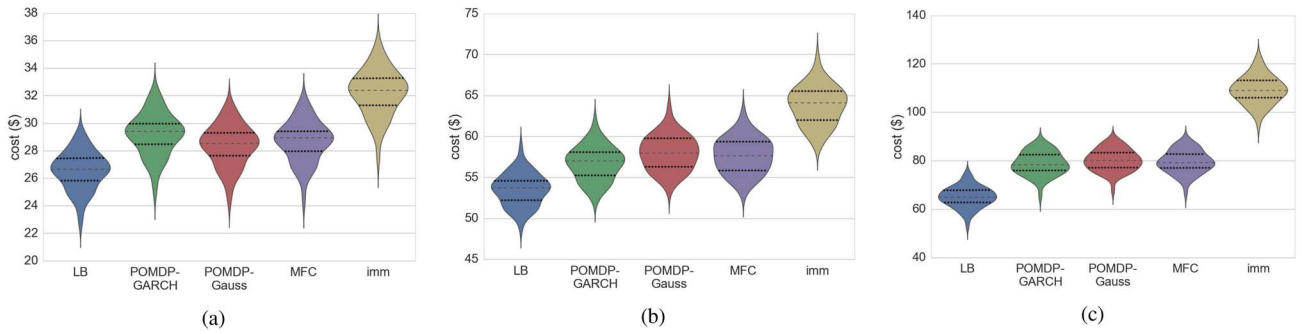

Fig. 6. A comparison of the monthly electricity bills for each optimization method for scenarios A, B, and C in subfigures (a), (b), and (c), respectively. The violin plots show the probability density of each cost. The dashed line corresponds to the median cost and the dotted lines correspond to the quartiles.

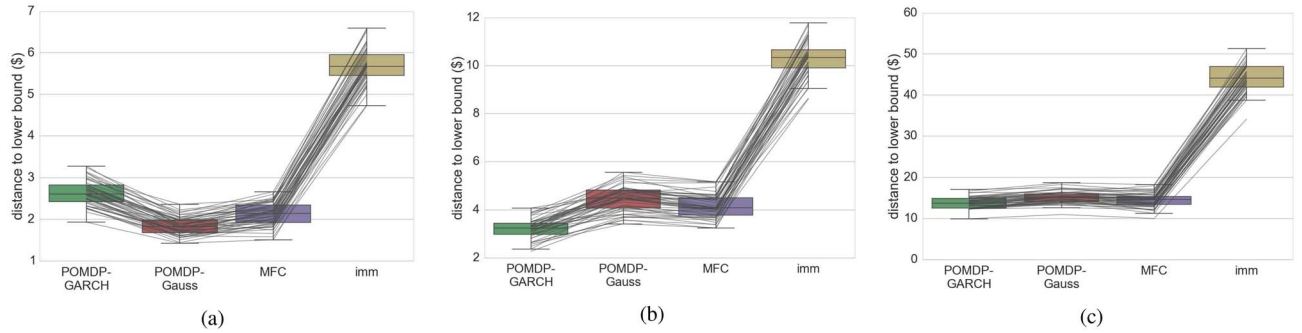

Fig. 7. A comparison of the distance to the lower bound for scenarios A, B, and C in subfigures (a), (b), and (c), respectively. The box plot shows the median, quartiles, and 10th percentiles. The lines between the box plots connect individual trials.

TABLE II
AVERAGE PRICE OF ELECTRICITY (CENTS/KWH)

Name	Energy (kWh)	LB	POMDP-GARCH	POMDP-Gauss	MFC	imm
A	1037	2.56	2.82	2.74	2.77	3.10
B	1576	3.40	3.60	3.69	3.67	4.04
C	1524	4.27	5.18	5.26	5.22	7.14

V. SIMULATION RESULTS

The different HEMS algorithms were evaluated against the three scenarios. Before conducting the simulations, we would expect the POMDP methods to improve over the myopic MFC in situations where the forecast and actual price diverge (as in scenario C). Table II compares the average price of electricity paid by the customer during the scenario time periods between the different HEMS methods. We see that the three intelligent HEMS methods are bounded by LB and status quo (imm). MFC performs slightly better than POMDP-GARCH in scenario A, but slightly worse than POMDP-Gauss. The GARCH technique is well suited to signals with high variance, leading to a slightly worse prediction in the non-variable scenario A. The POMDP-Gauss method is able to take advantage of the delayed gratification over the myopic MFC. In general, as the RMS-error of the RTP increases, the POMDP-GARCH performs better relative to the other techniques. As the variance of the RTP signal increases, the Gaussian assumption made by POMDP-Gauss begins to degrade performance. It is important to note that *all* methods, including the status quo, resulted in a lower per-kWh price than the tariff-based ComEd pricing (approximately 7.5 cents/kWh).

The monthly electricity bills across scenario A, B, and C are presented in Figs. 6(a), 6(b), and 6(c), respectively. Each violin plot shows the probability distribution of obtaining a specific monthly electricity bill (the wider the figure, the higher the probability density). The dashed line indicates the median cost and the dotted lines indicate the quartiles. The average monthly electricity bill on the tariff-based rate is $77.78, $118.20, and $114.30 for scenarios A, B, and C, respectively. As a customer, just opting-in to the RTP program with no change in energy usage offers the potential for significant savings. By combining a RTP program with a smart HEMS, the savings are improved. In general, the shapes of the distributions and coefficient-of-variation between the methods within the same scenario are similar, indicating a strong correlation between an individual trial and the monthly cost (within a trial, the appliance arrival patterns are the same and the overall energy usage is the same). This trial-cost correlation explains the large overlap between some of the distributions. The relative performance of the methods to the lower bound is presented in Fig. 7 where each box plot shows the median and quartiles, and the whiskers show the 10th percentile. The lines between the optimization methods connect the individual trials. In general, the relative performance of a trial is similar between the methods (e.g., the median trial for one method is close to the median trial for another). To choose the appropriate HEMS algorithm to use in real-time, a metaheuristic could be used as a selection

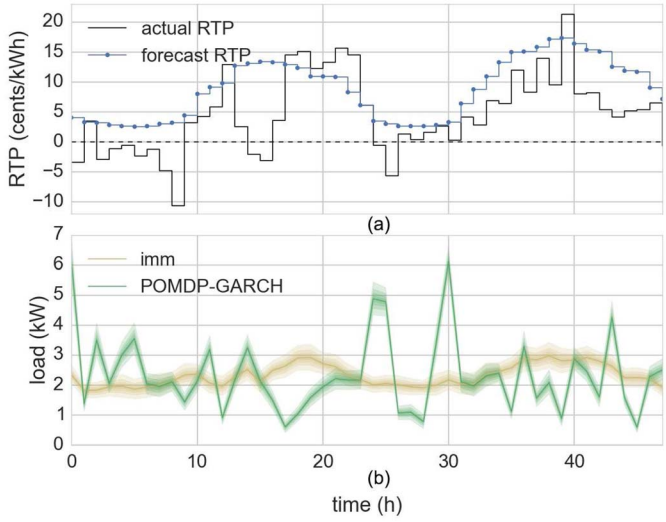

Fig. 8. (a) A 48-hour sample of the actual and forecast RTP. (b) Statistical time-series of the household load compared between the immediate and POMDP-GARCH methods.

procedure to determine the best performing predictor given the current system characteristics and conditions. An example selection procedure is to run each method simultaneously, and use the scheduling decision of the method that is currently performing best.

The change in energy usage in the household between the status quo and the POMDP-GARCH HEMS is presented in Fig. 8. Fig. 8(a) shows the difference between the actual and forecast RTP for a 48-hour period in scenario C. The curves in Fig. 8(b) present the HEMS response to the pricing signals. The cloud around the plot represents the statistical differences between trials. A drastic change in electricity usage occurs when switching to a HEMS algorithm—the POMDP-GARCH in this case—from the status quo. The POMDP-GARCH HEMS actively prevents using appliances during high-price times, but also predicts low-price times, such as at $t = 25$ hours. It is interesting to note that the peak load from the POMDP-GARCH HEMS is *higher* than the status quo. This is not a problem for a single household, but if many households within a given area act in a similar way it may cause problems such as the rebound effect [38]. To overcome the rebound effect, system-level-type constraints can be incorporated, such as maximum house load or peak-to-average ratio.

VI. Conclusion

Reducing the peak demand of the electric power system provides benefits by reducing the cost of electricity by lowering the number of expensive generators needed. By reducing the peak, we can reduce the capacity factor of dirty diesel-fired peaking generators. Moreover, as peak demand increases, the available transmission capacity will also need to increase. By reducing the peak demand, we can defer building new transmission lines; a costly, long-term project. Utilities are offering real-time pricing programs, passing ISO prices to customers, such as the ComEd residential real-time pricing program. Customers can take advantage of this real-time pricing to drastically reduce their monthly electric bill.

The partially observable Markov decision process is a promising non-myopic method for home energy management. On the high-end, the POMDP-GARCH HEMS resulted in a $30 cost savings over the status quo for the June 2008 scenario. Even for more modest months, savings of 10% were shown. It should be noted that these large savings were obtained just using flexible appliances. As more asset types with more capabilities, such as electric vehicles and HVAC, are added, the savings can be expected to increase.

Appendix

For this simulation study, we set $N_a = 10$, $\mu_p = 1$ kW, $\sigma_p = 0.25$ kW, $\mu_d = 1$ hour, $\theta_d = 0.5$, and $\theta_t = 0.05$. This corresponds to an average appliance load of 1 kW and duration of 1 hour. The coefficient-of-variation between appliance types (θ_d) is relatively large, while the variation within an appliance type is small (θ_t). This is because there is usually a large variation in the run times of different appliances, but when using the same appliance it usually runs for similar times. For scaling the load, we set $o_{\min} = 0.5$ kW and $o_{\max} = 6$ kW. This implies that at the ComEd system minimum, the house will be drawing approximately 500 W, and at the ComEd system maximum the house will be drawing approximately 6 kW. These parameters may be changed to simulate differently rated households.

References

[1] United States Energy Information Administration. (May 2014). *Annual Energy Outlook 2014*. [Online]. Available: http://www.eia.gov/forecasts/aeo/pdf/0383(2014).pdf

[2] P. Siano, "Demand response and smart grids—A survey," *Renew. Sustain. Energy Rev.*, vol. 30, pp. 461–478, Feb. 2014.

[3] A. Zipperer, P. A. Aloise-Young, and S. Suryanarayanan, "On the design of a survey for reconciling consumer behaviors with demand response in the smart home," in *Proc. North Amer. Power Symp. (NAPS)*, Manhattan, KS, USA, Sep. 2013, pp. 1–6.

[4] M. H. Albadi and E. F. El-Saadany, "A summary of demand response in electricity markets," *Elect. Power Syst. Res.*, vol. 78, no. 11, pp. 1989–1996, 2008.

[5] A. J. Coñejo, J. M. Morales, and L. Baringo, "Real-time demand response model," *IEEE Trans. Smart Grid*, vol. 1, no. 3, pp. 236–242, Dec. 2010.

[6] A.-H. Mohsenian-Rad and A. Leon-Garcia, "Optimal residential load control with price prediction in real-time electricity pricing environments," *IEEE Trans. Smart Grid*, vol. 1, no. 2, pp. 120–133, Sep. 2010.

[7] P. Du and N. Lu, "Appliance commitment for household load scheduling," *IEEE Trans. Smart Grid*, vol. 2, no. 2, pp. 411–419, Jun. 2011.

[8] H. Tischer and G. Verbic, "Towards a smart home energy management system—A dynamic programming approach," in *Proc. IEEE PES Innov. Smart Grid Technol. Asia (IGST)*, Perth, WA, Australia, Nov. 2011, pp. 1–7.

[9] Z. Chen, L. Wu, and Y. Fu, "Real-time price-based demand response management for residential appliances via stochastic optimization and robust optimization," *IEEE Trans. Smart Grid*, vol. 3, no. 4, pp. 1822–1831, Dec. 2012.

[10] T. Hubert and S. Grijalva, "Modeling for residential electricity optimization in dynamic pricing environments," *IEEE Trans. Smart Grid*, vol. 3, no. 4, pp. 2224–2231, Dec. 2012.

[11] M. Pipattanasomporn, M. Kuzlu, and S. Rahman, "An algorithm for intelligent home energy management and demand response analysis," *IEEE Trans. Smart Grid*, vol. 3, no. 4, pp. 2166–2173, Dec. 2012.

[12] K. M. Tsui and S. C. Chan, "Demand response optimization for smart home scheduling under real-time pricing," *IEEE Trans. Smart Grid*, vol. 3, no. 4, pp. 1812–1821, Dec. 2012.
[13] Y. Guo, M. Pan, Y. Fang, and P. P. Khargonekar, "Decentralized coordination of energy utilization for residential households in the smart grid," *IEEE Trans. Smart Grid*, vol. 4, no. 3, pp. 1341–1350, Sep. 2013.
[14] Z. Yu, L. Jia, M. C. Murphy-Hoye, A. Pratt, and L. Tong, "Modeling and stochastic control for home energy management," *IEEE Trans. Smart Grid*, vol. 4, no. 4, pp. 2244–2255, Dec. 2013.
[15] Z. Zhao, W. C. Lee, Y. Shin, and K.-B. Song, "An optimal power scheduling method for demand response in home energy management system," *IEEE Trans. Smart Grid*, vol. 4, no. 3, pp. 1391–1400, Sep. 2013.
[16] M. Beaudin, H. Zareipour, A. K. Bejestani, and A. Schellenberg, "Residential energy management using a two-horizon algorithm," *IEEE Trans. Smart Grid*, vol. 5, no. 4, pp. 1712–1723, Jul. 2014.
[17] T. M. Hansen, R. Roche, S. Suryanarayanan, A. A. Maciejewski, and H. J. Siegel, "Heuristic optimization for an aggregator-based resource allocation in the smart grid," *IEEE Trans. Smart Grid*, vol. 6, no. 4, pp. 1785–1794, Jul. 2015.
[18] Q.-S. Jia, J.-X. Shen, Z.-B. Xu, and X.-H. Guan, "Simulation-based policy improvement for energy management in commercial office buildings," *IEEE Trans. Smart Grid*, vol. 3, no. 4, pp. 2211–2223, Dec. 2012.
[19] S. P. Meyn, P. Barooah, A. Bušić, Y. Chen, and J. Ehren, "Ancillary service to the grid using intelligent deferrable loads," *IEEE Trans. Autom. Control*, vol. 60, no. 11, pp. 2847–2862, Nov. 2015.
[20] S. Misra *et al.*, "Residential energy management in smart grid: A Markov decision process-based approach," in *Proc. Green Comput. Commun. (GreenCom)*, Beijing, China, Aug. 2013, pp. 1152–1157.
[21] A. J. Conejo, F. J. Nogales, and J. M. Arroyo, "Price-taker bidding strategy under price uncertainty," *IEEE Trans. Power Syst.*, vol. 17, no. 4, pp. 1081–1088, Nov. 2002.
[22] A. Zipperer *et al.*, "Electric energy management in the smart home: Perspectives on enabling technologies and consumer behavior," *Proc. IEEE*, vol. 101, no. 11, pp. 2397–2408, Nov. 2013.
[23] W. Wang and Z. Lu, "Cyber security in the smart grid: Survey and challenges," *Comput. Netw.*, vol. 57, no. 5, pp. 1344–1371, Apr. 2013.
[24] M. Pipattanasomporn, M. Kuzlu, S. Rahman, and Y. Teklu, "Load profiles of selected major household appliances and their demand response opportunities," *IEEE Trans. Smart Grid*, vol. 5, no. 2, pp. 742–750, Mar. 2014.
[25] S. G. Eick, W. A. Massey, and W. Whitt, "$M_t/G/\infty$ queues with sinusoidal arrival rates," *Manag. Sci.*, vol. 39, no. 2, pp. 241–252, Feb. 1993.
[26] ComEd. *ComEd Residential Real-Time Pricing Program*. Accessed on Aug. 18, 2013. [Online]. Available: https://rrtp.comed.com/live-prices/
[27] M. A. A. Pedrasa, T. D. Spooner, and I. F. MacGill, "Coordinated scheduling of residential distributed energy resources to optimize smart home energy services," *IEEE Trans. Smart Grid*, vol. 1, no. 2, pp. 134–143, Sep. 2010.
[28] E. K. P. Chong, C. M. Kreucher, and A. O. Hero III, "Partially observable Markov decision process approximations for adaptive sensing," *Discrete Event Dyn. Syst.*, vol. 19, no. 3, pp. 377–422, Sep. 2009.
[29] R. E. Bellman, *Dynamic Programming*. Princeton, NJ, USA: Princeton Univ. Press, 1957.
[30] N. J. Gordon, D. J. Salmond, and A. F. M. Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," *IEE Proc. F Radar Signal Process.*, vol. 140, no. 2, pp. 107–113, Apr. 1993.
[31] R. C. Garcia, J. Contreras, M. van Akkeren, and J. B. C. Garcia, "A GARCH forecasting model to predict day-ahead electricity prices," *IEEE Trans. Power Syst.*, vol. 20, no. 2, pp. 867–874, May 2005.
[32] D. P. Bertsekas, *Dynamic Programming and Optimal Control*. Belmont, MA, USA: Athena Sci., 1995.
[33] S. Ali, H. J. Siegel, M. Maheswaran, D. Hengsen, and S. Ali, "Representing task and machine heterogeneities for heterogeneous computing systems," *Tamkang J. Sci. Eng.*, vol. 3, no. 3, pp. 195–208, Nov. 2000.
[34] V. Shestak, J. Smith, A. A. Maciejewski, and H. J. Siegel, "Stochastic robustness metric and its use for static resource allocations," *J. Parallel Distrib. Comput.*, vol. 68, no. 8, pp. 1157–1173, Aug. 2008.
[35] L. Wasserman, *All of Statistics: A Concise Course in Statistical Inference*. New York, NY, USA: Springer, 2005.
[36] M. Dong, P. C. M. Meira, W. Xu, and W. Freitas, "An event window based load monitoring technique for smart meters," *IEEE Trans. Smart Grid*, vol. 3, no. 2, pp. 787–796, Jun. 2012.
[37] PJM. *PJM Estimated Hourly Load*. Accessed on Apr. 30, 2015. [Online]. Available: http://www.pjm.com/markets-and-operations/energy/real-time/loadhryr.aspx
[38] R. Roche, S. Natarajan, A. Bhattacharyya, and S. Suryanarayanan, "A framework for co-simulation of AI tools with power systems analysis software," in *Proc. 23rd Int. Workshop Database Expert Syst. Appl. (DEXA)*, Vienna, Austria, Sep. 2012, pp. 350–354.

Timothy M. Hansen (S'07–GSM'11–M'15) received the B.S. (Hons.) degree from the Milwaukee School of Engineering, Milwaukee, WI, USA, in 2011, and the Ph.D. degree from Colorado State University, Fort Collins, CO, USA, in 2015.

He is currently an Assistant Professor with the Electrical Engineering and Computer Science Department, South Dakota State University, Brookings, SD, USA. His current research interests include robust resource allocation and optimization applied to cyber-physical systems, electric power engineering and markets, and high-performance computing.

Edwin K. P. Chong (F'04) received the B.E. (First Class Hons.) degree from the University of Adelaide, South Australia, in 1987, and the M.A. and Ph.D. degrees from Princeton University, in 1989 and 1991, respectively, where he held an IBM Fellowship.

He joined the School of Electrical and Computer Engineering, Purdue University, in 1991, where he was named a University Faculty Scholar in 1999. Since 2001, he has been a Professor of Electrical and Computer Engineering and a Professor of Mathematics, Colorado State University. He has co-authored the book entitled *An Introduction to Optimization* (4th Edition, Wiley-Interscience, 2013).

Prof. Chong was a recipient of the NSF CAREER Award in 1995, the ASEE Frederick Emmons Terman Award in 1998, and the IEEE Control Systems Society Distinguished Member Award in 2010, and a co-recipient of the 2004 Best Paper Award for a paper in the journal *Computer Networks*. He was the Founding Chairman of the IEEE Control Systems Society Technical Committee on Discrete Event Systems, and served as an IEEE Control Systems Society Distinguished Lecturer. He is currently a Senior Editor of the IEEE TRANSACTIONS ON AUTOMATIC CONTROL. He has served as a member of the IEEE Control Systems Society Board of Governors and as the Vice President for Financial Activities until 2014. He currently serves as the President-Elect. He was the General Chair for the 2011 Joint 50th IEEE Conference on Decision and Control and European Control Conference.

Siddharth Suryanarayanan (S'00–M'04–SM'10) was born in Chennai, India. He received the Ph.D. degree in electrical engineering from Arizona State University, in 2004.

He is currently an Associate Professor with the Department of Electrical and Computer Engineering, Colorado State University, where he supervises sponsored research and teaches in the area of electric power systems engineering.

Anthony A. Maciejewski (F'05) received the B.S.E.E., M.S., and Ph.D. degrees in electrical engineering from Ohio State University, Columbus, OH, USA, in 1982, 1984, and 1987, respectively.

From 1988 to 2001, he was a Professor of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA. He is currently a Professor and the Department Head of Electrical and Computer Engineering, Colorado State University, Fort Collins, CO, USA.

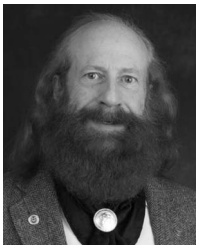

Howard Jay (H. J.) Siegel (F'90) received the double B.S. degrees in electrical engineering and management from the Massachusetts Institute of Technology, Cambridge, MA, USA, and the M.A., M.S.E., and Ph.D. degrees in electrical engineering and computer science from Princeton University, Princeton, NJ, USA, in 1972, 1974, and 1977, respectively.

Since 2001, he has been the Abell Endowed Chair Distinguished Professor of Electrical and Computer Engineering, Colorado State University, Fort Collins, CO, USA, where he is a Professor of Computer Science. From 1976 to 2001, he was a Professor with Purdue University, West Lafayette, IN, USA. He has co-authored over 430 technical papers. He has over 15 000 citations. He was a Consultant for industry and government. His current research interests include robust energy-aware resource management for heterogeneous parallel and distributed computing, the smart grid for electricity, parallel algorithms, parallel machine interconnection networks, and reconfigurable parallel computer systems.

Prof. Siegel was the Co-Editor-in-Chief of the *Journal of Parallel and Distributed Computing*. He was an Editorial Board Member of the IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS and the IEEE TRANSACTIONS ON COMPUTERS. He is a fellow of the Association for Computing Machinery.