Optimized Day-Ahead Pricing for Smart Grids with Device-Specific Scheduling Flexibility

Carlee Joe-Wong, *Student Member, IEEE,* Soumya Sen, *Member, IEEE,* Sangtae Ha, *Member, IEEE,* and Mung Chiang, *Fellow, IEEE*

***Abstract*—Smart grids are capable of two-way communication between individual user devices and the electricity provider, enabling providers to create a control-feedback loop using time-dependent pricing. By charging users more in peak and less in off-peak hours, the provider can induce users to shift their consumption to off-peak periods, thus relieving stress on the power grid and the cost incurred from large peak loads. We formulate the electricity provider's cost minimization problem in setting these prices by considering consumers' device-specific scheduling flexibility and the provider's cost structure of purchasing electricity from an electricity generator. Consumers' willingness to shift their device usage is modeled probabilistically, with parameters that can be estimated from real data. We develop an algorithm for computing *day-ahead* prices, and another algorithm for estimating and refining user reaction to the prices. Together, these two algorithms allow the provider to dynamically adjust the offered prices based on user behavior. Numerical simulations with data from an Ontario electricity provider show that our pricing algorithm can significantly reduce the cost incurred by the provider.**

***Index Terms*—Smart-Grid pricing; demand response; patience index; day-ahead pricing**

I. Introduction

A BASIC purpose of smart grids is to create an automated, widely distributed energy delivery network that uses smart meters to facilitate two-way flows of information and electricity between energy consumers and providers. This transformation enables greater support for demand response and provides more flexibility in demand shaping through time-dependent pricing (TDP).

In a smart grid infrastructure, electricity providers can send pricing information from their pricing database to the Energy Consumption Controller (ECC) unit located at the consumer's smart meters, as shown in Fig. 1. The ECC can monitor and control a consumer's energy consumption by scheduling device activities at periods of lower prices. An increasing number of devices, such as vacuum cleaners (e.g. Roomba), smart washing machines (e.g. Miele), and smart ovens (e.g. LG Thinq), are becoming more "intelligent" and can be scheduled, either manually or automatically by the ECC, to switch on or off depending on the prices at different times of the day. Such innovations further enable electricity providers to effectively use dynamic pricing to match their cost to revenues by flattening out peak demand and achieving better resource utilization. Enterprise markets offer additional opportunities in addition to consumer markets.

Several earlier works have studied TDP from a user's perspective of scheduling devices according to predictions of future prices. For instance, [1] proposes a mechanism for predicting prices one or two days in advance. Given these prices, household devices can be scheduled so as to balance impatience with the desire to save money. A related paper [2] considers the same problem, but with an emphasis on several users sharing a power source and simultaneously scheduling energy consumption in a distributed manner. A variation on this topic is considered in [3], which introduces an appliance commitment algorithm that schedules thermostatically controlled household loads based on price and consumption forecasts to meet an optimization objective. In this work, we focus on TDP from the energy provider's perspective (i.e., that of increasing profit and minimizing cost), as opposed to the consumer's optimization problem of scheduling his or her power consumption based on price projections.

Electricity providers' problem of determining prices according to user reaction has been studied in several previous works: for instance, [4] reviews the literature up to 2002 on modeling responses to dynamic prices and real trial studies. The more recent work [5] uses real data to quantitatively predict users' scheduling of energy consumption, while [6] considers a feedback loop between users and provider and proposes a real-time pricing algorithm from the perspective of price stability. Other papers such as [7] and [8] consider the total social welfare across users and providers, while [9] specializes a similar model to smart grids. The paper [10] treats the electricity market as an auction, with dynamic offers from providers selling electricity and real-time responses from users buying electricity. Similarly, [11] focuses on users' joint scheduling of energy usage, in the presence of either full or partial information about other users. In this work, we avoid game-theoretic and social welfare models, and instead provide a very practical framework that allows energy providers to set prices by *indirectly* estimating users' device specific scheduling flexibility from time-varying aggregate demand data. Such an approach reduces the required communication overhead and helps ensure scalability of the model.

Many prior papers consider real data in their modeling of user behavior and evaluation of time-dependent pricing. For instance, [12] fits a demand function model to real data. A similar approach is taken in [13], which analyzes the social

Manuscript received 1 October 2011; revised 14 February 2012.

C. Joe-Wong is with the Program in Applied and Computational Mathematics. Princeton University (e-mail: cjoe@princeton.edu).

S. Sen, S. Ha and M. Chiang are with the Department of Electrical Engineering, Princeton University (e-mails: {soumyas, sangtaeh, chiangm}@princeton.edu).

Digital Object Identifier 10.1109/JSAC.2012.120706.

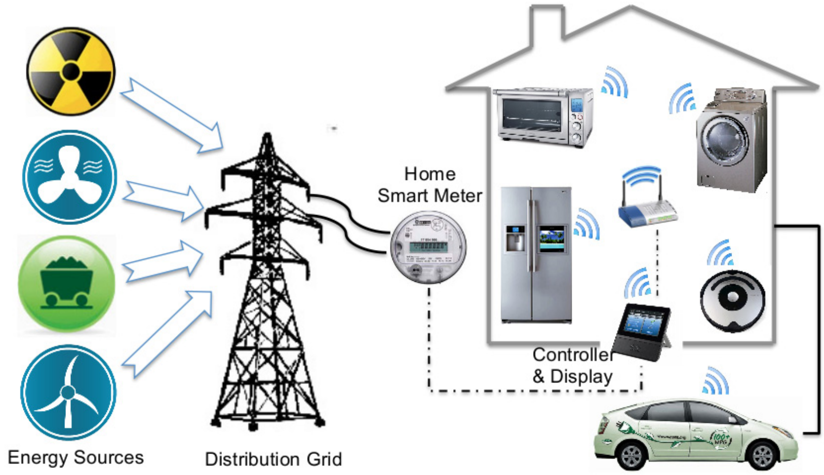

Fig. 1. Schematic of Smart Grid infrastructure with Home Smart Meter controlled devices.

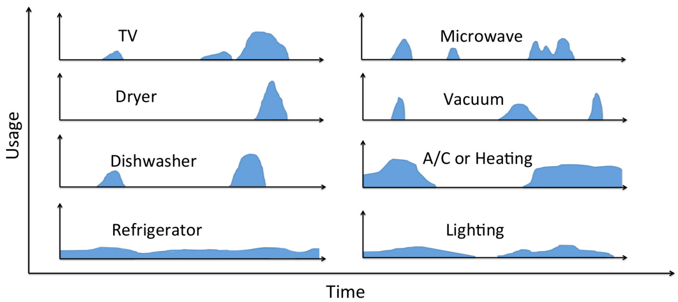

Fig. 2. Schematic of different devices' energy consumption throughout the day.

welfare from peak/off-peak pricing plans using simulations based on real data. TDP's benefits are illustrated in [14], which uses real data to illustrate that shifting usage physically insures against an overloaded network. Finally, the papers [15] and [16] also look at the social welfare gains from actual pricing trials. Table I summarizes the works discussed above.

There are several key differences between our approach and those of previous papers. First, most other works account for the *aggregate demand* across users at different times, but do not consider heterogeneity at the device level. In practice, different devices have very different time sensitivities in shifting their electricity consumption (e.g. smart washing machines and vacuum cleaners like Roombas can typically tolerate a longer delay in scheduling, while smart ovens that schedule cooking times have very little delay tolerance). Figure 2 illustrates this point by showing the time distribution of different devices' energy consumption. Devices like refrigerators consume a roughly constant amount of electricity, and thus cannot shift their usage. Devices like dryers or dishwashers, however, are only used once or twice a day; this usage can easily be shifted to different times of the day, often without user intervention. To account for this device heterogeneity, we develop a time-dependent pricing scheme that builds devices' different delay tolerances into the model of users' demand response.

Second, this work focuses on **day-ahead time-dependent pricing** as opposed to real time pricing, since the latter creates higher uncertainty for the consumer and is less attractive from a user adoption perspective. Consumers and many enterprise customers prefer day-ahead time-dependent pricing as it allows them to plan their activities in advance and also facilitates automated lightweight scheduling of devices by an ECC [22]. We consider the pricing problem of an electricity distributor selling energy directly to consumers, rather than that of an electricity generator.

Third, our formulation of the price optimization problem is shown to be highly tractable and scalable to large numbers of users and pricing periods. Moreover, the formulation relies on predictions of user behavior which can be determined relatively easily from previous observations. Our numerical results show that using our pricing algorithm can help electricity providers realize significant savings by flattening out electricity consumption over the day. The following points summarize our contributions:

- We consider **heterogeneity** in delay tolerances at the **device level** as opposed to modeling a user's energy consumption with utility functions. While modeling with utility functions is theoretically simple, estimating these parameters is typically difficult in practice. Instead, it is more practical for energy providers to monitor usage behavior and estimate users' delay tolerances for different devices through curve fitting on observed demand data as described in Section III.
- The analysis developed allows energy providers to estimate these delay tolerances across users based only on the *aggregate* data, instead of monitoring and estimating these parameters for each individual user. Our methods are thus easily **scalable** to a large number of users. They do not require any additional infrastructure changes, as the provider need only measure the total load on the network in order to calculate the prices.
- Our model allows demand under TDP to remain the same as that with flat-rate pricing. We justify this assumption by supposing that the provider only offers *discounts* from the previous usage-based price; thus, the price in any period will be no more than it is before TDP is introduced. In contrast, several earlier works propose

TABLE I
SUMMARY OF RELATED WORK.

Work	[4]	[5]	[6]	[7]	[8]	[9]	[10]	[11]	[12]	[13]	[14]	[15]	[16]	[17]	[18]	[19]	[20]	[21]
2/3	×	×	×	×	×	×		×	×	×	×		×	×	×	×	×	×
H	×		×		×	×	×	×				×						×
AF	×		×	×	×	×	×	×	×									×
RD	×	×						×	×	×	×	×	×	×	×	×	×	
SW	×		×	×	×	×	×	×		×	×	×	×					×
TR	×	×							×		×	×	×	×	×	×	×	

2/3: 2 or 3 period model H: Hour-long periods AF: Analytical Formulation
RD: Real Data SW: Social Welfare TR: Trial Results

models in which some demand (and associated revenue) will be lost under dynamic pricing.

- We consider the *optimized* price offerings from the **provider's perspective** of maximizing revenue, as opposed to many previous works that have solely focused on maximizing social welfare.
- In contrast to other works, our model incorporates users' *time-shifting* their electricity usage in time; the ECC can schedule devices according to the prices since they are known a day in advance. Most other papers do not consider day-ahead pricing, and do not incorporate baseline usage statistics from before TDP.
- Although we explicitly consider TDP for an *electricity distributor*, we assume that this distributor purchases electricity from a generating company. Distribution grid operators must purchase electricity from generators at different times (peak/off-peak), according to the prices in these periods. This generator is assumed to pass on to the distributor the variation in capital costs associated with different types of energy sources, i.e., base, intermediate and peak costs. We note that if the distributor and generator form part of the same company, then these costs are automatically passed on. Demand and capacity statistics from an energy provider in Ontario are used in the simulations.

The remainder of the paper is organized as follows. Section II introduces our pricing formulation and models of user behavior. We provide an algorithm for determining the optimal day-ahead prices, with the requisite online parameter estimation using the methods described in Section III. We then show our pricing algorithm's numerical results in Section IV. The parameters and calibration data in this section were drawn from real statistics provided by an Ontario operator. Finally, Section V concludes the paper by summarizing the results and suggesting avenues for further research.

II. Model Formulation

In this section, we introduce our model of time-dependent pricing as an optimized feedback loop between users and energy providers. A schematic of the feedback loop is shown in Fig. 3; the energy provider monitors the network load to estimate the consumer's willingness to shift his or her demand, and uses it to announce optimized prices for the next day. The user's response to these prices, via either manual or automated scheduling by the ECC, is then used to refine the provider's estimates of user behavior.

Our model considers a "day-ahead" pricing scheme [14], in which providers publish their prices one day in advance.

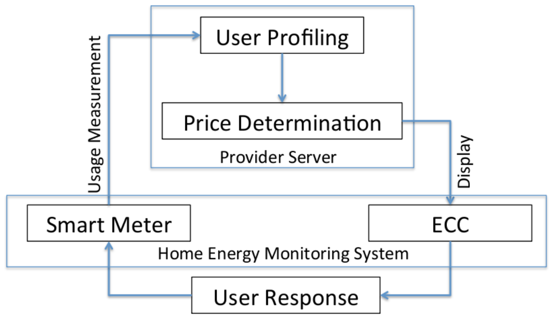

Fig. 3. Schematic of the feedback loop between the provider and users.

This scheme offers users more certainty than other common implementations of dynamic pricing [4], [22], such as "hour-ahead" or real time pricing. With day-ahead pricing, users can schedule their device usage for the upcoming day so as to optimize their amount spent and willingness to shift their device usage. Hour-ahead or real-time pricing would force the ECC to use a less optimal scheduling algorithm to solve an online knapsack problem. We note, however, that our algorithm can be easily adapted to hour-ahead pricing.

We suppose that there are n periods in a day, e.g., $n = 24$ for hourly prices. We assume that the provider faces some given demand in each period, which we call D_i, with i indexing the period. The demand D_i in each period is assumed to be roughly the same each day due to repeated daily patterns in electricity demands (e.g. period 1 has the same demand on Monday, Tuesday, etc.), so that the aggregate demand over each day is usually constant. We verify this assumption using real traces from an Ontario operator of hourly demand data over seven years [23]. Figure 4 shows the hourly demand over three consecutive days; it remains approximately the same from day to day. We use volatility measures from quantitative finance to study the day-to-day volatility of demand for the entire data set [24]. The mean volatility for all hours is found to be between 3.6% and 8.8%, with standard deviations between 1.5% and 5.4%. As shown in Fig. 5, the average daily volatility, measured over each week, is always less than 12%. This result further strengthens the case for cyclic demand patterns over days of the week and shows that the aggregate demand on each day can be assumed to be roughly constant.

The provider's goal is to incentivize users in the right way, so that they shift their energy consumption (e.g., a Roomba) to periods of lower demand. We model user behavior through the shifts in demand from the baseline D_i, induced by time-dependent prices. We assume that no usage is lost with the introduction of TDP, i.e., users consume the same amount

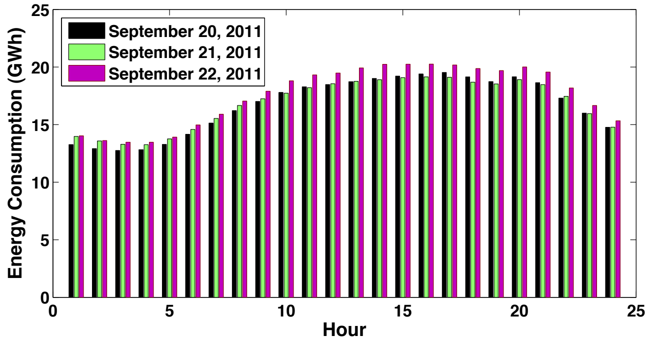

Fig. 4. Hourly electricity demand for IESO remains approximately the same over three consecutive days.

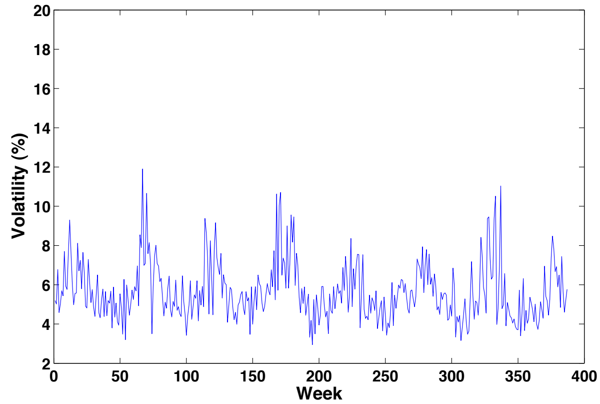

Fig. 5. Average over all hours of usage volatility in each hour, measured over each week.

over the entire day as they did before TDP (the distribution of the usage across different periods is merely shifted and not completely lost when using TDP). This is particularly true when device scheduling is done with an ECC, which automatically schedules the same household devices every day based on the consumer's requirements.

We use R to denote the flat usage-based rate that an energy provider charges to its consumers in the absence of TDP. The time-dependent optimization problem is then formulated in terms of *rewards*, or discounts applied to the nominal flat fee R. The reward in each period i, $i = 1, 2, \ldots, n$, is denoted by r_i. We note that given a day-ahead price offering, users can either delay a device (shift to a later period) or shift to an earlier period, with respect to the device's "original period" (the period in which the device would have been scheduled had there been no time-dependent rewards). The maximum reward offered is taken as R, meaning that the provider never offers negative prices to users.

The usage shifts from devices' original periods are calculated with *sensitivity functions* $s_j(r, t)$, which give each user's probability of shifting each device's usage by the amount of time t, given the reward r in the new period of usage. Sensitivity functions can vary across users and across devices, which we capture through appropriate parameterization of the sensitivity functions. The subscript j denotes this user-device parametrization. The sensitivity functions are assumed to be concave and increasing in reward (in accordance with the principle of diminishing marginal utility) and decreasing in time (i.e., users prefer to shift their usage as little as possible). Since the sensitivity functions are probabilistic, we normalize each s_j by the sum over all times t of $s_j(R, t)$. This normalization ensures that the amount of shifted usage calculated does not exceed the actual usage.

For simplicity, we assume that each period is one unit of time, and index the periods by i, $i = 1, 2, \ldots, n$. Then if a user shifts a device from period i to period k, i.e. shifts by $k - i$ periods, $k - i$ is assumed to be the number $b \in [1, n]$ which is congruent to $k - i$ modulo n. Note that if $k < i$, users have shifted from period i on one day to period k on the next day. We use $|k - i|$ to denote the maximum of $(k - i) \bmod n$ and $(i - k) \bmod n$. In other words, users will either advance their usage to the first period k before period i or delay their usage to the first period k after period i, whichever shifts their usage by a smaller amount of time. Since some devices (e.g. refrigerators) cannot be turned off, we also assume a baseline demand of d_i in each period i; this quantity is the amount of electricity used by devices that never shift their usage.

Using the above notation, we can formulate the electricity provider's optimization problem. The provider wishes to minimize the costs of both offering rewards and satisfying user demand (i.e., having a sufficient available capacity). We next find an expression for the cost of offering rewards.

Given the rewards $r_1, r_2, \ldots, r_n$ in each period, we first calculate the amount of demand shifted out of each period i into each period $k \neq i$. We take the amount of electricity required by each device originally in period i, multiply by the probability of shifting, and sum over all users and their devices to obtain

$$\sum_{j \in \mathcal{D}_i} v_j s_j(r_k, |k - i|), \tag{1}$$

where v_j is the amount of energy required by the user's device j. The set $\mathcal{D}_i$ is the set of all users' devices originally (i.e., before the introduction of TDP) in period i. To find the total amount of demand shifted into period i, we sum over k to obtain

$$\sum_{k \neq i} \sum_{j \in \mathcal{D}_k} v_j s_j(r_i, |k - i|).$$

Then the total cost of offering rewards is calculated as

$$\sum_{i=1}^{n} r_i \sum_{k \neq i} \sum_{j \in \mathcal{D}_k} v_j s_j(r_i, |k - i|) + r_i d_i; \tag{2}$$

the quantity $r_i d_i$ is the loss in revenue from the baseline demand d_i in period i.[1]

Next, we introduce an expression for the cost of satisfying user demand. Electricity generators generally operate multiple plants of different types, e.g. gas, hydroelectric (hydro), nuclear and coal [25], [26]. These plants may be categorized as base-, intermediate-, and peak-load. The base-load plants generally have a higher capital cost but low operating cost, and

[1]Depending on the sensitivity functions used, one may instead account for the baseline demand with an appropriately parameterized sensitivity function, i.e., one that does not allow any usage deferral, no matter the discount offered. To keep the sensitivity functions general, we do not do so here.

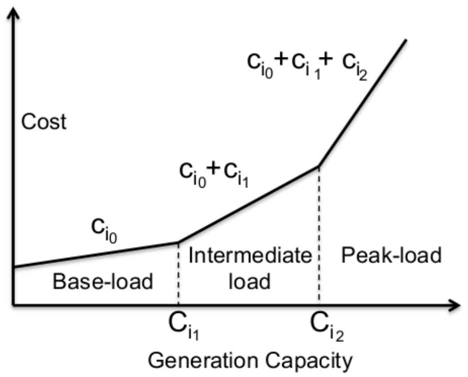

Fig. 6. Piecewise-linear cost structure of base-, intermediate-, and peak-load electricity plants.

thus run all of the time (e.g., nuclear and hydro). Intermediate-load plants (e.g., coal) have a higher operating cost, and peak-load plants (e.g., gas turbines) have the highest operating cost [25]. In any given period, if user demand exceeds the base-load capacity, the generator turns to the intermediate-load plants and then finally to peak-load plants to generate additional electricity. As noted above, we assume that the variation in these generation costs is passed on to the electricity distributor.

We model the cost of base-, intermediate- and peak-load plants linearly, with different marginal costs (slopes). These marginal costs are incurred from fuel and operational costs; while the operational costs can be assumed to be roughly constant, the fuel cost can vary significantly even on a short timescale. Thus, we let c_{i_1} denote the marginal additional cost of using intermediate- rather than base-load plants in period i, and c_{i_2} denote the marginal additional cost of using peak- rather than intermediate-load plants in period i. These marginal costs are instances of the random variables; we assume that their actual values are exogenously determined for use in the provider's optimization problem. Figure 6 shows the piecewise-linear cost structure for base-, intermediate- and peak-load plants; c_{i_0} denotes the slope of base-load electricity generation costs. We assume that any revenue gain from reselling surplus electricity is included in c_{i_0}.

Both plant capacity and the marginal cost of plant operation can vary from period to period. Thus, we use C_{i_1} and C_{i_2} to denote the base- and intermediate-load capacities respectively in period i. Again, these are instances of random variables drawn from exogenous (i.e., price-independent) distributions. Time-series prediction algorithms such as triple-exponential smoothing or auto-regression can be used to estimate the base- and intermediate-load capacities from historical data and exogenous factors [27], [28]. These predictions are then fed into the provider's cost minimization problem, which considers the additional cost from user demand exceeding base- and intermediate-load capacities.

Using (1), the amount of demand in each period i is

$$D_i - \sum_{j\in\mathcal{D}_i}\sum_{k\neq i} v_j s_j(r_k, |k-i|) + \sum_{k\neq i}\sum_{j\in\mathcal{D}_k} v_j s_j(r_i, |k-i|). \quad (3)$$

Then the cost of meeting user demand in each period i is

$$\sum_{l=1,2} c_{i_l}\left[D_i - \sum_{j\in\mathcal{D}_i}\sum_{k\neq i} v_j s_j(r_k, |k-i|) + \sum_{k\neq i}\sum_{j\in\mathcal{D}_k} v_j s_j(r_i, |k-i|) - C_{i_l}\right]^+, \quad (4)$$

where $[y]^+$ signifies the maximum of y and 0. Combining (2) and (4) then yields the following proposition:

Proposition 1: The provider's cost minimization optimization problem is

$$\min_{r_i} \sum_{i=1}^{n} r_i \sum_{k\neq i}\sum_{j\in\mathcal{D}_k} v_j s_j(r_i, |k-i|) + r_i d_i + \sum_{l=1,2} c_{i_l}\left[D_i - \sum_{k\neq i}\left(\sum_{j\in\mathcal{D}_i} v_j s_j(r_k, |k-i|) + \sum_{j\in\mathcal{D}_k} v_j s_j(r_i, |k-i|)\right) - C_{i_l}\right]^+ \quad (5)$$

$$\text{s.t. } r_i \geq 0,\ i = 1, 2, \ldots, n \quad (6)$$

$$\text{var. } r_i,\ i = 1, 2, \ldots, n \quad (7)$$

This optimization problem (5-7) is easily solvable even with large numbers of users and periods:

Proposition 2: The optimization problem (5-7) is a convex optimization problem, assuming that the sensitivity functions $s_j(r, t)$ are concave and increasing in r and decreasing in t.

Proof: See the Appendix. ■

We can get a sense of the range of rewards offered by assuming that the sensitivity functions are linear in rewards. In that case, taking the derivative of the objective function (5) with respect to the reward r_i yields

$$d_i + 2r_i \sum_{k\neq i}\sum_{j\in\mathcal{D}_k} \frac{\partial s_j}{\partial r_i}(r_i, |i-k|) v_j - \sum_{k\neq i}(c_{k_1} + c_{k_2}) \sum_{j\in\mathcal{D}_k} v_j \frac{\partial s_j}{\partial r_i}(r_i, |i-k|). \quad (8)$$

Setting this quantity equal to zero, we obtain the following proposition:

Proposition 3: If the sensitivity functions are linear in reward, then the maximum possible reward in period i is

$$\frac{\sum_{k\neq i}(c_{k_1} + c_{k_2}) \sum_{j\in\mathcal{D}_k} v_j \frac{\partial s_j}{\partial r_i}(r_i, |i-k|) - d_i}{2\sum_{k\neq i}\sum_{j\in\mathcal{D}_k} v_j \frac{\partial s_j}{\partial r_i}(r_i, |i-k|)}$$

$$\leq \frac{\max_k (c_{k_1} + c_{k_2})}{2} - \frac{d_i}{2\sum_{k\neq i}\sum_{j\in\mathcal{D}_k} v_j \frac{\partial s_j}{\partial r_i}(r_i, |i-k|)} \quad (9)$$

$$\leq \frac{\max_k (c_{k_1} + c_{k_2})}{2}. \quad (10)$$

Since the capacity and sensitivity functions change from day to day, the provider's cost-minimizing prices will also change. The provider thus requires a dynamic algorithm to adjust the prices offered in response to these changing optimization parameters. A simple adaptation of the optimization problem (5-7), inspired by a dynamic programming approach, is given in Algorithm 1.

Algorithm 1 COMPUTING DAY-AHEAD PRICES

```
Estimate the capacities for the next n periods.
Solve (5-7) for the optimal rewards using the capacities estimated in the previous step.
while the provider runs time-dependent pricing do
    Estimate the generating capacity for the next n periods.
    Solve (5-7) for the reward n periods in advance using the capacities estimated and taking the previously computed rewards for the next n − 1 periods as given.
    if it is period n then
        Estimate the sensitivity functions using the methods in Section III.
    end if
end while
```

Using Algorithm 1, the provider would initialize TDP by forecasting each period's capacity over the next day and computing the corresponding prices. For instance, if the provider initiates time-dependent pricing in period 1 on day 1, the rewards (equivalently, the prices) for periods 1 to n on day 1 are calculated just before this initial period 1. At the end of period 1 on this first day, the provider would estimate the generating capacity for period 1 on the next day. The provider then computes the reward for period 1 of the next day by optimizing the expected cost from period 2 on the first day to period 1 on the second day. The rewards for periods 2 to n on the first day are taken as given, as they were computed in the previous iteration. These estimation-calculation steps of calculating one new day-ahead reward in each period are repeated for as long as time-dependent pricing is run. Estimates of the underlying sensitivity functions can be refined as described in Section III below.

III. SENSITIVITY FUNCTION ESTIMATION

As discussed in Section II, implementing our pricing algorithm requires the electricity provider to estimate users' sensitivity functions. This section proposes a method for doing so, based on best-fit curves. We emphasize that our method uses only *aggregate* usage data; thus, the electricity provider need not keep track of individual usage (e.g. from smart meters) in real time.

To estimate the sensitivity functions, we assume a given functional form with tunable parameters. For instance, we use the following function in which the consumer's utility increases in reward amount but decreases exponentially in time:

$$s_\beta(r,t) = C_\beta \frac{r}{(t+1)^\beta}, \tag{11}$$

where C_β is a normalization constant depending on β, the adjustable parameter. For ease of notation, we refer to β as the patience index, though β actually parameterizes users' willingness to shift a device's usage to either before or after its original period. A higher value of β indicates a more impatient user, who is less willing to shift his or her usage for longer periods of time. Figure 7 shows the sensitivity function versus the duration of shifting and percent reward for different values of β; we see that for a lower value of β, the sensitivity function decreases much more slowly with the duration of shifting. Figure 8 gives an intuitive understanding of other values of β by showing the average duration of shifting, assuming flat-rate pricing. We stress that these sensitivity functions are simply mathematical approximations to user behavior; thus, while users' behavior may actually follow a highly nonlinear, non-differentiable pattern, we assume that this behavior can be well-approximated with functions such as those in (11).

The goal of the estimation algorithm is to estimate the particular values of β for different users and devices. For ease of notation, we can include all values of β for a given period into one *aggregate sensitivity function*, by summing the sensitivity functions for all devices in that period, weighted by the proportion of electricity consumption due to each device. Thus if α_j denotes the proportion of electricity consumption due to device j, we have the aggregate sensitivity function in period i:

$$S_i(r,t) = \sum_{j \in \mathcal{D}_i} \alpha_j s_j(r,t). \tag{12}$$

We note that the amount of usage deferred from period i to period k is then

$$B_{i,k} = (D_i - d_i) S_i(r_k, |k-i|)$$

and that given a set of rewards and times shifted, S_i depends entirely on the parameters α_j and β_j, where the β_j are patience parameters. We can group devices with similar patience parameters together, so that they become one $\alpha_j s_j(r,t)$ term in (12). Note that $D_i - d_i$ is the demand in period i that can actually be shifted.

Let A_i denote the difference between the demand D_i without time-dependent pricing and the demand with time-dependent pricing. Given a set of n prices over one day, each A_i can be expressed with the aggregate sensitivity functions:

$$A_i = \sum_{k \neq i} B_{k,i} - B_{i,k}.$$

Thus, each of the A_i, $i = 1, 2, \ldots, n$ can be expressed as a linear function of the $\frac{n(n-1)}{2}$ functions $B_{k,i}$. Starting from the expression for A_1 and continuing to A_{n-1}, these n linear equations can be sequentially solved for each $B_{1,2}, B_{2,3}, \ldots B_{n-1,n}$. One linear function of $(n-1)\left(\frac{n}{2}-1\right)$ $B_{i,k}$ variables then remains. We note that for each period i, $B_{i,k}$ appears in this expression for some value of k. This linear function can be numerically evaluated for each day in the provider's dataset. We can then use a best-fit algorithm (e.g. nonlinear least-squares) to estimate the parameters α_j and β_j which determine the $B_{i,k}$ variables and give the desired sensitivity functions. We note that this estimation algorithm can run offline every hour or every day and need not be done in real time.

We test the efficiency of this algorithm by specifying patience indices as in (11) for three periods with two sensitivity functions each. We calculate the amounts of usage shifted, randomly perturb them by an average of 10%, and estimate

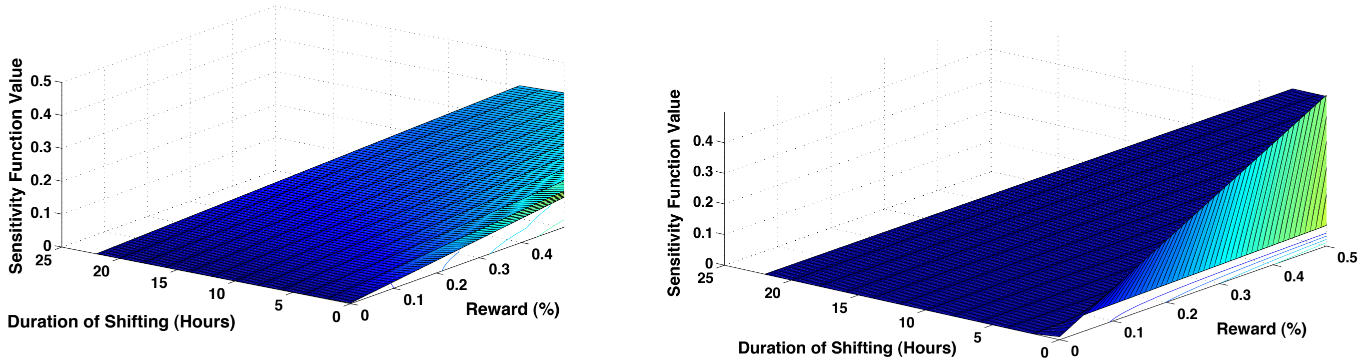

(a) Sensitivity function for $\beta = 0.5$ (a patient user). (b) Sensitivity function for $\beta = 5$ (an impatient user).

Fig. 7. Sensitivity function value versus duration of shifting and reward for different values of β. As β increases, a higher reward is necessary to maintain a given sensitivity function value.

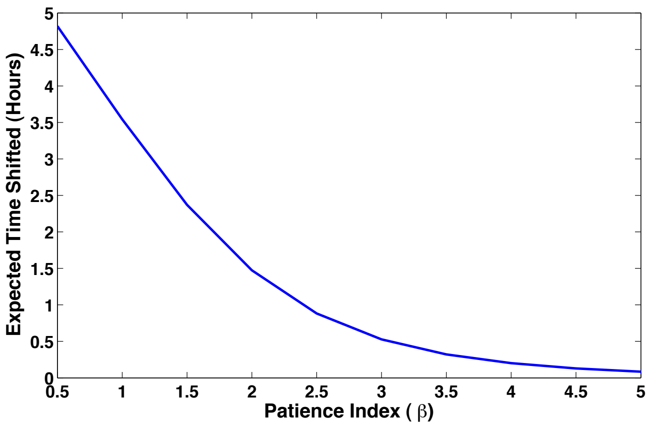

Fig. 8. Expected durations of shifting for several values of β under a uniform pricing scheme.

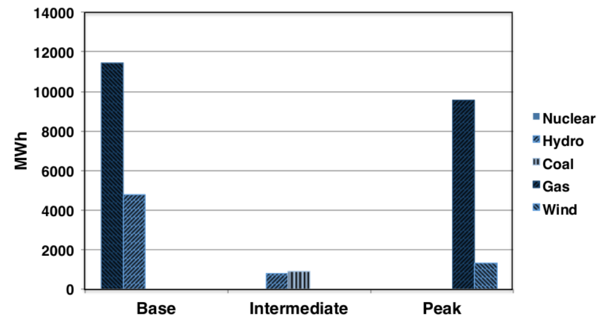

Fig. 9. Distribution of electricity supply sources.

TABLE II
ACTUAL AND ESTIMATED PARAMETER VALUES IN SIMULATION OF SENSITIVITY FUNCTION ESTIMATION.

Period	Actual Values			Estimated Values			Maximum Percent Error
	β_1	β_2	α_1	β_1	β_2	α_1	
1	1	2	0.17	1.03	2.48	0.46	11.8
2	1	2.33	0.5	1.02	2.49	0.45	9.0
3	1	2.67	0.83	0.90	2.15	0.71	0.5

the sensitivity function parameters from this noisy simulated data. The estimation errors between the estimated and actual parameter values for the sensitivity functions are shown in Table II. The maximum percent error is the percent error in using the estimated instead of actual parameters to evaluate the sensitivity functions.

IV. SIMULATIONS

In this section, we use numerical simulations with realistic parameters to show the feasibility of our pricing algorithm and demonstrate its ability to reduce the provider's cost of generating electricity and flatten electricity usage over the day. The parameters are based on real data from the Ontario Independent Electricity Systems Operator (hereafter referred to as IESO). We note that while our results will likely be qualitatively accurate in other markets, the Ontario market has several unique features. For instance, Ontario uses more nuclear energy than any other Canadian province [29], [30].

The base-load plants in Ontario are both nuclear and hydroelectric; while all nuclear plants are treated as base-load, only some hydroelectric plants are base-load (we assume 60% are base-load plants) [25]. The production capacity of each plant is taken as constant across different periods of a day for the purposes of simulation. The intermediate-load consists of coal (operating at 20% efficiency, as is consistent with the data in [25]) and the remaining hydroelectric plants. Finally, the peak plants are gas turbines, which are the most expensive to operate [26], [31]. The distribution of energy supply from different sources is shown Fig. 9. The slopes of the cost functions for base-, intermediate- and peak-load plants (refer to Fig. 6) are taken from the production estimates in [31]. The marginal costs of moving from intermediate- to peak-load and base- to intermediate-load plants are calculated to be \$62.46/MWh and \$18.54/MWh respectively. For simplicity, we assume that the electricity generator charges the distributor these prices; in practice, a premium would be added to the price that the electricity generator charges the distributor.

We assume that the maximum price (i.e., that before TDP is introduced) is \$110/MWh, which is estimated from the average peak price over the past several years offered by IESO [23]. We use a period length of one hour (i.e., 24 periods in a day). The demand prior to TDP is obtained by averaging the hourly demand experienced by IESO over each day of the past several months [23], perturbed up to 5% by

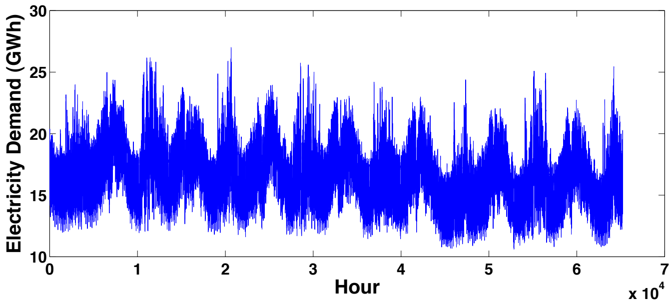

Fig. 10. Hourly electricity demand for IESO over the past 7 years.

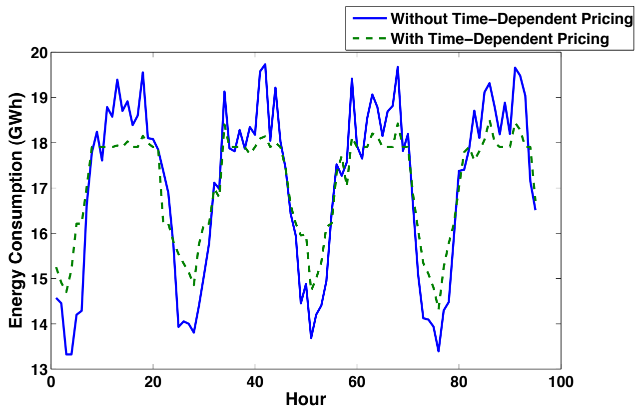

Fig. 11. Energy consumption over four days, with and without optimized TDP rewards.

a uniformly distributed random number. Figure 10 shows this data in graphical form.

The aggregate sensitivity function in each period is taken as a sum of functions of the form (11), with $\beta = 0.5, 1, \ldots, 5$. The α_j are chosen to intuitively reflect a reasonable distribution of users' willingness to shift devices. We assume a baseline demand (i.e., amount of electricity usage that cannot be shifted) of 1 GWh in each hour across all customers.[2]

Figure 11 shows the comparison in hourly electricity consumption with and without time-dependent pricing over the first four days of TDP. We see that consumption is reduced from a peak of 19.8 GWh to a peak of 18.5 GWh, which is only slightly above the intermediate-load capacity of 17.9 GWh. However, the demand curve is not fully flat because some devices (e.g. microwaves, lights) are time-sensitive, and hence much of the residual unevenness is to be expected.

The flattening of electricity usage helps reduce the cost of generating electricity. As we observe in Fig. 11, the rewards are calculated so as to bring electricity consumption below the intermediate-load capacity as much as possible. In contrast, the cost of moving from base- to intermediate-load plants is small enough that energy consumption in low-demand periods is not brought up to fully utilize the base-load capacity.

Figure 12 shows the rewards corresponding to Fig. 11's energy consumption pattern. We see that the rewards (discounts) are roughly cyclical, as might be expected, and that they are zero in peak periods. With the offered rewards shown,

[2]Due to the form of sensitivity functions, this baseline demand can also be interpreted as having sensitivity functions described in (11), with $\beta = \infty$.

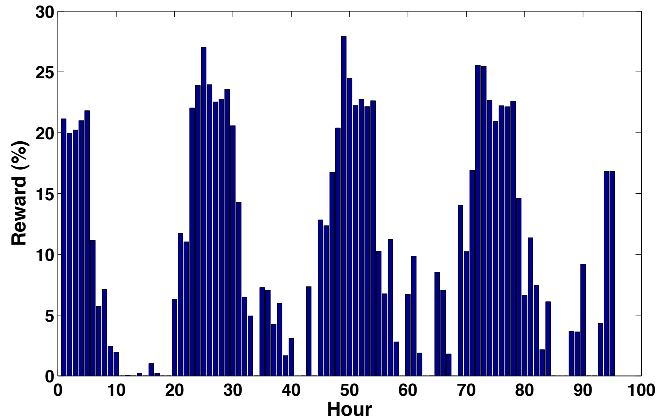

Fig. 12. Optimized rewards over four days, which yield Fig. 11's energy consumption pattern. The nominal price in Ontario is approximately 11¢/kWh.

the electricity provider's cost decreases from \$4624 to \$3396 as per the objective function in (5). Thus, offering rewards can reduce the electricity provider's cost by 27%. While the quantitative results of these simulations will vary from market to market, the qualitative results suggest that time-dependent pricing can indeed help electricity providers to even out consumption over the day and reduce the energy requirements from peak-load plants.

To illustrate our results for individual users and devices, we next simulate the behavior of two users with patience indices for different devices as listed in Table III. The baseline usage is taken to be lighting, with a corresponding patience index of ∞. User 1 is generally more patient than user 2, so we expect that user 1's devices will shift for a longer amount of time than user 2's. The device usage before TDP is based on energy consumption for typical devices in [32]. Figures 13a and 13c show the initial distribution of usage by device for users 1 and 2 respectively.

We next calculate the probability that each device will shift its usage, based on the patience indices given in Table III and the rewards for the fourth day calculated above. We then use this probability distribution to choose the period to which each device is shifted (if at all). The resulting distribution of usage by device is shown in Figs. 13b and 13d for users 1 and 2 respectively. The more patient user 1 shifts dryer, vacuum and entertainment usage, for as many as 10 hours for the dryer. User 2, however, shifts only vacuum usage, and that for only 3 hours. User 1 thus saves more from shifting (26%) than user 2 (16%).

As our model is probabilistic, other users, even those with the same patience indices, may have different results from those in Figs. 13b and 13d. To account for this variation, Figs. 13e and 13f respectively show one thousand users' total energy consumption with and without TDP. Half the users are as patient as user 1 above, and the other half are impatient like user 2. We see that the peak consumption is greatly reduced with TDP, from 2700 to 1900 GWh, i.e., almost a 30% reduction in peak usage. Moreover, the distribution of energy consumption over the day is visibly much flatter with TDP than without it. Indeed, the peak-to-average ratio of electricity

TABLE III
PATIENCE INDICES FOR BOTH USERS AND SEVERAL DEVICES.

User	Air Conditioning	Vacuum	Dryer	Television	Lighting
1	4	0.5	1	2	∞
2	4	1	2	4	∞
Lighting's infinite patience index means it is never shifted.					

usage decreases from 2.55 without TDP to 1.88 with TDP.

V. CONCLUSION

This paper introduces a new model for smart grid TDP, and in particular accounts for device-level heterogeneity in delay tolerances. By directly modeling users' willingness to shift energy consumption to lower-price periods, we formulate a highly tractable optimization problem to determine cost-minimizing time-dependent prices for electricity providers. Since these prices incentivize users to shift their energy consumption, we introduce a complete feedback loop between users and providers, allowing real-time estimates of user behavior and corresponding adjustments to the prices offered. Numerical results indicate that our optimized prices can help electricity providers significantly reduce electricity generation costs. We demonstrate these results for the estimated energy consumption and the offered rewards by using realistic parameters to simulate price generation for consecutive days.

The model and ideas presented in this work can be applied to several variations on time-dependent pricing. For instance, TDP can be used to determine when to obtain and store electricity, e.g. in batteries, which can help flatten demand [33]. One can also consider two sets of time-dependent prices: one for providers selling electricity to users, and one for users selling back individually-generated renewable energy, for instance from photovoltaic cells. Such ideas are gaining traction through such policies as feed-in tariffs [34].

ACKNOWLEDGEMENT

This work was in part supported by NSF grants CNS-0905086 and CNS-1117126, a Google research grant and a Princeton University Grand Challenges grant. C.J.-W. was supported by an NDSEG fellowship.

APPENDIX
PROOF OF PROP. 2

The main idea of the proof is to take the second derivative of the objective function (5) and show that it is positive-definite in the range of feasible rewards (i.e., when the marginal cost of offering the reward is lower than the reward itself). Moreover, this range of feasible rewards is a convex set. For simplicity, we assume one device requiring one unit of electricity in each period, with corresponding sensitivity function s_i in period i. Moreover, we suppress the time-dependence of the s_i.

For ease of notation, we denote the function $c_{i_l} \max(x, 0)$ by f_{i_l}. Moreover, the amount of demand

$$D_i - \sum_{k \neq i} s_i(r_k) + \sum_{k \neq i} s_k(r_i)$$

in each period is denoted by x_i.

Taking the first derivative of (5) yields

$$\sum_{k \neq i} (r_i s_k'(r_i) + s_k(r_i)) + d_i + \sum_{l=1,2} \sum_{k \neq i} f_{i_l}'(x_i - C_{i_l}) s_k'(r_i) - \sum_{l=1,2} \sum_{k \neq i} f_{k_l}'(x_k - C_{k_l}) s_k'(r_i). \quad (13)$$

We note that each f_{i_l}' is a constant, as the provider operates on a linear segment of the piecewise-linear function f_{i_l}. Thus, each f_{i_l}' is independent of the r_i; then the second derivative matrix of (5) is diagonal, with the ith entry the derivative of (13):

$$\sum_{k \neq i} (2s_k'(r_i) + r_i s_k''(r_i)) + \sum_{l=1,2} \sum_{k \neq i} f_{i_l}'(x_i - C_{i_l}) s_k''(r_i) - \sum_{l=1,2} \sum_{k \neq i} f_{k_l}'(x_k - C_{k_l}) s_k''(r_i). \quad (14)$$

It suffices to show that (14) is positive at for each i. We can regroup (14) as

$$\sum_{k \neq i} 2s_k'(r_i) + \left(r_i + \sum_{l=1,2} \sum_{k \neq i} f_{i_l}'(x_i - C_{i_l}) - f_{k_l}'(x_k - C_{k_l}) \right) s_k''(r_i).$$

Since the s_k are increasing, $\sum_{k \neq i} 2s_k'(r_i) > 0$. We thus must show that

$$r_i \leq \sum_{l=1,2} \sum_{k \neq i} f_{k_l}'(x_k - C_{k_l}) - f_{i_l}'(x_i - C_{i_l}),$$

which must be true, as the right-hand side is the marginal benefit of offering a reward, and r_i is the reward offered. ■

REFERENCES

[1] 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, 2010.

[2] A. H. Mohsenian-Rad, V. W. S. Wong, J. Jatskevich, and R. Schober, "Optimal and autonomous incentive-based energy consumption scheduling algorithm for smart grid," in *Innovative Smart Grid Technologies 2010*. IEEE, 2010, pp. 1–6.

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

[4] S. Borenstein, M. Jaske, and A. Rosenfeld, "Dynamic pricing, advanced metering, and demand response in electricity markets," *Center for the Study of Energy Markets*, 2002, Working Paper.

[5] A. Faruqui and L. Wood, "Quantifying the benefits of dynamic pricing in the mass market," Edison Electric Institute, Tech. Rep., 2008.

[6] M. Roozbehani, M. Dahleh, and S. Mitter, "Dynamic pricing and stabilization of supply and demand in modern electric power grids," in *2010 First IEEE International Conference on Smart Grid Communications*. IEEE, 2010, pp. 543–548.

[7] G. Brunekreeft, "Price capping and peak-load-pricing in network industries," *Diskussionsbeiträge des Instituts für Verkehrswissenschaft und Regionalpolitik, Universität Freiburg*, vol. 73, 2000.

[8] H. Chao, "Peak load pricing and capacity planning with demand and supply uncertainty," *The Bell Journal of Economics*, vol. 14, no. 1, pp. 179–190, 1983.

[9] P. Samadi, A. Mohsenian-Rad, R. Schober, V. W. S. Wong, and J. Jatskevich, "Optimal real-time pricing algorithm based on utility maximization for smart grid," in *2010 First IEEE International Conference on Smart Grid Communications*. IEEE, 2010, pp. 415–420.

[10] P. Vytelingum, S. D. Ramchurn, T. D. Voice, A. Rogers, and N. R. Jennings, "Trading agents for the smart electricity grid," in *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 2010, pp. 897–904.

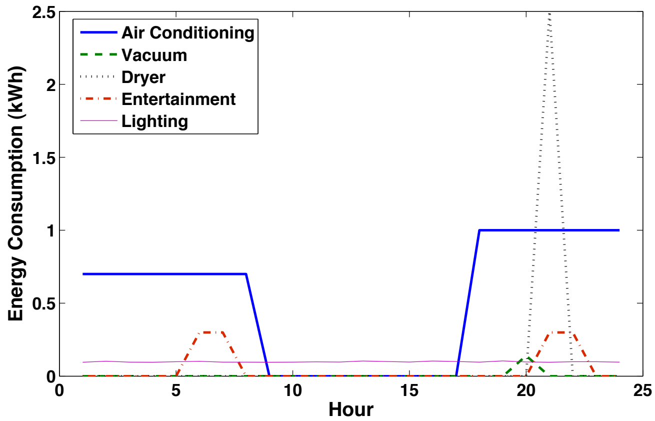

(a) User 1 (a patient user), before TDP.

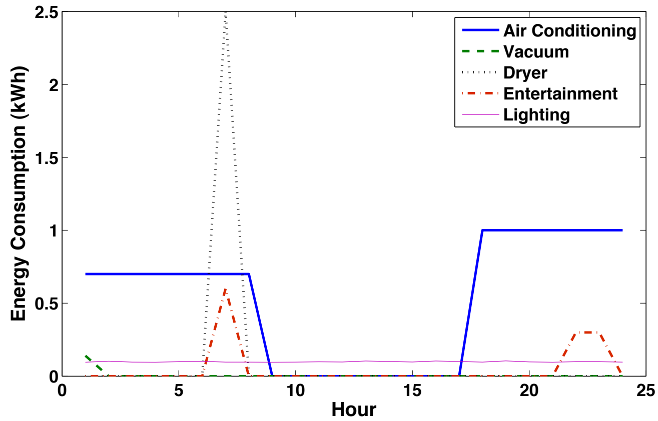

(b) User 1, with TDP.

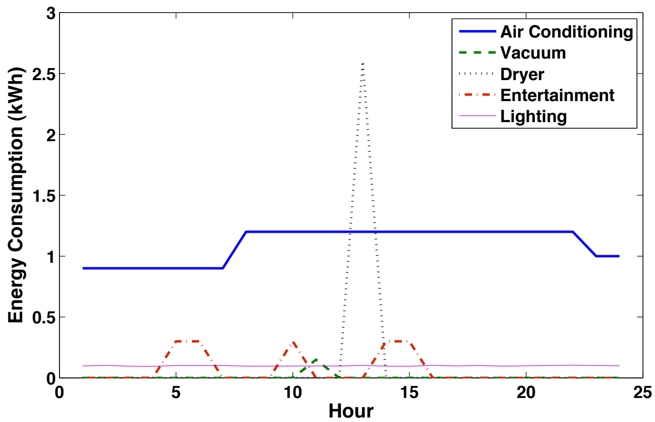

(c) User 2 (an impatient user), before TDP.

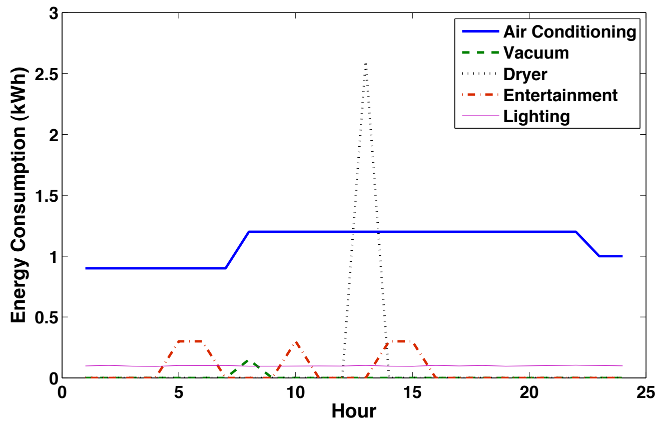

(d) User 2, with TDP.

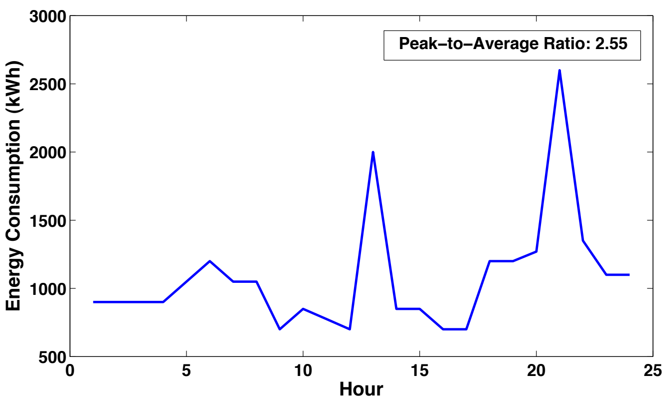

(e) Aggregated over 1000 users, before TDP.

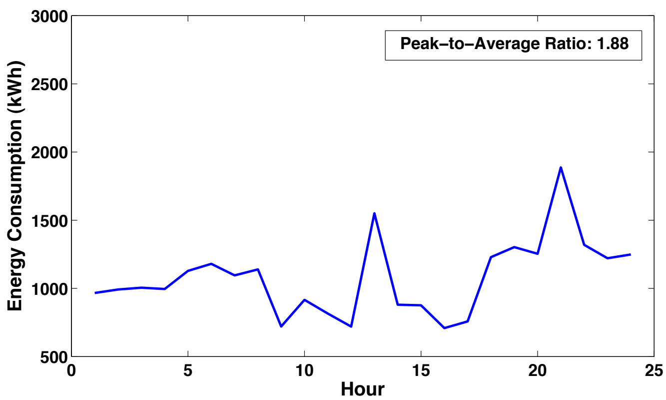

(f) Aggregated over 1000 users, with TDP.

Fig. 13. Energy consumption over one day, before and after time-dependent pricing (TDP) is introduced.

[11] S. Caron and G. Kesidis, "Incentive-based energy consumption scheduling algorithms for the smart grid," in *2010 First IEEE International Conference on Smart Grid Communications*. IEEE, 2010, pp. 391–396.

[12] J. A. Hausmann, M. Kinnucan, and D. McFaddden, "A two-level electricity demand model: Evaluation of the Connectict time-of-day pricing test," *Journal of Econometrics*, vol. 10, no. 3, pp. 263–289, 1979.

[13] S. Borenstein, "The long-run efficiency of real-time electricity pricing," *The Energy Journal*, vol. 26, no. 3, pp. 93–116, 2005.

[14] E. Hirst, "The financial and physical insurance benefits of price-responsive demand," *The Electricity Journal*, vol. 15, no. 4, pp. 66–73, 2002.

[15] J. Wells and D. Haas, *Electricity Markets: Consumers Could Benefit from Demand Programs, but Challenges Remain*. Darby, PA: DIANE Publishing, 2004.

[16] I. B. M. Global Business Services and eMeter Strategic Consulting, "Ontario Energy Board Smart Price Pilot final report," Ontario Energy Board, Tech. Rep., 2007.

[17] Charles River Associates, "Impact evaluation of the California statewide pricing pilot," Charles River Associates, Tech. Rep., 2005. [Online]. Available: http://www.calmac.org/publications/2005-03-24_SPP_FINAL_REP.pdf

[18] K. Herter, "Residential implementation of critical-peak pricing of elec-

tricity," *Energy Policy*, vol. 35, no. 4, pp. 2121–2130, 2007.
[19] I. Matsukawa, "Household response to optional peak-load pricing of electricity," *Journal of Regulatory Economics*, vol. 20, no. 3, pp. 249–267, 2001.
[20] F. A. Wolak, "Residential customer response to real-time pricing: the Anaheim critical-peak pricing experiment," Stanford University, Tech. Rep., 2006. [Online]. Available: http://www.stanford.edu/wolak
[21] N. Li, L. Chen, and S. Low, "Optimal demand response based on utility maximization in power networks," *IEEE Power and Energy Society General Meeting*, 2011.
[22] M. Scallet, "Power smart pricing blog," 2009, http://powersmartpricing.blogspot.com/2009/05/day-ahead-vs-real-time-pricing.html.
[23] IESO Ontario, "Market data," Sep. 2011, http://www.ieso.ca/imoweb/marketdata/marketData.asp.
[24] J. Liu, C. Tse, and K. He, "Fierce stock market fluctuation disrupts scalefree distribution," *Quantitative Finance*, vol. 11, no. 6, pp. 817–823, 2011.
[25] IESO Ontario, "Monthly generator disclosure report," Sep. 2011, http://www.ieso.ca/imoweb/marketdata/genDisclosure.asp.
[26] G. Masters, *Renewable and Efficient Electric Power Systems*. Wiley Online Library, 2004.
[27] J. Taylor, "Triple seasonal methods for short-term electricity demand forecasting," *European Journal of Operational Research*, vol. 204, no. 1, pp. 139–152, 2010.
[28] R. Weron and A. Misiorek, "Forecasting spot electricity prices: A comparison of parametric and semiparametric time series models," *International Journal of Forecasting*, vol. 24, no. 4, pp. 744–763, 2008.
[29] IESO Ontario, "Supply overview," Sep. 2011, http://www.ieso.ca/imoweb/media/md_supply.asp.
[30] Statistics Canada: Manufacturing and Energy Division, "Electric power generation, transmission and distribution," 2007, http://www.statcan.gc.ca/pub/57-202-x/57-202-x2007000-eng.pdf.
[31] J. Morgan, "Comparing energy costs of nuclear, coal, gas, wind and solar," 2010, http://nuclearfissionary.com/2010/04/02/comparing-energy-costs-of-nuclear-coal-gas-wind-and-solar/.
[32] City Water, Light and Power, "Appliance energy use chart," 2010, http://www.cwlp.com/energy_services/eso_services_programs/appliance_energy_use_chart.htm.
[33] P. Vytelingum, T. D. Voice, S. D. Ramchurn, A. Rogers, and N. R. Jennings, "Intelligent agents for the smart grid," in *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 2010, pp. 1649–1650.
[34] T. Couture and Y. Gagnon, "An analysis of feed-in tariff remuneration models: Implications for renewable energy investment," *Energy Policy*, vol. 38, no. 2, pp. 955–965, 2010.

Carlee Joe-Wong Carlee Joe-Wong (S'11) received her A.B. (magna cum laude) in mathematics from Princeton University in 2011, where she was awarded the George B. Wood Legacy Junior Prize. In 2011, she received the NSF Graduate Research Fellowship and the National Defense Science and Engineering Graduate Fellowship. She is currently a first-year graduate student in the Program in Applied and Computational Mathematics at Princeton University, where she received the Graduate School's Centennial Fellowship and is actively publishing in IEEE and ACM journals and conferences. Her research interests include network economics and optimal control theory.

Soumya Sen Soumya Sen (S'10, M'11) received the B.E. (Hons.) in electronics and instrumentation engineering from BITS-Pilani, India, in 2005, and both the M.S. and Ph.D. in electrical and systems engineering from the University of Pennsylvania in 2008 and 2011, respectively. He is currently a Postdoctoral Research Associate at the Princeton University. He has been actively involved in forging various industrial and academic partnerships, and served as a reviewer for several prestigious ACM and IEEE journals and conferences. His research interests are in Internet economics, e-commerce, social networks, network security, and communication network architectures and protocols.

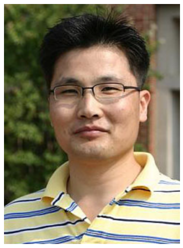

Sangtae Ha Sangtae Ha (S'07, M'09) is an Associate Research Scholar in the Department of Electrical Engineering at Princeton University, leading the establishment of the Princeton EDGE Lab, as its Associate Director. He received his Ph.D. in Computer Science from North Carolina State University in 2009 and has been an active contributor to the Linux kernel. During his Ph.D. years, he participated in inventing CUBIC, a new TCP congestion control algorithm. Since 2006, CUBIC has been the default TCP algorithm for Linux and is currently being used by more than 40% of Internet servers around the world and by several tens millions Linux users for daily Internet communication. Currently, he is working hard to make his pricing research into industry adoption. His research interests include pricing, greening, cloud storage, congestion control, peer-to-peer networking, and wireless networks.

Mung Chiang Mung Chiang (S'00, M'03, SM'08, F'12) is a Professor of Electrical Engineering at Princeton University, and an affiliated faculty in Applied and Computational Mathematics, and in Computer Science. He received his B.S. (Hons.), M.S., and Ph.D. degrees from Stanford University in 1999, 2000, and 2003, respectively, and was an Assistant Professor 2003-2008 and an Associate Professor 2008-2011 at Princeton University. His research on networking received the 2012 IEEE Kiyo Tomiyasu Award, a 2008 U.S. Presidential Early Career Award for Scientists and Engineers, several young investigator awards and several paper awards. His inventions resulted in a few technology transfers to commercial adoption, and he received a Technology Review TR35 Award in 2007 and founded the Princeton EDGE Lab in 2009. He serves as an IEEE Communications Society Distinguished Lecturer in 2012-2013, and is writing an undergraduate textbook "20 Questions About the Networked Life".