Control Mechanisms for Residential Electricity Demand in SmartGrids

Shalinee Kishore
Department of Electrical and Computer Engineering
Lehigh University, Bethlehem, PA 18015
Email: skishore@lehigh.edu

Lawrence V. Snyder
Department of Industrial and Systems Engineering
Lehigh University, Bethlehem, PA 18015
Email: lvs2@lehigh.edu

***Abstract*—We consider mechanisms to optimize electricity consumption both within a home and across multiple homes in a neighborhood. The homes are assumed to use energy management controllers (EMCs) to control the operation of some of their appliances. EMCs, which are a feature of the emerging SmartGrid, use both prices and user preferences to control power usage across the home. We first present a simple optimization model for determining the timing of appliance operation to take advantage of lower electricity rates during off-peak periods. We then demonstrate, using simulation, that the resulting solution may in fact be more peaky than the "non-scheduled" solution, thereby negating some of the benefits (for the utility) of off-peak pricing models. We then propose a distributed scheduling mechanism to reduce peak demand within a neighborhood of homes. The mechanism provides homes a guaranteed base level of power and allows them to compete for additional power to meet their needs. Finally, we introduce a more powerful EMC optimization model, based on dynamic programming, which, unlike our first optimization model, accounts for the potential for electricity capacity constraints.**

I. Introduction

The emerging SmartGrid will provide residential users flexibility in controlling their electricity costs. A primary driving force is the *smart meter*, which can deliver "real-time" electricity prices to homes, potentially every fifteen minutes [1]. The customer can make use of this information via an in-home *energy management controller* (EMC), which uses both prices and user preferences to control power usage across the home. The EMC may be standalone or embedded either in the smart meter or in appliances. At the same time, customers may participate in Direct Load Control (DLC) that allows utilities to control some power consumption within a home during peak usage (thereby bypassing or even replacing the EMC) as a part of an energy savings subscription plan.

Such methods for controlling electricity consumption are part of *demand response*, which relies on varying the price of electricity throughout the day in order to reduce peak demand. Reduced peak demand lowers electricity bills and benefits utilities by reducing complexity of grid stability, equipment overload, brownouts, and blackouts. It also enables utilities to comply with government mandates to cut peak demand, e.g., [2].

In this paper, we consider mechanisms to optimize electricity consumption both within a home containing *smart appliances* and across multiple such homes in a neighborhood. At the in-home level, we assume that appliances communicate with one another and with the EMC over a home-area network (HAN). Many home appliances offer some degree of flexibility regarding when they are operated. For example, homeowners may be indifferent as to when the dishwasher or clothes dryer run, as long as the cycle is complete by a prescribed time. We present a simple optimization model that exploits this flexibility in order to determine the optimal timing of appliance operation, taking into account both electricity price fluctuations and user preferences. This optimization model represents the logic embedded in a simple EMC and aims to minimize a customer's cost, measured as a combination of the dollars paid to the utility and the inconvenience of delayed appliance use. Using simulation, we demonstrate that, if multiple homes each optimize their appliance usage to take advantage of off-peak energy prices, the problem of demand peaks is not alleviated; rather, the demand peak simply shifts to the off-peak period, creating a new "rebound" peak [3] that is even more exaggerated. Thus, from the perspective of the utility, this optimal solution (for the customer) reduces the effectiveness of off-peak pricing models.

Next, we propose a distributed scheduling mechanism to reduce peak demand across a collection of local homes. Our scheme relies on a common control channel among a collection of homes that permits EMCs in each home to communicate with each other. (The security and privacy concerns of exchanging power consumption data across homes is beyond the scope of this paper but is an important consideration in the practical implementation of this approach. Further, the control channel required for this scheme can be derived from the Automated Meter Reading Infrastructure.) The utility chooses a maximum allowable peak power consumption level for the homes while also guaranteeing a minimum available level of power for each home at all times. If a home wishes to use more power than its minimum guaranteed level, then its EMC uses our proposed mechanism to contend with other homes' EMCs for the remaining available power. The proposed distributed scheduling scheme is designed to ensure that the total consumption across the neighborhood of homes does not exceed the target peak demand. Therefore, the utility can reduce the peakiness of its demand load (without simply rationing each home's usage), no matter what pricing scheme is used.

Finally, we introduce a new EMC optimization model for in-home optimization. This model accounts for the fact that, due to the proposed distributed scheduling scheme, a home may not receive as much power as desired, and therefore the scheduling problem becomes much harder. It is therefore much more realistic than the first optimization model. On the other hand, it is more computationally intensive.

Unlike prior work, we focus on demand response at the residential level for both in-home and across-homes electricity consumption. The work on in-home energy management generally focuses on HVAC applications [4], [5], [6], [7] or on architectures to support EMC-type systems [8]. The problem of scheduling electricity consumption across multiple homes could be addressed using the approach in [3], where homes are analogous to "energy aggregators." The solution proposed therein was a central hub to coordinate communications and energy control across these aggregators. In contrast, we propose here a decentralized demand response scheme for multiple local homes. We note that, although our proposed models are presented in terms of "homes" and "neighborhoods," they would apply equally well to commercial and industrial entities, and groupings thereof.

The remainder of the paper is organized as follows: Section II describes our simple optimization-based approach to in-home power scheduling and demonstrates the peak-shifting phenomenon described above. Section III discusses neighborhood-level power scheduling. We present our dynamic programming algorithm in Section IV and conclude our discussion in Section V.

II. In-Home Scheduling

A. Optimization Model

In this section, we consider the problem of optimizing appliance timing within a single home. The planning horizon for the model consists of T discrete time periods. The home has N appliances. Each appliance n, when on, consumes an amount c_n of power (measured in kW). (When off, the appliance is assumed to consume 0 kW, though this assumption can easily be relaxed.) User requests for appliance n are random, as is the length of time the appliance remains on. If appliance n is off in period t, then the probability that it is requested in period $t+1$ is given by λ_{nt}. Similarly, if appliance n is on in period t, then the probability that it completes its operation and turns off in period $t+1$ is given by μ_{nt}. Note that these probability parameters may vary over time and across appliances. A special case occurs when $\lambda_{nt} = \lambda_n$ and $\mu_{nt} = \mu_n$ for all t, in which case the duration of appliance n's off and on times are geometrically distributed with probability parameters λ_n and μ_n, respectively.

When the user requests appliance n, the EMC may turn it on immediately, or it may choose to delay turning it on. The user specifies a maximum allowable delay of d_n time periods for appliance n, with $d_n \geq 0$. (If $d_n = 0$, then the appliance is "non-schedulable" and must be turned on immediately when the user requests it.) Each period of delay incurs a cost of $\psi_n^1 \geq 0$, which represents the inconvenience to the user brought about by the delay and is measured in \$/hr. (The superscript 1 is meant to distinguish this parameter from a related parameter ψ_n^2 that will be introduced in Section III-B.)

The cost of electricity in period t is denoted by π_t and measured in \$/kWh. We assume for simplicity that electricity prices are deterministic (but dynamic) throughout the planning horizon, but if the prices are in fact stochastic, then π_t can simply be replaced by its mean.

If appliance n is requested in period t, the decision of when to turn it on is simple: we must simply find the s that solves

$$\min_{t \leq s \leq t+d_n} (s-t)\psi_n^1 + \sum_{r=s}^{T} \left(\prod_{i=s}^{r-1} (1-\mu_{ni}) \right) \pi_r c_n. \quad (1)$$

The first term represents the delay cost (in dollars) incurred by waiting until period s to turn the appliance on. The second term (also measured in dollars) represents the expected energy cost while the appliance is on. The product term within the second term calculates the probability that the appliance is still on in period r. (We take the product $\prod_{i=s}^{r-1}$ to equal 1 if $r = s$.) This minimization problem can be solved in $O(T^2)$ time since the product over i can be updated within the loop that calculates the sum over r, using one operation per iteration. Therefore, the problem of optimizing all N appliances can be solved in $O(NT^2)$ time.

This algorithm works since each appliance can be optimized individually. This approach would not work if, for example, electricity prices were a non-linear function of the load or if (as in Section III) there are constraints on the power used by a given home.

B. Simulation

We simulated a neighborhood consisting of 50 homes, each containing 3 appliances that are scheduled by an EMC. The appliance parameters are summarized in Table I.[1] The request probability λ_{nt} was assumed to vary throughout a 24-hour period, with a minimum of "Min λ_{nt}" and a maximum, attained at 6:00 PM, of "Max λ_{nt}." On the other hand, the μ_{nt} values were assumed to be stationary. So, for example, the dishwasher has a peak usage level of $\lambda = 0.0704$, corresponding to a mean inter-request time of 13.7 hours; its lowest usage level is $\lambda = 0.010$, corresponding to a mean inter-request time of 100 hours; and the dishwasher operates for a mean of 3 hours. Power usage values c_n are adapted from [9].

Electricity prices are assumed to take two levels, corresponding to peak and off-peak hours. During the peak period, from 10:00 AM to 10:00 PM, electricity costs \$0.21/kWh, and at all other times it costs \$0.014/kWh. (These are actual rates from Con Edison's time-of-use pricing model in New York City [10].)

The system was simulated for 5 days (120 hours), with the first and last day omitted from the results as warm-up and warm-down intervals. Although Table I reports parameters in terms of hours, the simulation and scheduling algorithm

[1]For now, ignore the column labeled "ψ_n^2." This will be used in Section III-C.

TABLE I
APPLIANCE PARAMETERS FOR 50-HOME SIMULATION. ALL PARAMETERS ARE REPORTED IN TERMS OF HOUR-LONG TIME PERIODS.

Name	c_n	d_n	ψ_n^1	ψ_n^2	Min λ_{nt}	Max λ_{nt}	μ_{nt}
Dishwasher	1.8	6	0.10	2.5	0.0100	0.0704	0.283
Clothes Dryer	3.4	4	0.25	2.5	0.0392	0.1193	0.632
Water Heater	5.0	2	0.40	5.0	0.0952	0.2078	0.865

used 10-minute time periods, and all parameters were adjusted accordingly. We simulated the system twice, once assuming that an EMC schedules the appliances using the model discussed in Section II-A and once assuming that no scheduling is performed and appliances are turned on as soon as they are requested. The two simulations used the same sample path of appliance request times and "on" durations.

The average cost per home in the non-scheduled system over days 2–4 is \$7.29, compared to an average cost in the scheduled system of \$6.24 for energy costs and \$0.68 in delay costs. Therefore, the EMC scheduler saves the average home 14.5% in energy costs, or 5.1% when delay costs are also factored in. The EMC scheduling model is therefore effective in reducing consumers' costs.

Unfortunately, it also defeats the purpose of the utility's off-peak pricing scheme. Fig. 1 plots the electricity consumption, in kW, across all 50 homes, for days 2–4 of the simulation. The shaded bands represent peak pricing periods. Although the scheduled system reduces the peak 6:00 PM demand to a small extent, it also creates a new, even larger "rebound" peak, immediately after the off-peak prices begin. The maximum load is larger in the scheduled system than in the non-scheduled system (83.2 vs. 68.0 kW), as is the standard deviation of the load across periods (13.2 vs. 11.3 kW). We note that this rebound peak occurs even though the request intensity during off-peak hours is low. Therefore, this simulation demonstrates that, at least under certain assumptions about consumer usage patterns and electricity prices, the off-peak pricing model fails to achieve its goal of reducing load peakiness, and may even worsen the problem. In the next section, we propose a mechanism that the utility can use to ensure a more level load throughout the day.

III. NEIGHBORHOOD-LEVEL SCHEDULING

We next consider a decentralized approach to support neighborhood-level load scheduling. To do so, we assume that there exists a communication network between the EMCs in each neighborhood home. To simplify discussion here, we assume that all EMCs in the neighborhood are "one hop" away and can transmit/receive each other's signal sent over a common control channel. This channel may be supported by an underlying smart metering infrastructure or some other local communication network. Access to the control channel is granted to the EMCs by following the protocols of the supporting network. We assume here that the control channel has high capacity and supports the exchange of EMC scheduling packets with high priority and high reliability.

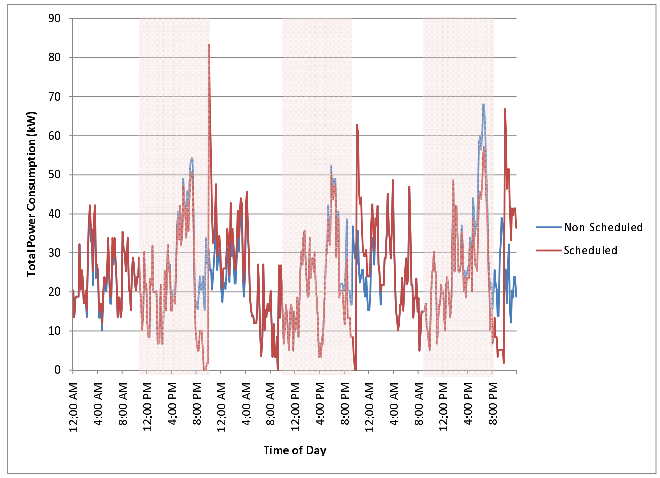

Fig. 1. Total power usage by period. Simulation includes 50 homes, 3 appliances each. Chart includes 3 days (432 10-minute periods) of results and excludes 1 day each of simulation warm-up and warm-down. Shaded bands represent peak pricing periods.

We describe below a scheduling-level exchange of packets over this channel that enables neighborhood EMCs to compete for power while collectively maintaining a relatively low and level peak demand. In contrast to purely pricing-based resource allocation, our proposed scheme does not require local EMCs to compete in a retail-level capacity market (where power is auctioned and bought). Instead, it extends access control methods typically used in communication networks to randomize an EMC's access to the local power capacity for a neighborhood. Specifically, we assume that time is divided into *scheduling slots* and a peak total power demand for the neighborhood is set (by the utility) as $P_{\max,t}$ for slot t.[2] The scheme described here aims to meet $P_{\max,t}$, which the utility can fix to be below the typical peak demand for time t.

To describe our scheme, we will first assume (in Section III-A) that homes must compete for *all* of the available power. Presumably, such a scheme would be unacceptable to consumers, who would want a guarantee that at least some of their power needs can be met at all times. Therefore, in Section III-B, we modify the scheme to guarantee each home a certain level of power at all times; homes may then compete for additional power.

A. No Guaranteed Minimum

We assume that whenever a new load request is generated within a home, the EMC will seek to meet the load requirement by coordinating with other neighborhood homes over the common control channel. During scheduling slot t, all loads that are currently being supported are termed as *active*. All EMCs that have an active load will continuously monitor the common control channel. When a new load is requested within a home, its EMC will first require information about current active loads for the neighborhood. To do so, it will transmit a

[2] The duration of a scheduling slot is roughly a few minutes and there may be several scheduling slots within a "time period" described in Section II.

probe packet over the control channel in the next scheduling slot. To support transmissions from multiple EMCs in a given scheduling slot, we assume the slot is itself divided into M "minislots." We assume each EMC selects a random minislot (uniformly distributed in the range $[1, M]$) to send out a probe message. With sufficiently large M and slot durations, an EMC's probe transmission can be transmitted with negligible probability of collision with another EMC's transmission.[3]

Once the first EMC's (say EMC i's) probe transmission goes through successfully, all other contending EMCs cancel their probe transmission and wait another random number of minislots. This time the number of minislots chosen for the backoff is uniformly distributed in the range from m to M, where m minislots is a duration long enough to obtain responses from the EMCs currently supporting active loads and one additional transmission from EMC i. This is because the successful probe message will require all EMCs supporting active loads to respond with a short *response packet* containing the power levels of their supported loads. We assume that in a neighborhood of K homes, the m minislots are long enough to support K such packets. EMCs from the K homes are assigned a transmission order during initial network formation and EMCs with active loads respond to the probe packet in this order over the m minislots.

Based on the response from the other EMCs, the EMC requesting the new load then computes the current total power usage and determines if its desired power demand can be supported within the total allowable neighborhood load of $P_{\max,t}$. If so, the load joins the active set and sends a short *admission packet* containing its power consumption over the control channel; otherwise, the EMC enters random backoff *at the scheduling layer* (for this load request) and re-attempts the inquiry procedure after its backoff timer expires. The scheduling layer backoff mechanism recognizes that this EMC's load cannot be supported in this current scheduling slot and waits for a random number of scheduling slots before reattempting a probe message and admission to the active set. We assume that at the scheduling layer, an EMC entering backoff for a given load will select a random number of scheduling slots uniformly distributed in the range $[1, w]$ for some integer w.

Other EMCs attempting a new load admission in the current time slot (that had entered a second backoff after the initial probe message) continue to monitor the control channel over these m minislots. Based on the probe response packets and the presence/absence of an admission packet, these EMCs know the current total power usage for the neighborhood at the end of the m minislots. If their requested power can be supported, they continue to monitor the channel until their backoff timer (for the minislots) expires and then simply join the active set (if the total usage permits their admission within the constraint of $P_{\max,t}$) and send out an admission packet indicating the power level of this new active load.

[3]Even if there are multiple EMCs transmitting in the same minislot, we assume the contention resolution mechanism of the communication channel will permit only one EMC to successfully transmit. The other EMCs will then backoff.

Fig. 2. Distributed scheduling algorithm flowchart.

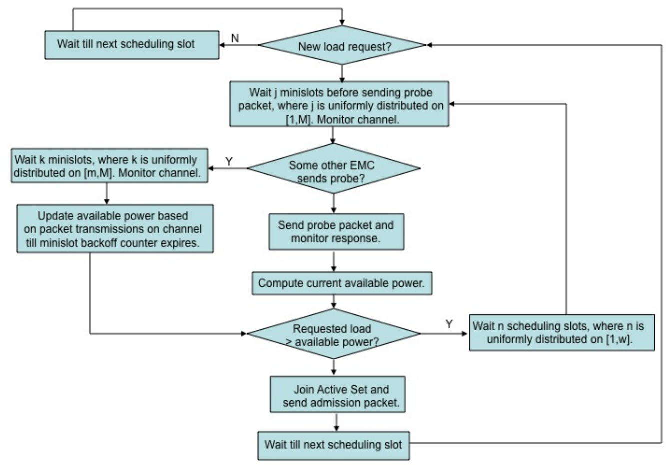

The steps of this distributed scheduling algorithm are presented in Fig. 2. The figure shows the procedure that must be followed by each EMC in the neighborhood. At the same time, those EMCs that have active loads must monitor the channel to respond to any probe messages sent out in each scheduling slot. As noted above, the transmission times for these response packets are prescribed when the neighborhood EMCs are initially connected on the control channel.

B. Guaranteed Minimum

The mechanism described in the previous section would be unpalatable to consumers since it may result in them receiving little or no power at certain times. We now modify the mechanism to guarantee each home a minimum level of power at all times, while also allowing a home to compete for additional power. Specifically, we assume now that each neighborhood home is allotted a base level of electrical power. For ease of exposition, we assume that each home is allotted the same base power level P_b, although our mechanism can easily be adapted if this level differs from home to home. The process for determining the base power level is beyond the scope of this paper, but we note that it may be set in a number of ways. For example, it may simply be imposed by the utility, or the utility may offer consumers a menu of options, with larger P_b values incurring higher prices.

At any given time, a home may require more or less than its base power level. If the requested load for a home is less than P_b, then its EMC does not need to engage in distributed scheduling with the other neighborhood EMCs since its load is "low." However, if the requested load of a home exceeds P_b (i.e., its load is "high"), then its EMC must coordinate and compete with other homes so that they collectively meet the $P_{\max,t}$ requirement. In our proposed mechanism, in each scheduling slot, an EMC for a high-load home will first use as much as possible of its base power level P_b and then compete with other EMCs to obtain additional power to meet the rest of the load. If it is unable to obtain enough power to meet the

Fig. 3. Total power usage by period under distributed scheduling mechanism with $P_{\max,t} = 40.0$ and $P_b = 4.0$.

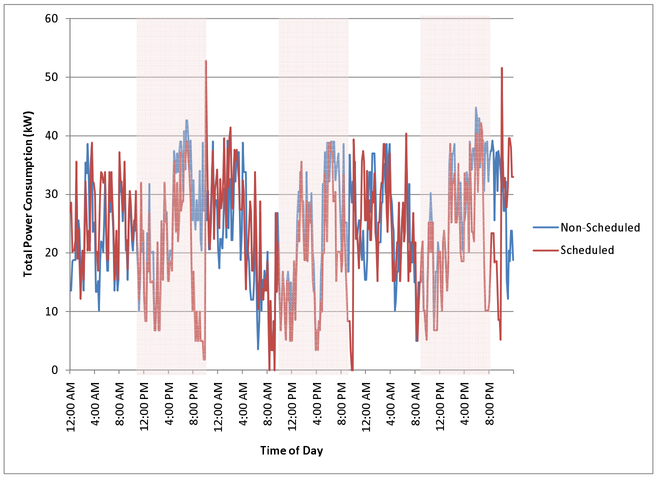

entire demand, the EMC will prioritize appliances based on their delay penalties ψ_n^1.

Under this mechanism, it is impossible to *guarantee* that all appliances are turned on within their maximum delay d_n, since there is no guarantee that sufficient power is available at all times. Therefore, we introduce a second delay penalty ψ_n^2 that is incurred in each period of delay after d_n has elapsed. That is, if appliance n is requested in period t and turned on in period $t + \Delta$, $\Delta > d_n$, then it incurs a delay cost of $\psi_n^1 d_n + \psi_n^2(\Delta - d_n)$. The mechanism could also be constructed so that additional electricity may be purchased on demand, at a premium; if this premium is smaller than the additional delay cost, then the transaction is worthwhile.

C. Simulation

Using the same data as in the simulation in Section II-B, we simulated the 50-home neighborhood, now operating under the distributed scheduling mechanism described in Sections III-A and III-B. The ψ_n^2 values (not used in the previous simulation) are given in Table I. In our simulation, we assume that the scheduling slots and the optimization time periods are equal (10 minutes) and coincide with each other. For the sake of simplicity, we also assume that $w = 1$; that is, EMCs do not enter a random backoff when their desired load cannot be met, but rather, they simply try again in the next scheduling slot. In addition, we do not simulate the transmission of probe packets, since the communication channel is resolved at a time scale that is much shorter than that of the scheduling slots/optimization periods.

Fig. 3 plots energy usage over time assuming $P_{\max,t} = 40.0$ for all t and $P_b = 4.0$. Note that the peakiness is greatly reduced and the load is significantly more level than in Fig. 1.[4]

We also simulated the system under a range of $P_{\max,t}$ and P_b values. Fig. 4 plots the average total cost per home,

[4]In some periods, the total load exceeds $P_{\max,t}$. This is because $50 \cdot P_b > P_{\max,t}$, so when many homes use close to their base power level, the total load may exceed $P_{\max,t}$.

Fig. 4. Average total cost per home for scheduled system for various values of $P_{\max,t}$ and P_b.

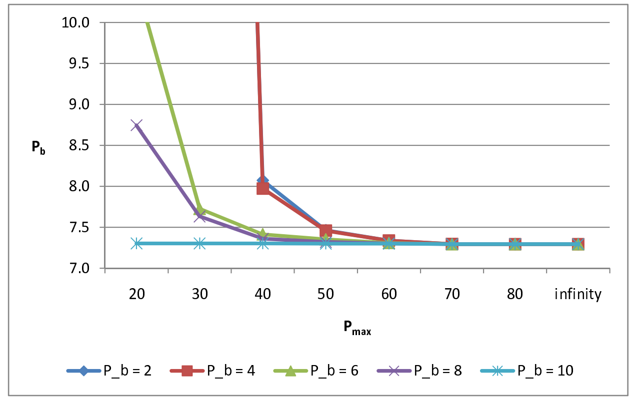

for the scheduled system, over days 2–4 of the simulation, for $P_{\max,t} \in \{20, 30, \ldots, 80, \infty\}$ and $P_b \in \{2, 4, \ldots, 10\}$. ($P_{\max,t} = \infty$ represents the case in which the distributed scheduling mechanism is not used and all homes may use as much power as desired.) Note that the consumer cost increases only slightly or not at all as $P_{\max,t}$ decreases from ∞. Therefore, significant reductions in peak demand are possible with minimal inconvenience/cost to the consumer. (Of course, in some cases, the peak demand exceeds $P_{\max,t}$, as demonstrated in Fig. 3.)

As $P_{\max,t}$ continues to increase, the cost begins to increase sharply as some appliances become delayed beyond d_n and therefore begin to incur the much higher delay cost of ψ_n^2. This occurs because the optimization model employed by the EMC implicitly assumes that the home can use as much power as desired. If the EMC delays operation of an appliance, only to find that insufficient power is available when the appliance should be turned on, significant increases in cost may occur. To avoid this, the EMC optimization model should account for potential constraints on the available power. The algorithm described in the next section is designed to do this.

IV. Dynamic Programming Algorithm

We now introduce a dynamic programming (DP) algorithm to optimize appliance operation times, subject to a capacity constraint on the total power available. This model assumes a single, fixed capacity level and ensures that the total scheduled power consumption never exceeds the capacity. Of course, under the distributed scheduling scheme from Section III, the capacity is stochastic and changes over time. Our DP could be adapted to handle the stochastic capacity, but this would require knowing the probability distribution of the capacity in each period, and these distributions depend on the EMCs' decisions, which in turn depend on the distributions, and so on. The resulting equilibrium decisions (if any) would be optimal for the in-home scheduling problem, but determining them is significantly more difficult and is outside the scope of this paper. Of course, even though the EMC optimizes based on a fixed capacity level, it could adjust its decisions in real time as the stochastic capacity is realized.

We assume that the maximum available capacity in period t is given by b_t. For example, we might set $b_t = \alpha P_b$, where $\alpha \geq 1$. Our DP ensures that the load used by the home in each period never exceeds b_t. It also ensures that each appliance is turned on before its maximum delay d_n has elapsed; therefore, the parameter ψ_n^2 is not relevant in this section.

Let S_{tn} be the state of appliance n in period t. The possible values of S_{tn} are $\{-1, 0, 1, \ldots, d_n\}$, where state -1 represents the appliance being on, state 0 represents it being off and not requested, and state i, $1 \leq i \leq d_n$, represents it having been requested for i periods. Let $\mathbf{S}_t = (S_{tn})_{n=1}^N$ be the state vector for period t. Note that this state space may be quite large. For example, if there are N appliances and each has $d_n = d$, then there are $(d+2)^N$ possible state vectors $\mathbf{S}_t$. Therefore, this DP approach is practical only for small-sized problems.

Let $f_t(\mathbf{S}_t)$ be the optimal cost for periods $t, t+1, \ldots, T$ if the system begins period t in state $\mathbf{S}_t$. In a given period, let $\mathcal{O}$ be the set of appliances that are currently on, let $\mathcal{R}$ be the set of appliances that have been requested but not yet turned on, and let $\bar{\mathcal{R}} \subseteq \mathcal{R}$ be the set of requested appliances that have reached their maximum delay d_n. We must choose which set $\mathcal{T} \subseteq \mathcal{R}$ of requested appliances to turn on. Let $f_{T+1}(\mathbf{S}_t) \equiv 0$ for all $\mathbf{S}_t$. Then $f_t(\mathbf{S}_t)$ can be expressed recursively as

$$f_t(\mathbf{S}_t) = \min_{\mathcal{T} \subseteq \mathcal{R}} \left\{ \sum_{n \in \mathcal{O} \cup \mathcal{T}} \pi_t c_n + \sum_{n \in \mathcal{R} \setminus \mathcal{T}} \psi_n^1 + \mathbb{E}[f_{t+1}(\mathbf{S}_{t+1})] \,\middle|\, \bar{\mathcal{R}} \subseteq \mathcal{T}, \sum_{n \in \mathcal{O} \cup \mathcal{T}} c_n \leq b_t \right\}. \quad (2)$$

The minimization is taken over all possible subsets $\mathcal{T}$ of $\mathcal{R}$. The first two terms inside the braces calculate the current-period cost—the energy cost for appliances that already have been or are about to be turned on, and the delay cost for the appliances that have been requested but not yet turned on. The third term calculates the expected cost in all future periods; the expectation is taken over all possible states $\mathbf{S}_{t+1}$, after accounting both for the decisions made in the current period (which appliances to turn on) and the random state transitions (appliances turning off or being requested) that occur at the beginning of the next period. The constraints on the last line ensure that all appliances that have reached their maximum delay are turned on and that the total load, among all appliances that are already on or are selected to be turned on, does not exceed the capacity.

This DP presents three main computational challenges, all due to the large state space: (1) the number of states $\mathbf{S}_t$ for which we must compute $f_t(\mathbf{S}_t)$ is large; (2) the number of subsets $\mathcal{T} \subseteq \mathcal{R}$ in the minimization may be large, depending on $|\mathcal{R}|$; and (3) the expectation $\mathbb{E}_{\mathbf{S}_{t+1}}[\cdot]$ is difficult to calculate due to the large number of possible states to transition to. For small values of N and d_n, the DP may be solved exactly. However, for larger values, approximate dynamic programming (ADP) [11] approaches may be employed. For example, in our implementation we use sampling to calculate $\mathbb{E}_{\mathbf{S}_{t+1}}[\cdot]$ and to determine which states $\mathbf{S}_t$ we must compute $f_t(\cdot)$ for. We also reduce the possible subsets $\mathcal{T}$ under consideration in the minimization by considering only $\mathcal{T} = \bar{\mathcal{R}}$, $\mathcal{T} = \bar{\mathcal{R}} \cup \{n\}$ for each $n \in \mathcal{R}$, and $\mathcal{T} = \mathcal{R}$. Techniques such as these greatly improve the algorithm's execution speed, though of course the resulting algorithm no longer guarantees the optimal solution.

If the user does not specify a maximum allowable demand d_n, but only specifies the delay penalty ψ_n^1, then the DP can be simplified considerably. In this case, we can reduce the state space by collapsing the states $\{1, 2, \ldots\}$ into a single state and ignoring the constraint $\bar{\mathcal{R}} \subseteq \mathcal{T}$.

V. Conclusion

In this paper, we propose a power scheduling protocol in a SmartGrid system, as well as two optimization methods for choosing the timing of appliance operation within a home in order to take advantage of lower off-peak energy prices. Our distributed scheduling mechanism guarantees homes a base power level while allowing them to compete for the remaining available power. Simulation results demonstrate that off-peak pricing models may exacerbate, rather than alleviate, the problem of demand peakiness, and that our distributed scheduling protocol can overcome this problem.

Acknowledgment

The authors gratefully acknowledge partial support from a Faculty Innovation Grant from Lehigh University.

References

[1] K. Wacks. GridWise architecture council / NIST home-to-grid domain expert working group, April 2009. http://collaborate.nist.gov/twiki-sggrid/pub/SmartGrid/H2G/H2G-Req5.pdf.

[2] Pennsylvania Act 129, 2008. http://www.puc.state.pa.us/electric/Act_129_info.aspx.

[3] M. LeMay, R. Nelli, G. Gross, and C.A. Gunter. An integrated architecture for demand response communications and control. *Proc. of Hawaii International Conference on System Sciences*, pages 174–183, 2008.

[4] Duy Long Ha, Stephane Ploix, Eric Zamai, and Mireille Jacomino. A home automation system to improve household energy control. In *12th IFAC Symposium on Information Control Problems in Manufacturing*, Ecole des Mines, Saint Etienne France, May 2006.

[5] Duy Long Ha, Stephane Ploix, Eric Zamai, and Mireille Jacomino. Tabu search for the optimization of household energy consumption. In *2006 IEEE International Conference on Information Reuse and Integration IEEE IRI 2006: Heuristic Systems Engineering*, Waikoloa, Hawaii: Etats-Unis d'Amerique, 2006.

[6] Shadi Abras, Stephane Ploix, Sylvie Pesty, and Mireille Jacomino. A multi-agent design for a home automation system dedicated to power management. In *Artificial Intelligence and Innovations 2007: from Theory to Applications*, pages 233–241. Springer, Boston, 2007.

[7] D. L. Ha, M. H. Le, and S. Ploix. An approach for home load energy management problem in uncertain context. *Proc. of IEEE International Conference on Industrial Engineering and Engineering Management*, pages 336–339, 2008.

[8] Y. K. Penya. Last-generation applied artificial intelligence for energy management in building automation. In *Proceedings of the 5th IFAC International Conference on Fieldbus Systems and their Applications, FeT 2003*, pages 79–83, Aveiro, Portugal, July 2003.

[9] U. S. Department of Energy. Estimating appliance and home electronic energy use, May 2010. http://www.energysavers.gov/your_home/appliances/index.cfm/mytopic=10040.

[10] K. Belson. Rewarding those who wait to flip the switch. *New York Times*, 2008. July 21.

[11] W. B. Powell. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley-Interscience, New York, 2007.