Exploring Power-Voltage Relationship for Distributed Peak Demand Flattening in Microgrids

Zhichuan Huang
University of Maryland,
Baltimore County

David Corrigan
State University of New York,
Binghamton

Ting Zhu
University of Maryland,
Baltimore County

Hongyao Luo
University of Maryland,
Baltimore County

Xiaoxiong Zhan
State University of New York,
Binghamton

Yu Gu
IBM Research-Austin

ABSTRACT

Due to limited energy storage units in microgrids, how to regulate peak demand is one of the main challenges. Thus, researchers propose different techniques to flatten peak demand in individual residential buildings. However, if each home in the grid flattens peak demand only with its own power consumption information, it is possible that peak demand of the microgrid would not be flattened but only shifted to another period. Therefore, it is critical for homes to cooperate with each other to flatten peak demand. In this paper, we utilize the power-voltage relationship in individual homes to enable that each home can infer the information of power consumption in the community by locally monitoring the voltage value on the common power line. The inferred information is then used for homes to flatten peak demand of the microgrids in a distributed manner. Furthermore, we leverage existing thermal appliances (e.g., water heaters) as thermal "batteries" in individual homes instead of purchasing batteries to flatten peak demand. We evaluate our system's performance by conducting experiments and extensive empirical data driven simulations. Evaluation results indicate that our design enables homes to effectively flatten peak demand by more than 29% without affecting homeowners' behaviors.

Categories and Subject Descriptors

J.7 [**Computer Applications**]: Computers in Other Systems—*Command and control*

General Terms

Measurement, Design, Management

1. INTRODUCTION

The microgrids enables a small number of homes to be interconnected and share electricity generation and energy storage. The microgrid is getting more popular because it can be disconnected from the power grid and provide high reliability. Given extremely limited energy storage units in microgrids, energy is normally consumed as generated. To meet short periods of peak demand, a significant portion of energy generation infrastructure in power grids is reserved with high operational costs. In traditional power grids, utility companies spend extremely expensive investment to meet with high peak demand which rarely occurs. For example, Energex (an Australian electric power distribution company) used 13% of its $8.8 billion infrastructure for only 100 hours to match peak demand within the whole year of 2010 [3]. In microgrids, with less amount of homes, it is more likely that large portion of homes may consume energy at the same time. Thus, the peak demand compared to average power consumption can be much higher in microgrids, which introduces higher operation costs for peak demand. To flatten peak demand in microgrids, researchers have proposed two different approaches: i) workload shifting [6] and ii) energy storage (e.g., batteries) [12, 17]. For workload shifting, homeowners have to reschedule their appliances' workloads to avoid peak energy consumption periods, which is inconvenient for the homeowners. For energy storage approach, large batteries are charged when the demand is low and discharged when the demand is high [12]. However, large batteries are expensive and have high energy conversion loss during the charging and discharging stages. Furthermore, in these approaches, each home flattens peak demand only with its own power consumption information, which may result in peak demand shifting but not flattening. Thus, it is critical to let homes in the grid cooperate with each other to flatten the peak demand. One possible solution is to measure peak demand with a meter in the transformer and then broadcast the information by a central controller to all the homes. However, the central controller would have to broadcast the information frequently to enable real-time control in each home, which introduces high communication overhead. The major challenge is to notify each home the aggregated power consumption of the community in a distributed manner without leaking private information of individual homes.

To address this challenge, we investigate the relationship between aggregated power consumption of the community and voltage at each home. The main observation is that voltage at each home decreases when aggregated power con-

ICCPS '15 April 14 - 16 2015, Seattle, WA, USA

http://dx.doi.org/10.1145/2735960.2735968.

sumption increases. To generalize this observation, we build the power-voltage relationship model. Note that the frequency is used in traditional power grid to infer aggregated power consumption. However, the frequency is only sensitive when there is huge power consumption change in the power grid and the frequency measurement equipments (e.g., PMUs) are expensive. Thus, we investigate the relationship between aggregated power consumption and voltage instead of frequency in microgrids. Based on this model, each home can infer aggregated power consumption of the community in real-time by monitoring its own voltage value and power consumption through the common power line. Because homes can only infer the aggregated power consumption, the privacy of individual homes' power consumption can be protected. With the aggregated power consumption, each home can either schedule its workloads or utilize battery to avoid high power consumption of the community.

To further avoid the high investment of batteries, we leverage the existing thermal appliances (e.g., water heaters) as thermal batteries in individual homes to cooperatively flatten peak demand. In our paper, we use water heaters in small residential buildings as a case study to demonstrate the feasibility. The main reasons are i) unlike the additional high cost investment required by large batteries or ice storage machines [7], water heaters are existing appliances in homes; ii) according to the technical report from the U.S. Energy Information Administration [1], 18% of total household energy is used by water heaters in the U.S.; and iii) energy can be "stored" in the water heater for a relatively long period. Thus, we can utilize the water heaters as thermal "battery" to flatten the peak demand. Note our design can also be applied to utilize other thermal appliances (e.g., HVAC) to flatten peak demand.

Unlike traditional batteries, we can only store energy in water heaters. It is very difficult to converse the thermal energy back to the electrical energy to power other appliances during the peak demand periods. Furthermore, we need to ensure our design will not affect the homeowners' hot water usage activities. To address these challenges, we deploy temperature sensors and water flow sensors to monitor the water heater's status (e.g., water temperature and water usage) in real-time and predict future hot water usage activities based on historical data. Based on the predicted hot water usage activities, each home can locally schedule the water heater's heating events based on power consumption in the community to avoid the community level high aggregated power consumption. The main contributions of this paper are as follows:

• To the best of our knowledge, this is the first distributed approach that utilizes the power-voltage relationship in each home to flatten peak demand in microgrids. By locally monitoring voltage in every individual home, we can infer the aggregated power consumption in a community and shift workloads accordingly to flatten the peak demand.

• We i) build a comprehensive model to represent the energy loss and energy consumption of the water heater; and ii) design a distributed water heater scheduling algorithm, which locally schedules the heating time of water heater without affecting the home owners' hot water usage and achieves peak demand flattening in the whole community.

• We leverage the existing water heaters as thermal batteries to flatten peak demand. Our approach does not need additional investment for batteries and avoid the potential pollution associated with batteries. Our approach can also be applied to homes with other thermal appliances (e.g. HVAC).

• We extensively evaluate our design by conducting experiments in two homes and simulations with 10 homes in a community. The results indicate our design is effective, practical, and outperforms other approaches in flattening peak demand. Specifically, our design enables homes to effectively flatten peak demand by more than 29% without affecting homeowners' hot water usage.

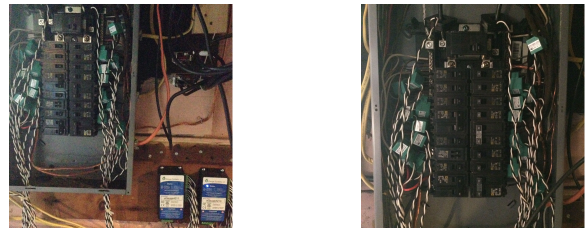

Figure 1: eGauge Deployment at two homes

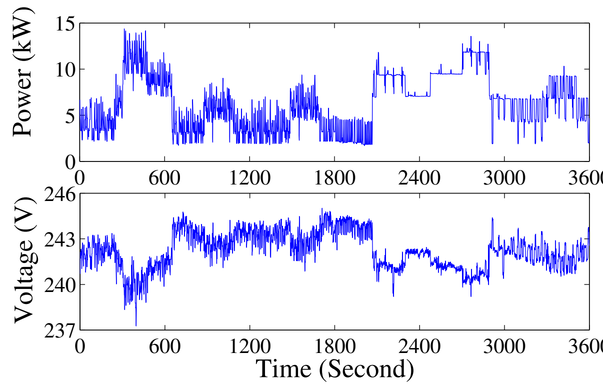

Figure 2: Power and Voltage at homes H_1

2. MOTIVATIONS

Previous work on reducing peak demand either focuses on a single home [12] or proposes a centralized solution for a whole community [11] [20]. Little research focuses on a distributed design to flatten peak demand by allowing individual homes cooperatively control their workloads. The main challenge is how to detect high power consumption periods and share power consumption information among homes in a distributed manner. To address this design challenge, we describe our empirical study on the relationship between power consumption and voltage values in a community. Specifically, we observed the following phenomena, which serves as the foundation of this work.

Observation: *Since homes are connected in parallel to the same distribution power line, if power consumption of one home increases, the voltage at all homes will decrease.*

Our observation shows that it is possible for each home to infer aggregated power consumption in a community by monitoring its local voltage value changes.

2.1 Experiment Setup

In our experiment, we deployed eGauge power meters [2] at two individual homes (H_1 and H_2) to collect the energy consumption related data (e.g., power, voltage, current, and etc.) every second. The experiment setup is shown in Figure 1. These two homes are directly connected to the same transformer. Note that there are also other homes connected to the transformer. To reduce the impact of other homes,

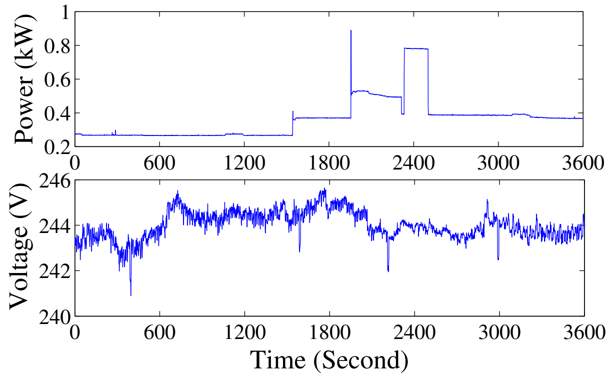

Figure 3: Power and Voltage at homes H_2

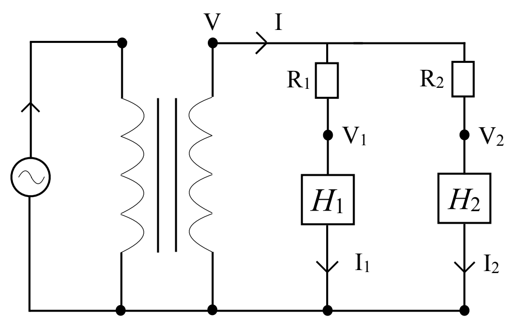

Figure 4: Equivalent Electrical Circuit Model for 2 Homes

we conducted our experiments at midnight. In the experiment, we use the empirical data from these two homes to investigate the relationship between the power and voltage.

2.2 Empirical Results

In the experiment, we intentionally maintain low power consumption at H_2 and find that voltage at homes H_1 and H_2 is closely related to power consumption at home H_1, especially when H_1's power consumption changes dramatically. We also get the similar experiment results by maintaining low power consumption at H_1 and change H_2's power consumption. Figures 2 and 3 show the changes of power consumption and voltage values within the same one hour at homes H_1 and H_2, respectively. Because power consumption at H_2 is relative low compared to H_1, voltage change at H_2 is also related to power consumption change of H_1 for the most time. During the time interval 2,000 second to 2,500 second, power consumption at H_2 varies from 0.4 to $1kW$. Then voltage change at H_2 is less related to H_1's power consumption because the change of H_2's power consumption cannot be ignored. To systematically analyze the relationship between power consumption and voltage of homes, we build the power-voltage relationship model based on electrical circuit characteristics of the power grids.

2.3 Power-Voltage Relationship Model

In this subsection, we describe our power-voltage model for two homes and evaluate the model by conducting experiments. Then we extend our model to support N homes.

2.3.1 Power-Voltage Model for Two Homes

Figure 4 shows the equivalent electrical circuit model of two homes (H_1 and H_2) connecting to a transformer in parallel. R_1 and R_2 represents the resistances of the power line from the transformer to each home. V_1, V_2, I_1, and I_2 can be measured at each home; while V is the output voltage of the transformer and I is the total current coming out of the transformer. Based on this circuit model, we have

$$V = V_1 + I_1R_1 = V_2 + I_2R_2 \tag{1}$$

According to the Electric Power Distribution Handbook [16], a transformer can be considered as a constant kVA device for a voltage from 100% to 105%. If the power consumption of one home increases, the total current I increases and voltage V drops. Since the values of R_1 and R_2 are fixed, V_1 and V_2 also drop accordingly. And if the power consumption increase is caused by home H_1, which means I_1 is also increased, then the drop of V_1 is more significant than the drop of V_2. This is reflected in Figures 2 and 3, at time equals 2,068 seconds, power consumption of H_1 is increased from $3.286kW$ to $7.384kW$. In the mean time, voltage V_1 drops from $243.44V$ to $241.716V$ and V_2 drops from $244.374V$ to $243.758V$. The voltage change at H_1 is ($1.724V$), which is larger than H_2's voltage change (i.e., $0.616V$).

2.3.2 Empirical Verification of the Model

To further verify the accuracy of our model, we conducted extensive experiments by manually controlling power consumption at H_1 or H_2. When we generated power consumption at one home, we kept power consumption at the other home stable. Note that in the local neighborhood, there are around 6 homes under the same transformer. Since we can only control two homes, there exists interference of voltage change that caused by other homes' power consumption change. To reduce this type of interference, the experiments were conducted at midnight. Then we use the relationship between current change (I_1 or I_2) and voltage differences ($V_1 - V_2$) to verify the power-voltage relationship in our model. Based on Equation 1, we have

$$V_1 - V_2 = I_2R_2 - I_1R_1 \tag{2}$$

Since R_1 and R_2 are fixed values, the voltage differences between V_1 and V_2 should be theoretically in linear relationship with the current I_1 and I_2. The results of our experiment are shown in Figure 5. Generally, the relationship between voltage differences and current is linear. Specifically, when the current is low, voltage difference is highly impacted by not only power consumption of these two homes but also other homes. Thus, the variance of voltage difference is relative high. When the current is high, current of two homes dominates the total current of the community, then variance of voltage difference is relative low. Overall, the experimental result shows that our model can accurately capture the power-voltage relationship for homes under the same transformer very well.

2.3.3 Power-Voltage Model for N Homes

In Section 2.3.1, our circuit model assumes only two homes are under the same transformer. However, in reality, there exist more than two homes under the same transformer. The circuit model for N homes is more complicated because power consumption of each home may affect voltage values at other homes differently. In general, homes are connected to a transformer as shown in Figure 6. Without loss of generality, we assume that the transformer stays in the middle of the street and homes are connected to the transformer from two directions. We find that home H_{2i-1}'s voltage value depends on i) the transformer's output voltage (V);

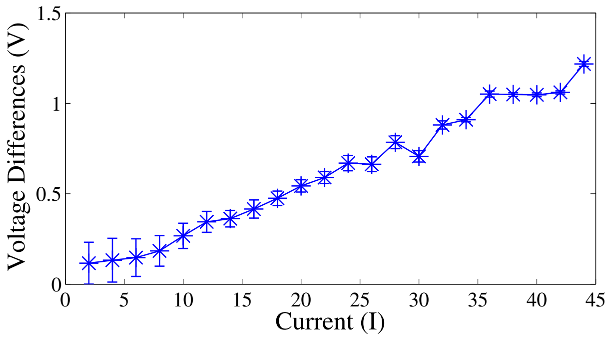

Figure 5: Voltage Differences v.s. Current

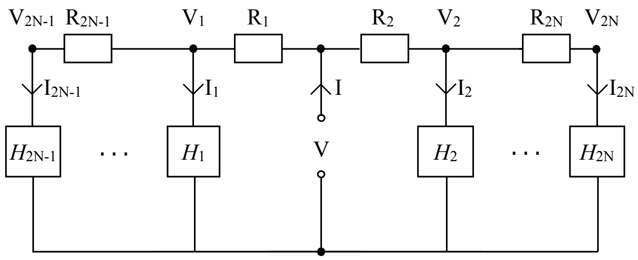

Figure 6: Circuit Model for N Homes

ii) the current from the transformer to H_{2i-1}; and iii) resistances of the power line from the transformer to H_{2i-1}. For example, H_1's voltage value only depends on the transformer's output voltage (V), the current (I_1) through H_1, and the resistance (i.e., R_1). Based on the above analysis, the voltage values at homes H_{2i-1} and H_{2i} can be calculated by using Equations (3) and (4), respectively.

$$V_{2i-1} = V - \sum_{j=1}^{i}\sum_{k=j}^{N} I_{2k-1}R_{2j-1} \quad i = 1, 2, ..., N \tag{3}$$

$$V_{2i} = V - \sum_{j=1}^{i}\sum_{k=j}^{N} I_{2k}R_{2j} \quad i = 1, 2, ..., N \tag{4}$$

Based on the Equation (3) and (4), because R_{2j-1} and R_{2j} are fixed value, the voltage drop from transformer to each home is in linear relationship of currents through the power line.

2.4 Opportunity and Challenges on Peak Demand Flattening

While power-voltage relationship can be used in many areas, this paper focuses on distributed peak demand flattening of a community such as microgrids. Existing approaches of peak demand flattening are either centralized solutions or conducted in a single building and do not consider the global impact to the whole community. To enable distributed peak demand flattening in a community, each home needs to know the aggregated power consumption in the community. Based on the power-voltage relationship, if one home monitors voltage and power consumption of its own, then it can infer the aggregated power consumption in the whole community. Thus, the power-voltage relationship provides a great opportunity for homes in a community to cooperate with each other for peak demand flattening.

However, the challenge is that each home may not have the equal information of aggregated power consumption. Based on Equation (3), V_{2i-1} is determined by I_1, I_3, $\cdots$, and I_{2N-1}, which means when power consumption of one home increases, each home may have different voltage drops. Furthermore, if more than one home increase or decrease power consumption, the voltage change at different home may be very different. Another challenge is that even every home is aware of aggregated power consumption, they need to decide how to schedule workload to avoid simultaneous low (or high) power consumption in the community. This problem is similar to the packet collision caused by the simultaneous transmissions in wireless networks. To address these design challenges, we introduce the distributed peak demand flattening (DPDF) system (detailed in Section 3).

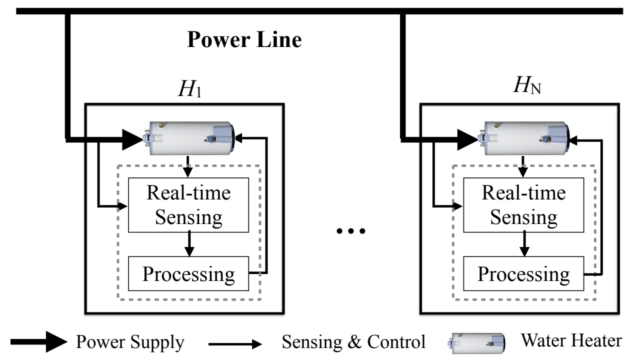

Figure 7: System Architecture

3. SYSTEM OVERVIEW

In summary, our system works as follows: Each home locally monitors its own i) power consumption, ii) voltage value, and iii) water heater related data in real-time. Based on the water heater model, the amount of heating time is calculated so that the water heater can satisfy the future hot water usage without affecting the homeowners' comfort. In the meantime, each home infers aggregated power consumption in the community based on the power-voltage relationship model. According to the community's aggregated power consumption and energy consumption demand for heating, each home can turn on the water heater when the community's aggregated power consumption is low and postpone water heater's heating tasks when the community's aggregated power consumption is high. By doing this, we can achieve distributed peak demand flattening in real-time.

Figure 7 shows the distributed peak demand flattening (DPDF) system architecture, which contains two main components at each home: real-time sensing and a processing unit. The real-time sensing component gathers real-time data from different types of sensors: temperature sensors, water flow meters, and eGauge power meters. Based on these sensor readings, the processing unit decides when to turn on the water heater and how long the water heater should be turned on. The decisions of processing unit are applied to turn on/off the water heater.

The detailed interactions of real-time sensing and processing is shown in Figure 8. In order to estimate the energy consumption of the water heater, the processing unit utilizes three types of sensing data in the left top of figure: i) water temperature data is used as an input to the standby loss model for calculating energy loss of the water heater (discussed in Section 4.1.1); ii) readings from the water flow meter are used to calculate hot water usage in real-time;

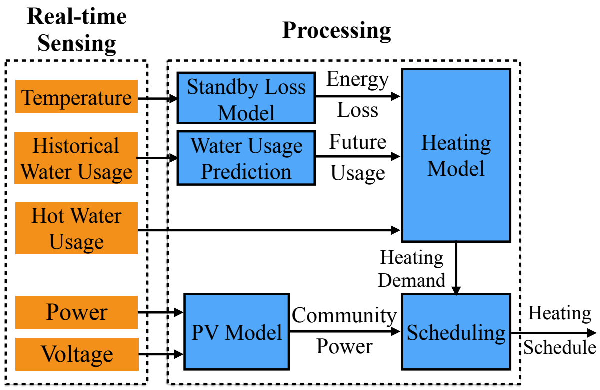

Figure 8: Interactions of Real-time Sensing and Processing

Notations	Definitions
V	Voltage at transformer
I	Current at transformer
V_i	Voltage at home H_i
I_i	Current at home H_i
R_i	Resistance from H_i to others or transformer
ϵ	Standby heat loss co-efficiency
$T_h(t)$	Temperature of hot water outlet pipe at time t
$T_c(t)$	Temperature of cold water inlet pipe at time t
$T_e(t)$	Temperature of environment at time t
$t_{standby}$	Standby time
t_{heat}	Heating time
P_{heat}	Power of water heater
E_{loss}	Energy loss for standby
E_{use}	Energy loss for hot water usage
$f(t)$	Hot water usage at time t
$S(t)$	Water tank status at time t

Table 1: Definitions of notations

and iii) historical hot water usage is used to predict future hot water usage (discussed in Section 4.1.2). Based on the standby loss, real-time hot water usage and predicted future hot water usage, the energy consumption demand for heating is calculated based on the heating model to satisfy the homeowners' future hot water usage (discussed in Section 4.1.3).

On the other hand, readings from the power line (power and voltage in the bottom of the figure) is applied to estimate the real-time aggregated power consumption in the community. Based on the power-voltage relationship model we developed in Section 2.3.3, each home can infer power consumption in the community with locally monitored power consumption and voltage data. Based on the energy consumption demand for heating and real-time power consumption in the community, each home calculates its own heating schedule (when and how long to turn on the water heater) to flatten peak demand (discussed in Section 4.2).

4. DESIGN

In this section, we present the details of water heater models (i.e., standby loss model and hot water usage prediction model) and describe how to calculate heating duration based on the models. Then combined with power-voltage relationship model, we propose a distributed design for each home to cooperatively flatten peak demand in real-time.

4.1 Water Heater Modeling

(a) Water tank (b) Flow meter (c) Temp. sensor

Figure 9: Flow meter and temperature sensors deployment

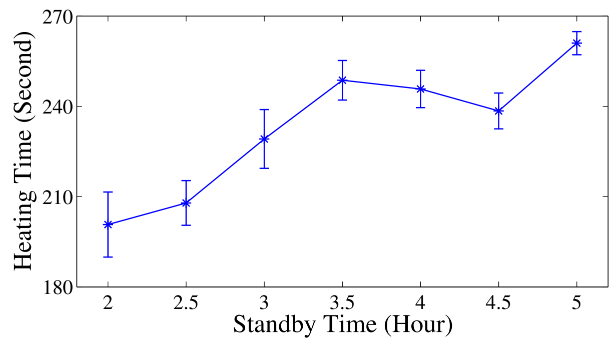

Figure 10: Standby time VS. Heating time

The water heater deployed in homes is usually running in the normal control mode. In the normal control mode, the water heater keeps the temperature of hot water within a certain range. When the temperature of hot water drops below the low threshold, the water heater starts heating until the temperature reaches the high threshold. And the water heater works with a constant pressure water tank to store the water, which means the outlet hot water is the same as the inlet cold water. The energy loss of the water heater is caused by two parts: standby loss and hot water usage. The notations used in this paper is described in Table 1.

4.1.1 *Standby Loss Model*

The standby loss is related to following factors: the temperatures of hot water and environment ($T_h(t)$ and $T_e(t)$), the standby time duration ($t_{standby}$) and the co-efficiency (ϵ).

$T_h(t)$ and $T_e(t)$ can be measured directly by deployed temperature sensors. ϵ can be found from the water heater supplier or industrial standards. But this is a theoretical parameter of the material which may have significant difference from the ground truth value. With our real-time sensing data, we use the empirical data to test the ϵ of our tank by running some experiments. Then we can calculate energy loss for standby E_{loss} as:

$$E_{loss} = \epsilon * \sum_{t_{standby}} (T_h(t) - T_e(t)) \tag{5}$$

Specifically, if there is no hot water usage, the energy loss comes from the standby loss. Hence, we have the equation $E_{loss} = E_{heat}$ and E_{heat} is the amount of energy hot water gets from water heater. Given the water heater working power P_{heat}, heating efficiency η and heating time t_{heat}, we can have $E_{loss} = P_{heat} * t_{heat} * \eta$.

To verify our model of water heater standby loss, we set up experiments for the water heater and experiment setup is shown in Figure 9. In our experiment, we turn off the water heater from 2 to 5 hours and then turn on the water heater to find how long the water heater needs to be turned on for heating. During the experiment, there is no hot water usage, so the energy loss is all from standby loss. The result is shown in Figure 10. With the standby time increases, the time for the water heater to be turned on increases linearly. However, there are some variances in the data. This is because the temperature sensor in the water tank is fixed in the middle of the water tank. And the water heater turns off when the water temperature readings from the sensors reaches the low temperature threshold. However, at that time, the temperature of water in the top of water tank may be different. Thus it may take different time for the water heater to be turned on. Overall, our model is accurate enough to predict standby loss in the water tank.

4.1.2 Hot Water Usage Prediction Model

For energy loss caused by hot water usage, the hot water usage data from the flow meter is the key factor. Given the real-time hot water usage $f(t)$ at time t, the temperatures of hot and cold water ($T_h(t)$ and $T_c(t)$) and the time duration of hot water usage, we can calculate energy loss of hot water usage from time t to $t+w$ $E_{use}(t,w)$ as:

$$E_{use}(t,w) = c * \sum_{x=t}^{x=t+w} (T_h(x) - T_c(x)) * f(x) \qquad (6)$$

The water heat capacity c, is a constant value as $4.2J/(mol \cdot K)$. Then based on working power of the water heater, we can calculate heating time duration for hot water usage.

To flatten peak demand in the community, each home needs to predict future hot water usage based on historical hot water usage. Generally, our design can apply any hot water usage prediction algorithm. For the purpose of this study, we employ a simple model based on Exponentially Weighted Moving Average (EWMA). The EWMA exploits the diurnal nature of hot water usage activities. On a typical day, it expects the hot water usage to be similar to the hot water usage of previous days with slight deviations in weather and daily activities. More sophisticated models that consider changing weather conditions, or other information can also be applied to our design. Note that thermal storage capabilities of the water in the tank allow leniency. As long as temperature of hot water is kept between the low and high thresholds, water heater function will not be affected as far as the users are concerned.

Our prediction algorithm works with a prediction window size w. Given window size w, prediction algorithm predicts hot water usage from current time t to $t+w$. Note that our prediction uses a sliding window, which means it predicts hot water usage every second instead of every w for hot water usage for the next window. The detailed prediction algorithm is as follow: i) it calculates the average hot water usage in past seven days; ii) it predicts today's hot water usage from t to $t+w$ based on historical average hot water usage and prediction errors in previous prediction. With a smaller window size, the prediction accuracy is better but it limits the time for the water heater to shift heating schedule for flattening peak demand. With a larger window size, it allows longer time for water heater to shift heating schedule but is more likely to have lower prediction accuracy.

4.1.3 Heating Duration Model

Based on the standby loss and predicted future hot water usage, we can then calculate heating duration. In our design, we need to calculate heating duration with the hot water usage after the last water heater event, future hot water usage, and standby loss. To update heating duration in real-time, we need to maintain water tank's status $S(t)$ in real-time. $S(t)$ is used to describe the current status of water tank at time t. For example, $S(t) < 0$ means the status of water tank is less than the targeted status and the water heater needs to be turned on for the hot water usage after last water heater event and standby loss; $S(t) > 0$ means water heater has extra energy for future hot water usage. $S(t)$ can be updated every second as follow:

$$S(t{+}1) = S(t) - \epsilon(T_h(t) - T_e(t)) - c(T_h(t) - T_c(t))f(t) + \eta P_{heat} \qquad (7)$$

P_{heat} is the power consumption of water heater and η is power efficiency of water heater. With real-time status of water tank, we can calculate heating demand with predicted future hot water usage. Given the prediction window size w, we calculate energy loss for standby and hot water usage in the next window, then we have heating duration as follow:

$$t_{heat} = (\epsilon(T_h(t) - T_e(t)) + c(T_h(t) - T_c(t)) \sum_{t}^{t+w} f(t)) / (P_{heat} * \eta) \qquad (8)$$

4.2 Water Heater Scheduler

With the water heater model and predicted future hot water usage, we introduce how to schedule the water heater here. The key idea is to turn on the water heater when power consumption is low and turn off the water heater when power consumption is high. However, if every home decides to turn off water heater simultaneously to flatten peak demand, the power consumption in the community will decrease significantly. Thus in turn will trigger homes to turn on water heater again. Finally, it causes an endless loop for each home to turn on/off water heater at the same time, which is not helpful to flatten peak demand in the community.

In our design, we take an adaptive feedback control to enable homes to flatten peak demand cooperatively. When each home detects it can flatten peak demand by turning on/off water heater, it does not turn on/off water heater immediately but with some backoff time. The detailed state diagram of water heater scheduler is shown in Figure 11. At every time t, each home calculates the heating demand t_{heat} based on Equation (8) and estimate power consumption in the community. When the power consumption is stable, each home keeps water heater with previous state. When it detects either high or low demand, each home calculates backoff time B. If $B = 0$, home immediately turns on the water heater when demand is low and turn off the water heater when demand is high. Otherwise, it either set up a new timer B or update an existing timer until timer expires.

4.2.1 Power Consumption in the Community

Based on Section 2.3.3, voltage at H_i depends on power consumption of all the homes under the same transformer. For power consumption change P_{heat} at H_i, if power consumption of other homes under the same transformer is the same, voltage change at H_i would be:

$$\triangle V_i = \triangle I_i * \sum_{j=1}^{j=i} R_j \approx \lambda * P_{heat} * d_i \qquad (9)$$

Because the voltage change $\triangle V_i$ is relatively small com-

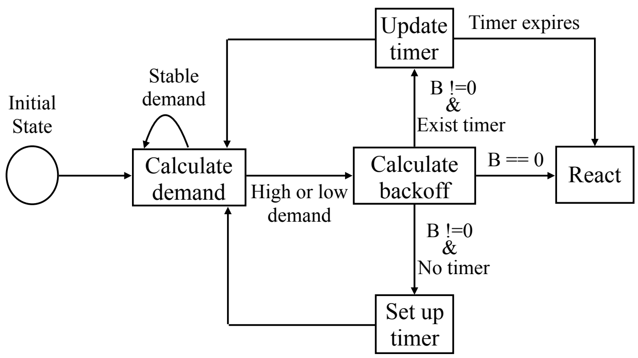

Figure 11: State diagram of water heater scheduling

pared to V_i, thus we can assume P_{heat} is linear to $\triangle I_i$. The resistances of power line depends on the distance of power line, thus $\sum_{j=1}^{i} R_j$ is also linear to d_i. Then we can have λ, which is a constant value and d_i is the distance between H_i and the transformer. With larger d_i, power transmission lines from transformer to H_i will be longer and introduce more voltage drop. Because voltage and power change is in linear relationship at each home, each home can flatten peak demand by flattening detected voltage. Thus, for every second, it compares the current voltage and average voltage $\overline{V_i}$. $\overline{V_i}$ is the average voltage of the past day in our design. If current voltage is too high or low compared to $\overline{V_i}$, it can flatten peak demand by turning on/off the water heater; otherwise, it keeps the water heater with previous status.

4.2.2 *Backoff Time Design*

In this section, we discuss how each home utilizes its own power consumption and voltage to decide backoff time B, which is generated to let the proper number of homes react to power consumption change. The backoff time depends on two parts: i) distance d_i between H_i and the transformer. If d_i is small, based on Equation (9), its power consumption change has less impact to the voltage, thus B is shorter. ii) voltage difference between current voltage value and average voltage value. If voltage difference is high, it means aggregated power consumption is either too low or too high. Thus more homes should be involved to flatten the peak demand, then B is shorter. Finally we calculate B as:

$$B = \frac{d_i}{|V_i - \overline{V_i}|/\triangle V_i} \tag{10}$$

For example, if $|V_i - \overline{V_i}|/\triangle V_i$=2, it means two homes are needed to turn on/off water heater to flatten peak demand and backoff for homes will be $0.5, 1, 1.5, \cdots$. Then two homes will react to turn on/off water heater in the next second. For the next second, each home can detect that peak demand is flattened and does not turn off the water heater anymore. Thus our design enables homes to flatten the peak demand cooperatively with real-time feedback.

4.2.3 *The Detailed Algorithm*

Combining all the design components, water heater decision algorithm can be specified by the pseudo code shown in Algorithm 1. The algorithm runs in every home and decides whether to turn on/off water heater for the next second. Each home first calculates $\triangle V_i$, B based on Equation (9) and (10) (Line 1). Then if $t_{heat} >= w$, which means to fulfill future hot water usage, the water heater needs to be

Algorithm 1 Water Heater Scheduler Algorithm

Input: t_{heat}, d_i, V_i, $\overline{V_i}$, P_i
Output: Water heater decision H.

1: Calculate $\triangle V_i$, B based on Equation (9) and (10)
2: **if** $t_{heat} >= w$ **then**
3: $\quad H = ON$;
4: **else**
5: $\quad$ **if** $V_i < \overline{V_i}$ & $H = ON$ & $0 < \triangle V_i < 2(\overline{V_i} - V_i)$ **then**
6: $\quad\quad$ **if** $B = 0$ **then**
7: $\quad\quad\quad H = OFF$;
8: $\quad\quad$ **end if**
9: $\quad\quad$ **if** $Timer == 0$ **then**
10: $\quad\quad\quad Timer = B$;
11: $\quad\quad$ **else**
12: $\quad\quad\quad Timer - -$;
13: $\quad\quad$ **end if**
14: $\quad$ **else if** $V_i > \overline{V_i}$ & $H = OFF$ & $0 < \triangle V_i < 2(V_i - \overline{V_i})$ **then**
15: $\quad\quad$ **if** $B == 0$ **then**
16: $\quad\quad\quad H = ON$;
17: $\quad\quad$ **end if**
18: $\quad\quad$ **if** $Timer == 0$ **then**
19: $\quad\quad\quad Timer = B$;
20: $\quad\quad$ **else**
21: $\quad\quad\quad Timer - -$;
22: $\quad\quad$ **end if**
23: $\quad$ **else**
24: $\quad\quad Timer = 0$;
25: $\quad$ **end if**
26: **end if**

turned on for the whole next window (Lines 2-3). Otherwise we can decide whether to turn on/off the water heater to flatten peak demand. Then we check two scenarios: i) power consumption is high and the water heater is turned on (Lines 4-13); and ii) power consumption is low and the water heater is turned off (Lines 14-23). If turning on/off the water heater can flatten the peak demand, we set up a backoff timer to turn on/off the water heater. If it does not satisfy any scenarios, we reset backoff timer (Lines 24-28).

5. SYSTEM EVALUATIONS

In this section, we evaluate the performance of our system. We collect empirical data of (i) hot water usage, (ii) power consumption from 10 homes. To verify the performance of our design, we collect both total power consumption and power consumption of water heaters in homes. The hot water usage and power consumption is measured by temperature sensors and eGauge sensors respectively for every second. The details of power consumption data collection can be found in [10]. The experiment setup is shown in Figures 1 and 9. The water flow and the residue power consumption (the total power consumption minus the power consumption of the water heater) at one home for one week are shown in Figures 12 and 13. During our experiment, the maximum value of water flow we measured is 13 Liters/Min. The maximum power consumption in Figure 13 is around $12kW$. Since this work is the first one to investigate distributed peak demand flattening in the community, the state-of-the-art research is complementary, but provides no appropriate baseline for comparison. Therefore, we compare our design with current water heater working mode as the baseline. In current water heater working mode, the water heater is turned on when the temperature of the water heater drops

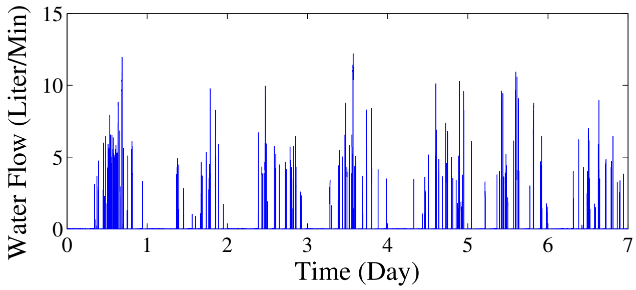

Figure 12: Water flow at one home for one week

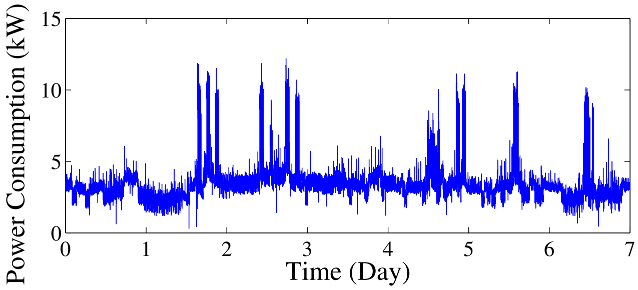

Figure 13: Residue power consumption (power consumption of water heater is removed) at one home for one week

under the low threshold and turned off when temperature reaches the high threshold.

5.1 Basic Evaluation Results

In this section, we evaluate the effectiveness of our system, which includes power consumption and standby loss of water heater in the community. All results are simulated with the seven days' empirical data of hot water usage and power consumption. Because our design needs to maintain past average voltage, in simulation of DPDF, we also use baseline in the first day. Thus we mainly compare the performance of DPDF and baseline for next six days.

5.1.1 Peak demand

To clearly illustrate the difference between baseline and DPDF, we only show the power consumption of baseline and DPDF for one day in Figure 14. For power consumption without water heater, the peak demand is mainly from 8am to 12am and 6pm to 8pm. In the mean time, hot water usage is also at the similar time. For baseline, when it detects hot water usage, it turns on water heater immediately. Furthermore, because different homes are highly possible to use hot water in similar time of evening, the peak demand rises from 53.46kW to 90.50kW (at around 7pm and 11pm). The power of the water heater in our simulation is 5.29kW, thus at least 8 homes turn on water heaters at the same time, which introduces extremely high peak demand. For DPDF, because it predicts hot water usage in future, it can turn on water heater earlier or later to flatten peak demand. The results show that DPDF's peak demand is only 64.04kW, which is 29% lower compared to baseline.

5.1.2 Water Heater Scheduling

To better understand how homes schedule water heater event, we show detailed water heater energy consumption events of 10 homes. Due to limited space, we only show water heater energy consumption events over two days in Figure 15. The upper figure is the water heater energy consumption events of the baseline. The X-axis is the time in

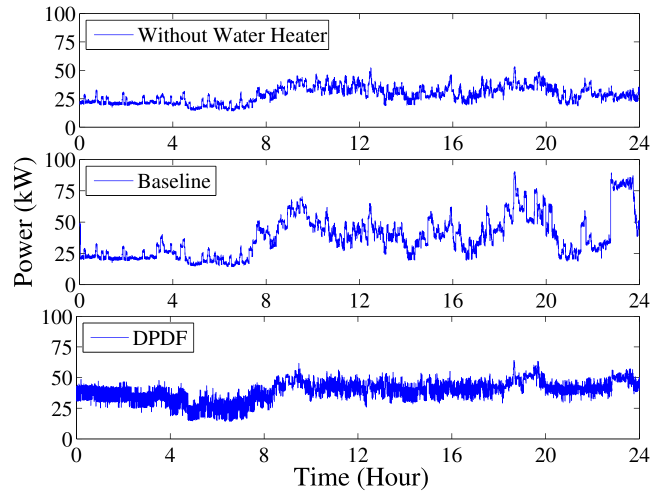

Figure 14: Power consumption in community for one day

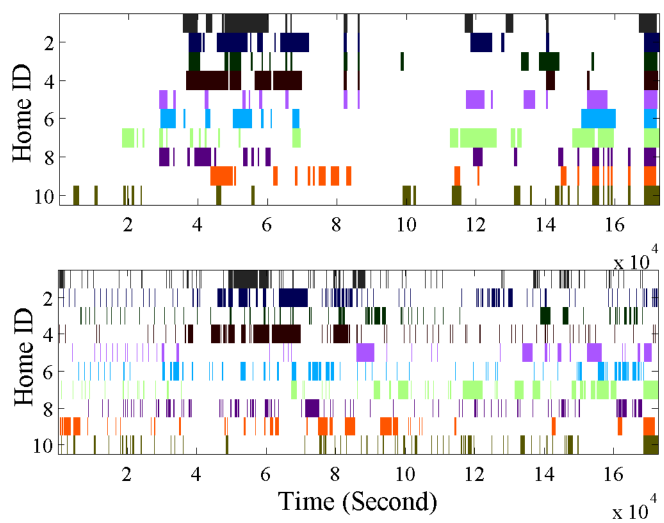

Figure 15: Water heater events over two days (Upper figure is baseline and lower figure is DPDF)

second, and water heater energy consumption events of 10 homes is shown in 10 different colors. White color means water heater is turned off. For the baseline, most of the water heater energy consumption events last for longer periods. This is because the water heaters are turned on right after hot water usage. After people take a show or bath for ten minutes, the water heater will be turned on for around 1 hour to reach the high temperature threshold. Different homes are highly possible to turn on water heaters at the same time because most people take a show or bath in the evening. For DPDF, the water heater energy consumption events last much shorter. This is because each home can detect voltage at the power line to infer power consumption in the community. If power consumption in the community is high, it can shift water heater energy consumption events to flatten peak demand. Thus the water heater energy consumption events of different homes will not collide.

5.1.3 Standby Loss

Because DPDF allows homes to turn on water heater earlier or later. The temperature of the water heater may not stay in a small range compared to baseline. And if homes turn on the water heater too early, it will introduce high

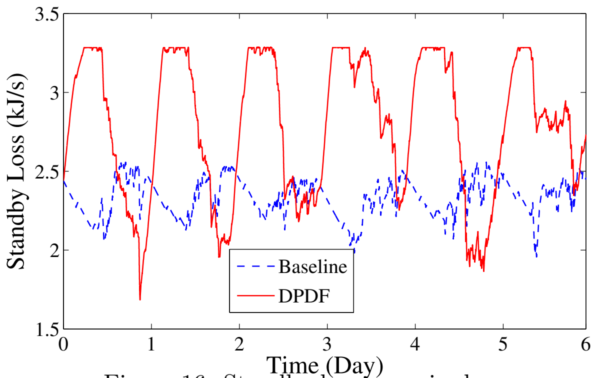

Figure 16: Standby loss over six days

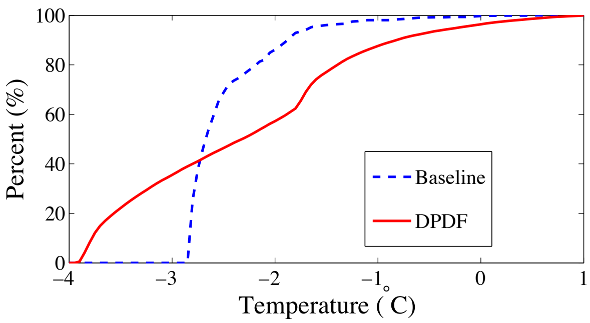

Figure 17: Distribution of temperature difference

standby loss in water heater. To investigate the impact by shifting water heater event, we show the real-time standby loss over six days in Figure 16. Based on our standby loss model, the standby loss is determined by hot water and environment temperature. We conduct our experiments for a week. The temperature sensor deployed in basement (where water heater locates) shows that the environmental temperature is almost the same through the week. Thus the standby loss is mainly determined by the hot water temperature in the water tank. For the baseline, it keeps the hot water temperature between the low threshold and high threshold, thus the standby loss is also kept in a small range. For DPDF, it allows the water heater to turn on earlier or later to flatten peak demand, thus the standby loss varies more than the baseline. However, the average standby loss of DPDF ($2.77kJ/s$) is almost the same compared baseline ($2.39kJ/s$).

5.1.4 *Hot Water Usage*

Though DPDF allows homes to turn on water heater earlier or later, DPDF can also fulfill hot water usage correctly. In Figure 17, we show the distribution of difference between the targeted temperature and the hot water temperature when there exists hot water usage events. The targeted temperature of hot water in our experiment is set as $50°C$. For the baseline, it turns on water heater immediately after hot water usage, then the temperature of hot water is always a little lower than the targeted hot water usage (mainly $2°C$ to $3°C$ lower than the targeted temperature in Figure 17). For DPDF, when a home predicts future hot water usage, it can turn on water heater earlier to better fulfill the hot water usage, thus the hot water temperature can be higher than the targeted temperature for some time. And to flatten peak demand, homes may also turn on water heater later. In Figure 17, for most of time, the hot water temperature in the water tank is only $2°C$ to $4°C$ lower than the targeted temperature. Thus the impact on people's hot water usage

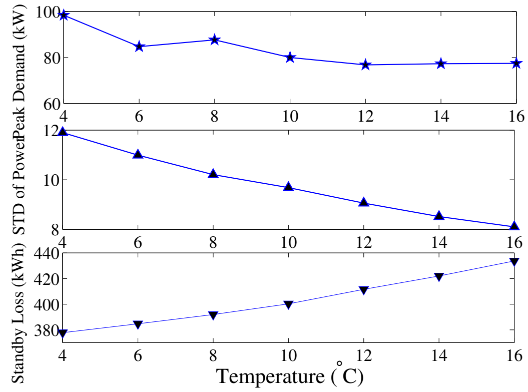

Figure 18: Impact of temperature threshold

behaviors is very low.

5.2 Advanced Evaluation Results

Because our system is designed for different environments, such as different window sizes and water tank sizes at different homes, it is crucial to investigate the system's behavior and sensitivity under diverse settings.

5.2.1 *Impact of Temperature Threshold*

In our design, homes turn on/off the water heaters based on local voltage to flatten peak demand in the community. To ensure the temperature of hot water is always acceptable for people's usage, we set a temperature threshold of hot water. If the temperature of hot water is lower than the threshold, the water heaters will be forced to be turned on to fulfill hot water usage.

The peak demand, STD (standard deviation) and standby loss of different temperature thresholds are shown in Figure 18. The X-axis is the temperature differences between temperature threshold and user's targeted hot water temperature. For peak demand, it decreases when difference between the temperature threshold and the targeted temperature increases from $4°C$ to $10°C$. This is because with a higher temperature difference, homes can flatten peak demand by turning on/off water heaters for longer periods, which reduces peak demand. However, peak demand with a temperature difference from $10°C$ to $16°C$ stays almost the same. A possible explanation is with a higher temperature difference between the threshold and the targeted temperature, homes may turn on water heaters longer to fulfill future hot water usage. However, hot water usage prediction accuracy also drops for longer periods, which increases the risk of flattening peak demand for future periods. For STD, it decreases with a larger difference because even peak demand is not flattened, power consumption is flatten by shifting the water heater events. For standby loss, because homes turn on the water heaters longer to fulfill the future hot water usage with a higher temperature difference, hot water of high temperature stays in the water tank longer, which causes a higher standby loss.

5.2.2 *Impact of Water Tank Size*

The water tank size also has direct impact on the peak demand. The peak demand, STD (standard deviation) and standby loss of different water tank sizes are shown in Figure 19. For peak demand, it decreases when the tank size increases. This is because with a larger tank size, homes can shift more water heater events earlier to flatten the peak de-

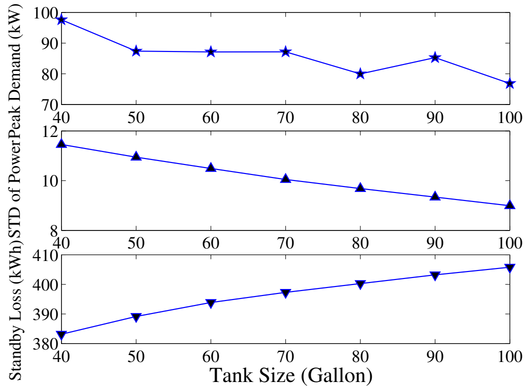

Figure 19: Impact of tank size

mand in future when current power consumption in the community is low. And homes can also turn off water heaters longer to avoid detected high power consumption and turn on water heaters later. STD also decreases when the tank size increases. Thus, increasing the tank size can flatten the peak demand much better. Standby loss increases from $384kWh$ to $406kWh$ when tank size increases from 40 gallons to 100 gallons. Considering the standby loss is for 10 homes in six days, for a single home, standby loss increases only $0.36kWh$ per day. Thus, our design can flatten the peak demand much better with a relatively low standby overhead when the tank size is larger.

6. RELATED WORK

Our work is related to the following areas of previous work.

Peak Load Flattening: There have been various works in the research community to flatten the peak demand [13]. In [8], a real-time distributed load control algorithm is proposed to reduce the variance of aggregated load by shifting the power consumption of deferrable loads to periods with high renewable generation. Distributing load at a shared electric vehicle charging lot is investigated to reduce peak consumption [9]. In [14], large energy systems (including HVAC) are scheduled to reduce peak demand. In our work, we propose to flatten peak demand of the community in a distributed way by leveraging water heaters.

Workload Scheduling: Several techniques have also been proposed for improving building energy efficiency by monitoring the occupancy [4]. In [5], authors focus on scheduling the actuation of the HVAC system by leveraging the existing WIFI infrastructure along with the mobile phones of the occupants for fine grained occupancy information. In [15] authors try to efficiently heat the houses based on occupancy prediction. In [19], incentive-driven energy sharing is proposed in microgrid for workload scheduling. to In [18], a decentralized optimal load control mechanism is proposed to provide contingency reserve in the presence of sudden demand-supply mismatch. In our work, we schedule water heater events to flatten peak demand in the community and our evaluation results show our scheduling algorithm does not affect users' hot water usage.

7. CONCLUSION

In this paper, we leverage the existing appliances (i.e., water heaters) in regular homes and introduce a distributed real-time system that allows homes in a community to cooperatively flatten the peak demand without affecting the homeowners' behavior. Specifically, we i) investigated the relationship between power consumption and voltage changes; ii) developed an accurate power-voltage relationship model; iii) designed a hot water usage prediction algorithm and a distributed water heater scheduling algorithm. To evaluate the performance of our system, we conducted real-world experiments and large scale simulations with empirical data. Evaluation results indicate that our design can effectively flatten the peak demand by more than 29% without affecting homeowners' hot water usages.

Acknowledgment

This work was supported by NSF CNS-1503590.

8. REFERENCES

[1] EIA's residential energy consumption survey. 2009.

[2] eGauge. http://www.egauge.net. 2014.

[3] Energex. https://www.energex.com.au/. 2014.

[4] Y. Agarwal, B. Balaji, S. Dutta, R. Gupta, and T. Weng. Duty-cycling buildings aggressively: The next frontier in HVAC control. In *IPSN*, 2011.

[5] B. Balaji, J. Xu, A. Nwokafor, R. Gupta, and Y. Agarwal. Sentinel: Occupancy based HVAC actuation using existing wifi infrastructure within commercial buildings. In *SenSys*, 2013.

[6] S. Barker, A. Mishra, D. Irwin, P. Shenoy, and J. Albrecht. Smartcap: Flattening peak electricity demand in smart homes. In *PerCom*, 2012.

[7] D. Du Bois. Ice Energy's "Ice Bear" Keeps Off-Peak Kilowatts in Cold Storage to Reduce HVAC's Peak Power Costs. In *Energy Priorities. p. 1. Retrieved 2007-01-31*.

[8] L. Gan, A. Wierman, U. Topcu, N. Chen, and S. H. Low. Real-time deferrable load control: Handling the uncertainties of renewable generation. In *e-Energy*, 2013.

[9] J. Huang, V. Gupta, and Y.-F. Huang. Scheduling algorithms for phev charging in shared parking lots. In *American Control Conference*, 2012.

[10] Z. Huang, H. Luo, D. Skoda, T. Zhu, and Y. Gu. E-sketch: Gathering large-scale energy consumption data based on consumption patterns. In *IEEE Big Data*, 2014.

[11] Z. Huang, T. Zhu, Y. Gu, D. Irwin, A. Mishra, and P. Shenoy. Minimizing electricity costs by sharing energy in sustainable microgrids. In *BuildSys*, 2014.

[12] A. Mishra, D. Irwin, P. Shenoy, J. Kurose, and T. Zhu. Smartcharge: cutting the electricity bill in smart homes with energy storage. In *e-Energy*, 2012.

[13] A. Mishra, D. Irwin, P. Shenoy, and T. Zhu. Scaling distributed energy storage for grid peak reduction. In *e-Energy*, 2013.

[14] T. Nghiem, M. Behl, R. Mangharam, and G. Pappas. Scalable scheduling of building control systems for peak demand reduction. In *American Control Conference*, 2012.

[15] J. Scott, A. Bernheim Brush, J. Krumm, B. Meyers, M. Hazas, S. Hodges, and N. Villar. Preheat: Controlling home heating using occupancy prediction. In *UbiComp*, 2011.

[16] T. Short. *Electric Power Distribution Handbook*. Taylor & Francis, 2003.

[17] Y. Xu, D. Irwin, and P. Shenoy. Incentivizing advanced load scheduling in smart homes. In *BuildSys*, 2013.

[18] C. Zhao, U. Topcu, and S. Low. Optimal load control via frequency measurement and neighborhood area communication. *IEEE Transactions on Power Systems*, 28(4):3576–3587, 2013.

[19] W. Zhong, Z. Huang, T. Zhu, Y. Gu, Q. Zhang, P. Yi, D. Jiang, and S. Xiao. ides: Incentive-driven distributed energy sharing in sustainable microgrids. In *IGCC*, 2014.

[20] T. Zhu, Z. Huang, A. Sharma, J. Su, D. Irwin, A. Mishra, D. Menasche, and P. Shenoy. Sharing renewable energy in smart microgrids. In *ICCPS*, 2013.