Privacy-Aware Peak Load Reduction in Smart Homes

Aarushi Sarbhai
School of Computering
University of Utah
Salt Lake City, USA
aarushis@cs.utah.edu

Jacobus Van der Merwe
School of Computering
University of Utah
Salt Lake City, USA
aarushis@cs.utah.edu

Sneha Kasera
School of Computering
University of Utah
Salt Lake City, USA
aarushis@cs.utah.edu

***Abstract*—Smart meters record power consumption data at every minute or even every second. This fine-grained data on electricity usage exposes private information about the residents of the house like the number of occupants, times of occupancy, appliance information, and much more. A solution to obscure this data is to add a battery to each home and use it strategically to manipulate the readings observed at the smart meter. Deploying such a solution at a large scale can result in sudden peaks in the energy usage. This is an alarming concern for the electric utility companies as this may cause outages, making the grid unstable. This paper is the first to expose this shortcoming and propose algorithms to mitigate the problem while maintaining the privacy of the residents. Furthermore, this paper shows that the proposed algorithms are more effective in preserving privacy than existing ones while reducing the peak load.**

I. Introduction

Extracting appliance information by monitoring the net load consumed, known as *Non-Intrusive Load Monitoring* [1], was first introduced in the early 1990s [2]. The appliance data is now easily available through smart meters that are automated to record and report power consumption in real time. This fine-grained data is used by electric utility companies to estimate power generation, distribution, and predict the demand pattern [3]. Certain third parties also have access to this data either directly from the meter or from utility companies that often do not consult the customer about how the data will be used. Third parties use this data for in-home displays, energy management solutions, or load control equipment [4].

Meter data can be disaggregated to identify specific appliance load signatures and observe schedules [1, 5] resulting in many privacy concerns. Appliance signatures can also be derived from patterns in other parameters such as the fundamental frequency, voltage and current, and harmonic currents [2]. Further, this data can be combined with data from other sources, like other installed sensors, to track an individual's position [1]. Further, when combined with publicly available demographic pointers of the residents, activities, behaviors, and preferences can be identified.

An existing approach to preserve privacy is to use a rechargeable battery [6] as a local storage entity in each home between the meter and the appliances as shown in Fig. 1. Even though many algorithms have been proposed to strategically

Supported by the Department of Energy under Grant No. DE-SC0013102.

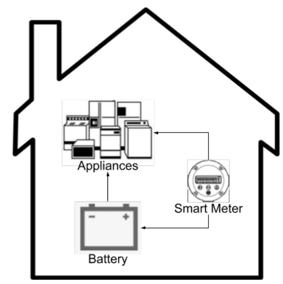

Fig. 1: Smart Home Configuration with Battery.

charge and discharge the battery to vary the readings recorded by the meter [6–8], these privacy solutions are not seen in practice. We argue that this is because the objectives of smart metering and privacy are contradictory [9] and a hypothesis that has not been considered until now is its cascading effect on power generation. Hence we analyzed the cumulative load profiles of multiple homes and observe sudden surges in the demand load. This is a concern because the utility company needs to be able to prepare for the worst case to prevent outages. If the possible peaks are rare but significantly higher than the average load observed, the generated power would be underutilized.

To prevent such scenarios, we propose to exploit Demand-Response (DR) communication schemes similar to the ones currently being used to mitigate high load at the utility companies. We leverage the infrastructure for DR systems to provide real-time feedback from the utility to individual homes. This feedback is sent when the utility detects "congestion" in the network caused by aggregation of simultaneous load consumption across the grid. When the homes receive this feedback message, the battery is discharged to supply the load in the homes. However, the battery is a local buffer of limited capacity and if all homes back-off together then once the load reduces and the feedback stops, all batteries would try to charge together resulting in a period of high load. To avoid this, we propose two algorithms to evenly distribute this load. Our novel algorithms ensure the safety of the grid while using our improved privacy solution.

We use an open-source power grid simulation tool, GridLAB-D [10], to evaluate our algorithms. GridLAB-D allows us to control different components of the smart grid and to capture time-series power distribution data. We use a dataset from the UMass Smart project [25] to run our simulations with real load patterns. To evaluate the impact of battery size on our algorithms we use three battery size models that size the battery depending on each home's load pattern. We find that larger batteries are able to hide the load signatures more effectively but they also cause high peaks in the load profile which is undesirable. Our evaluations show that our algorithms are more effective in preserving privacy with smaller batteries than existing ones and appropriately reduce the peak load by distributing it over time.

The goal of this work is to build and analyze the effect of the battery on the privacy of the residents in smart homes and the peak load at the electric utility providers. Towards this goal, we make the following contributions:

- We design an algorithm to control battery usage in each home to minimize private data disclosure.
- We develop novel algorithms that use a dynamic Demand-Response framework to reduce the aggregate peak load across a group of homes.
- We use data-driven simulations to evaluate our algorithms and show that they are effective in preserving privacy while reducing aggregate peak load.

II. Threat Model

A. Threat to Privacy

The adversary is either capable of accessing the readings directly from the power utility company, monitoring the consumption recorded at the smart meter, or by intercepting transmissions to and from the utility companies. We also consider the utility company to be adversarial in terms of violating smart home privacy, for load optimization services and targeted commercial advertisements.

B. Threat to Reliability

Fig. 2 shows the load profile that a utility provider would observe when a group of homes is using a battery-based privacy solution. We observed this pattern in our experiments described in Section VII. If the utility company is unable to predict surges in demand from homes, as the peak observed at 420 minutes, they might not be able to prepare for a sudden peak ahead of time. This makes the smart grid unreliable and prone to power outages in that area. Very high peaks can also cause fuse breaks at substations resulting in prolonged outages. We assume that the utility company would want to provide a reliable service, thus, we do not consider it to be a threat to reliability.

III. Background and Related Work

A. Smart Power Grid

The utility companies need to plan the generation and distribution pattern throughout the day [3]. The general pattern is largely predictable and an increase in demand is usually compensated by shifting the distribution, quick backup generators (which have higher start-up costs), or by buying from other companies.

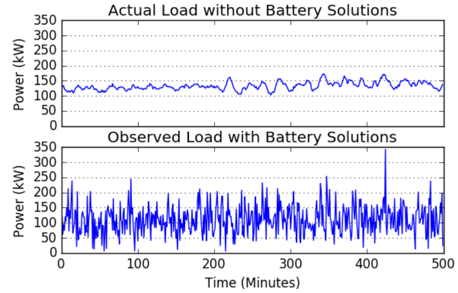

Fig. 2: Load Profile Observed at Utility Node

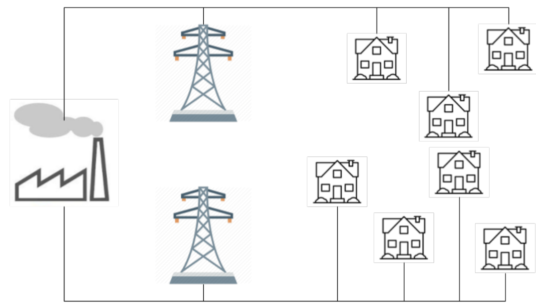

Fig. 3: Hierarchical Network Structure of the Smart Grid

Smart meters are connected to the electric utility companies over a two-tier network through substations as shown in Fig. 3. A substation (also known as collector) is connected to a group of homes [11]. They monitor the demand load over time, recording the electric usage at the homes at a fine granularity of a few minutes or seconds. Some also have the capability to log power-relevant events like outages and peaks [12].

B. Battery-based Solutions to Preserve Privacy

In this model, the demand of the devices in the home is partially met by the reserve charge in the battery and partially by drawing power directly from the grid. Therefore, the values recorded by the smart meter are a sum of the actual demand of the devices and the battery usage. The battery can be used to store and supply power to the devices in the home at strategic times to hide appliance loads from the smart meter. Early approaches to preserve privacy were to flatten the load to hide the variations in the usage patterns [6, 7]. Other ones focused on adding patterns through the use of the battery that could not be correlated to the usage patterns of specific devices [8]. The algorithms proposed in this paper follow the second class of algorithms.

These algorithms base their calculations on the actual demand in the house. Due to this intrinsic property of the algorithms and the physical restrictions of the battery (capacity and charging/discharging rate), at certain times, the meter readings can be directly correlated to the actual demand. If an adversary observes points at which the load changed, it would be possible to correlate those with activities in the house [8].

In a variation of the algorithm in [8], the battery is charged and discharged probabilistically depending on the state of charge in the battery. This makes it difficult to monitor the battery state and draw correlations to the actual demand.

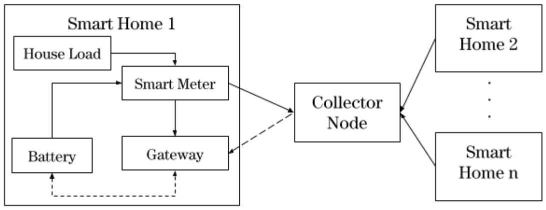

Fig. 4: Data and Control Flow Paths in Smart Grid

However, this algorithm is also restricted by the battery specifications. In this paper, we design an algorithm that removes the constraint of always discharging or charging the battery at a constant rate.

C. Peak Load Reduction Solutions

Demand Response (DR) schemes provide incentives to customers to lower electricity usage at times of high wholesale market prices or when the reliability of the system is in question [13]. Such schemes are decided ahead of time and incentives are predefined [14, 15]. They also require the customer to actively change their load according to the program they have opted for. In these cases, utility companies often use third parties to keep a check on the customers.

Direct Load Control (DLC) is a more dynamic and intrusive form of DR where the utility can remotely control the customer's electrical equipment [16–18]. This technique is common in an industrial setting where control over appliances like the air conditioning system is more effective [11, 19, 20]. However, some operators [21] offer residential DLC solutions.

Battery-based peak-aware charging algorithms are proposed in [22, 23] where factors such as cost generators and transmission rates are used to minimize the cost for residents. Their model is based on predicting the next-day demand using machine learning to reduce peak load. However, these solutions cannot be directly used in conjunction with the class of battery-based control algorithms we explore in this work.

IV. System Model

The overall system model is shown in Fig. 4. In this model, apart from the power connections to draw electricity, the substation or collector nodes are connected to a group of smart homes by data paths to communicate the information recorded by the meter and for the control messages to be sent to the homes by the utility indicating the load level. In addition to smart meters, these homes also have a battery connected through a gateway node. The function of the gateway is to implement the local control algorithm, which decides, in real-time, when and how the battery will be used to support the demand in the home. The gateway is assumed to have charge controller-like [24] capabilities to be able to control the battery at the level required.

V. Control Algorithms

In this section, we first improve on existing privacy-preserving algorithms [6–8] by operating the battery independent of any external factors, such as the load in the home. Additionally, unlike the earlier work, we charge the battery at a random rate. Next, we propose two algorithms that distribute the load over time to reduce undesired peaks that might be observed in the aggregate load at the utility.

A. Modified Privacy-Preserving Algorithm

Our *Random Charging* scheme is outlined in Algorithm 1. Charging the battery probabilistically depends on the state of charge in the battery. Unlike existing work, in our algorithm, the battery is charged at a random rate, which is bounded by the maximum charging rate. On the other hand, in our algorithm, when a discharge decision is made, the battery is discharged at the maximum rate to compensate for the demand of the devices and hide the usage pattern.

Algorithm 1 Random Charging (RC)

Input: Battery's state of charge $soc \epsilon$ [0,1]
Output: Battery rate: charging(+); discharging(-)

```
function RC(soc)
    At each timestamp (t):
    if Pr[random(0,1) > soc] then
        return (random(0,1) * rate_charging)
    else
        return (-1 * rate_discharging)
    end if
end function
```

Due to the random nature of our algorithm, we are able to construct a robust system that can largely maintain an equilibrium in the load variations at each home. However, occasional high peaks are observed in the aggregated load across multiple homes. We show a quantitative analysis of this behavior in Section VII-F2.

B. Feedback-Driven Algorithms for Peak Load Reduction

When a high load period coincides across multiple homes, a momentary spike is observed. This observed peak is unpredictable unlike the expected progressive pattern observed when no privacy solutions are used in homes. These momentary spikes can potentially cause outages in areas of high demand. To tackle these load spikes, we need a "congestion avoidance" strategy to distribute the load over time. In our model, a collector sends a feedback message to the homes when it is experiencing a surge in demand. There are no feedback messages sent when the utility is able to handle the load.

If all the homes receiving this feedback message reduce their load at the same time by using the battery, the load on the utility will drop and there will be no feedback message. At this point, all the batteries would attempt to charge with a high probability, as the state of charge would be low, resulting in an increase in the load observed at the utility. To prevent this, we study different approaches to incorporate the feedback message to add some variability in the battery usage according to the number of consecutive feedback messages received to avoid synchronization among homes.

1) Linear Response: This approach is outlined in Algorithm 2. When the feedback is first received, the probability to incorporate its effect is half. If the feedback persists, it would indicate that the load needs to reduce further. Thus, the longer the message is received, the higher is the chance of the homes to discharge the battery.

The gateway now bases its decision on two factors, the current state of change in the battery and high load indicated through the feedback messages. The probability of charging the battery is lower when the state of charge in the battery is high. The probability of charging the battery further reduces as more number of consecutive feedback messages are received.

Algorithm 2 Random Charging with Linear Response (LR)

Input: Battery's state of charge(soc) ϵ [0,1], feedback($fb(t)$) ϵ {0,1}
// normal load, fb(t) = 0
// undesired high load, fb(t) = 1
Output: Battery rate: charging(+); discharging(-)
function RC($soc, fb(t), fb(t-1)$) // Called at time(t):
    if $fb(t) == 0$ **then**
        $n = 0$
    else if $fb(t-1) == 0$ **then**
        $n = 1$
    else
        $n = n + 1$ // linear response
    end if
    if Pr[random(0,1) $>$ soc] and Pr[random(0,n) == 0] **then**
        return $random(0,1) * rate_{charging}$
    else
        return $(-1 * rate_{discharging})$
    end if
end function

2) Quick Response: In comparison to the previous Linear Response algorithm, this algorithm (Algorithm 3) reduces the probability of charging the battery by responding to the feedback quicker. Here, the probability of discharging the battery increases multiplicatively.

To prevent high load directly after the feedback messages stop, a recovery period is added that slowly relaxes the constraint induced by the feedback message. If a similar multiplicative decrease is applied to reduce this influence, then this quick decrease can cause a successive peak. This would happen when the load momentarily reduces and the feedback message stops earlier than the previous case, but since the conditions have not significantly changed, the load goes back when the restriction is removed too quickly. This creates an oscillating unstable system. By adding a slow decay, an optimally distributed equilibrium is possible.

VI. Battery Sizing

A very small battery would not be able to hide the load or act on the utility feedback effectively. On the other hand, a large battery will cost more and, as shown in Section VII-F2, increase the peaks in the load profile. Thus, determining the optimal battery size for each home is critical to find a balance. Battery specifications are decided based on analysis drawn from the load profile of the individual homes. Fig. 5 shows how much of the load the battery sizes can handle.

Algorithm 3 Random Charging with Quick Response (QR)

Input: Battery's state of charge(soc) ϵ [0,1], feedback($fb(t)$) ϵ {0,1}
// normal load, fb(t) = 0
// undesired high load, fb(t) = 1
Output: Battery rate: charging(+); discharging(-)
function RC($soc, fb(t)$) // Called at time(t):
    if $fb(t) == 0$ and $n! = 0$ **then**
        $n = n - 1$ // linear back-off
    else
        $n = n * 2$ // multiplicative response
    end if
    if Pr[random(0,1) $>$ soc] and Pr[random(0,n) == 0] **then**
        return $random(0,1) * rate_{charging}$
    else
        return $(-1 * rate_{discharging})$
    end if
end function

1) Bulky Battery Sizes: This model is designed such that the battery is sufficient to supply the load at high peak times. The capacity is chosen such that it is able to supply enough power required during the peak intervals. The discharging/charging rate is set such that any peak can be masked at a given timestamp. The sizes of these batteries are close to those used in off-grid systems.

2) Moderate Battery Sizes: The algorithms controlling the battery use are very different from the behavior assumed in the previous model, i.e., the battery is not always used to supply all the load during peak times. A more befitting model should supply some of the load during peaks. This model sizes the capacity of the battery to be able to mask instantaneous peak loads while the discharging/charging rate is adjusted accordingly.

3) Compact Battery Sizes: To test the performance bounds, this model selects battery sizes that have enough capacity to compensate for the average load observed. The discharging/charging rate is adjusted to match the respective battery.

VII. Evaluation

To evaluate the proposed algorithms, we simulate our system model in the GridLAB-D simulation tool [10]. The experiments are driven by the readings of 40 homes taken from the UMass Smart project [25] for January-June 2016. The meter readings generated by the simulation were used to compare the performance of the algorithms across all battery sizing models. The focus of the analysis is to study how effectively our algorithms are able to conceal points that reveal the activities in the homes in comparison with the existing work. To determine whether these solutions can be widely used in practice, we study the collective load distribution to identify if there are any events that may cause grid instability due to high load peaks.

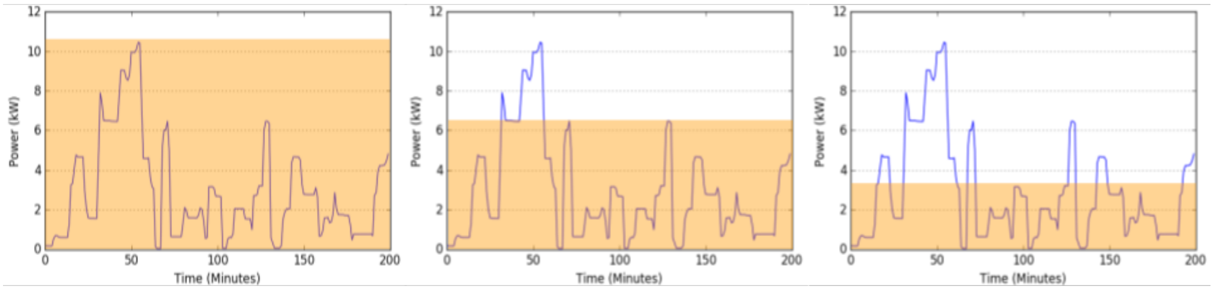

(a) Bulky Battery (b) Moderate Battery (c) Compact Battery

Fig. 5: Battery Sizing Models

A. Quantifying Privacy

An algorithm is privacy preserving if it can hide individual device signatures, load patterns, and activity bursts. Similar to existing work [6–8], we use mutual information (MI) to capture the correlation between the actual load in the home and the observed load at the meter by considering both these time-series values as distributions of random variables. MI is a measure of how much information about the actual demand load consumed by the devices is exposed by the observed load recorded at the smart meter. A lower mutual information score signifies low correlation between the actual and observed load.

1) MI with Independence Assumption: The first step in calculating mutual information is to discretize both the time-series data, i.e., the observed (x) and actual (y) load. In this first case, each value is assumed to be independent of the previous values in the same series to evaluate mutual information in the most general case. Next, the joint (p(x,y)) and marginal (p(x), p(y)) distributions are calculated using the discretized time-series data, which is finally used to calculate the MI according to (1).

$$I(X;Y) = \sum_{t} \sum_{(x,y)} p(x(t), y(t)) log \frac{p(x(t), y(t))}{p(x(t))p(y(t))} \tag{1}$$

2) MI with Correlation Assumption: The assumption that each value observed is independent of the previous is not true in the real setting as we are more likely to see bursts of activity and inactivity in homes. To evaluate the mutual information in a more realistic setting, we compare if the changes in the load profiles of the observed load are correlated to the changes in the load profile of the actual load. We see that the maximum correlation is observed at a lag of one minute by calculating the autocorrelation function for each time-series measurements. We considered a range of shifts in time, from a lag of one minute to a lag of one day.

The mutual information score is altered to account for the information lost due to the dependence within the observed load values. Therefore, we introduce a single-step Markov chain to the probabilities of the samples. This is similar to the metric used in [8]. With this assumption, we are able to capture if there is a dependence in the corresponding changes in the distributions whereas the previous case in Section VII-A1 only checks for the dependence in the corresponding values in both these distributions. It is essential to compare the privacy based on adding this assumption since changes in the load profile can also expose information. The mutual information is now calculated for these new random variables that are the observed and actual loads given the value at one timestamp before the current value. The mutual information equation is modified as per (1).

$$\begin{aligned} I(X;Y) &= \sum_{t} \sum_{(x\prime,y\prime)} p(x\prime(t), y\prime(t)) log \frac{p(x\prime(t), y\prime(t))}{p(x\prime(t))p(y\prime(t))} \\ &= \sum_{t} \sum_{(x,y)} p(f(t-1))p(f(t))|(f(t-1))). \\ &log \frac{p(f(t-1))p(f(t))|(f(t-1)))}{p(x(t)|x(t-1))p(y(t)|y(t-1))} \\ &- p(f(t-1)) log \frac{p(f(t-1))}{p(x(t-1))p(y(t-1))} \\ &where, f(t) = (x(t), y(t)) \end{aligned} \tag{2}$$

B. Peak Load Analysis

The other area of concern is the peaks observed at the utility node. The motivation to use the feedback from the utility is to avoid the scenario when we might observe these peaks by distributing the load more uniformly.

1) Load Distribution: We analyze the cumulative density function (CDF) of the load profiles as observed at the utility node to compare the load distribution. The aim is to bring the load distribution as close to the distribution of the actual load to have the least impact on load prediction.

2) Power Utilization: The peak average power ratio (PAPR) drives the power generation estimation [9]. The PAPR value is computed by the fraction of peak load by the root mean square of the load over time. A smaller PAPR value indicated higher utilization of the power generated.

$$PAPR = \frac{load\ observed_{max}}{load\ observed_{rms}} \tag{3}$$

C. Simulation Tool and Experimental Setup

The GridLAB-D time-series power distribution system simulation and analysis tool [10] provides different module abstractions such as power flow, demand response, and market pricing, which can be used to simulate the different aspects of the model. The system can be executed for different time scales, ranging from sub-seconds to many years.

The simulated model of 40 Homes connected through 4 substation or *Collector* nodes is shown in Fig. 6. Each *Collector* node is connected to 10 *Smart Home* nodes.

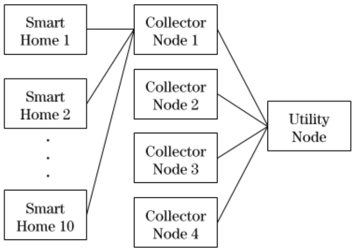

Fig. 6: Experiment Structure

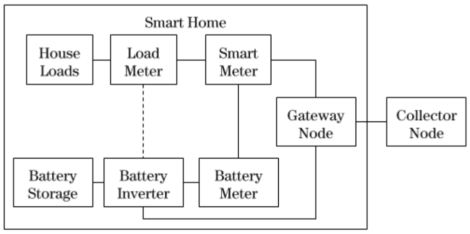

Fig. 7: Smart Home Node

We enhance the simulator by adding classes to record the different states needed by the different control algorithms; for instance, the feedback from the collector nodes, which is used in the proposed algorithms, was stored in a new parameter. The battery control algorithms were specified in a MATLAB function that was linked to the main model file.

The power generation is controlled at the utility node and the collector nodes simply convert and forward as per demand. Apart from this, the only operation outside the smart homes was recording the load at the collector nodes and communicating the feedback message for high loads back to the homes, when required.

D. Smart Home Node

Fig. 7 shows the *Smart Home* node that is realized by a collection of appliance nodes connected together. The point of connection between the home and the *Collector* node is through a *Gateway* node. The *Gateway* also receives feedback messages from the *Collector* node when it is experiencing high load. The load in the smart home is defined at the *House Load* object.

The *Smart Meter* connects the home to the substations and records the net power usage of the smart home. *Load Meter* records the power required by the devices in the home and *Battery Meter* records just the power consumed and supplied by the battery.

A *Battery* storage has been added to each home through an *Inverter* that connects to the *Smart Meter*. This configuration allows us to control the battery usage. When the battery is charging from the grid, it draws power along this path and when it is discharging, it supplied power through the same path in the reverse direction. While discharging, the demand in the house is partially or completely, supplied by the battery. To obtain such algorithmic control, the simulator has to interact with MATLAB where the local control algorithm is implemented. The *Inverter* is able to sense the demand in the home from the *Load Meter* needed for the control algorithm.

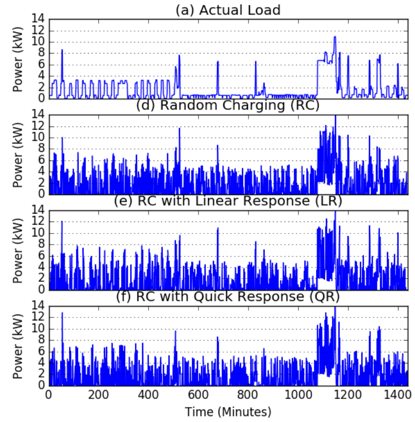

Fig. 8: Load Profile at Sample Home with Moderate Battery Sizes

E. Dataset

The UMass Smart project [25] contains data for 114 single-family apartments for the period 2014-2016. The aggregate load is recorded for every minute by a smart meter at the apartment. Of these 114 apartments, 40 were selected for our simulation based on diverse load profiles. From the dataset, we selected an even split of homes that had an average load higher and lower than the mean across all homes in the dataset. These homes also had a varying percentage of peaks to ensure widely heterogeneous usage patterns. The duration of these peaks was also a differentiating factor in selecting the homes.

F. Experimental Results

The load variation observed at the smart meter from the use of different control algorithms is directly comparable. We compare our proposed algorithms with the existing work i.e., Best Effort (BE) [6], Non-Intrusive Load Leveling (NILL) [7], and Stepping Framework (SF) [8]. Events in the actual load can be directly correlated to the observed load when limited by the battery specifications but load signatures are better hidden as the battery sizes increase. Fig. 8 shows a sample of the data recorded over a day and how the load varies with these algorithms.

1) Privacy Comparison: In this section, we show the variation in MI for the simulated homes. A bin size of 500 W was selected during the discretization step while calculating the MI. This size was empirically chosen as it was found that the MI was less sensitive to bin sizes in this range. Earlier work [8] also found this bin size to be appropriate to represent the signal well.

Fig. 9a summarizes the MI for all homes with different battery sizes under all the algorithms. It is also interesting to note how gradually information leakage is observed when the battery use is suppressed by the incorporation of the feedback message. Given strict constraints by the battery in the compact battery model, the proposed algorithms are still able to hide signatures better than the existing work. With the use of the moderate sized batteries, the proposed algorithms consistently perform better. When comparing the bulky battery model an interesting conclusion can be drawn. The BE algorithm outperforms all other algorithms, and in

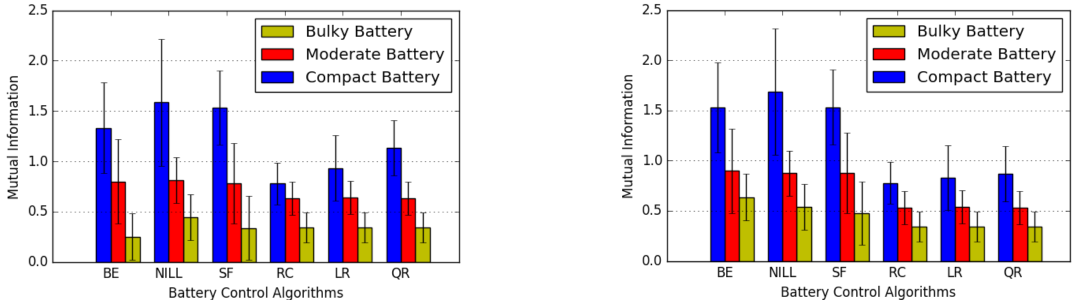

(a) Independence Assumption

(b) Correlation Assumption

Fig. 9: MI across Battery Sizes.

its best case successfully hides almost all data points. This does however come at a cost examined in the next section. The existing work also shows high variance in MI as for an interval when the battery has the capacity, it can largely hide all load changes, but this is followed by a period of maximum information exposure when the battery is unable to support the demand due to this overuse. On the other hand, the proposed algorithms, achieve a consistent level of privacy with less variation even across homes.

To summarize, when comparing the variation in MI for the various algorithms for the different battery sizes it is observed that the proposed algorithms are less restricted by the battery specifications and are able to hide load signatures better, making it harder to identify targeted load signatures. It is also interesting to note that the variations proposed in this paper are able to achieve similar MI values with compact batteries as the existing algorithms are with bulky batteries.

Fig. 9b shows that not only do the proposed algorithms leak less information than the existing work, the difference in MI between them has also increased. A drastic change is observed in the case of the bulky battery sizes from Fig. 9a where the best effort algorithm exposes more information. This is due to its high correlation in observed values, even a small change in the load profile can be easily correlated. The proposed algorithms expose much less information. The effect of the feedback message is also minimal at the privacy front and the impact of the battery specifications is minimal.

2) Load Comparison: Fig. 10a shows the load profiles as observed at a utility node that could cause instability. The proposed feedback-driven algorithms are able to realize more uniform load distribution and reduce the chance of experiencing these peak events. The change in the observed load profile is shown in Fig. 10b, with reduced peaks.

As shown in Fig. 11a, 11b, 11c, RC has a very wide variation in its pattern whereas a desirable load distribution pattern would be one that is as close to the actual load as possible with reduced peaks. Thus, the goal is to achieve a load distribution that has a shape like the distribution of the actual load and also see lower maximum values than those with the RC algorithm. By utilizing the feedback-driven algorithms, the distribution of the load improves and the chance of observing concerning peaks reduces.

With the compact battery sizes, as shown in Fig. 11a, the load distribution with the feedback-driven algorithms is very close to that of the actual load. When using the moderate sized batteries the load distribution by the feedback-driven algorithms is still better at compressing the gap between that of the actual load and the observed load with the RC algorithm, seen in Fig. 11b. The slight shift only indicated those values that were lower than that in the actual load, which is not a concerning difference. Fig. 11c shows a similar load distribution but also highlights an important observation when using bulky battery sizes. The large batteries, when used in this manner, cause very high peaks, which make them undesirable.

The PAPR values are estimated using the cumulative load, as observed by the utility node. Table I shows the impact of the feedback on the efficiency of power generation. With the feedback-driven algorithms, the load utilization is higher than the RC algorithm. The battery also has an effect on this parameter. As seen in Table I, the bulky battery sizes increase the max peak observed at the homes, resulting in higher PAPR values.

TABLE I: PAPR Values for Proposed Algorithms

Battery	Random Charging	Linear Response	Quick Response
Compact	2.710	2.277	2.213
Moderate	2.814	2.145	2.113
Bulky	7.536	7.458	7.224

VIII. Conclusion

We reveal the shortcomings in existing privacy-preserving algorithms and the threat such battery-based solutions pose to the grid's stability and reliability. To overcome the pitfalls in the existing privacy solutions we propose a Random Charging algorithm. We also propose two algorithms that maintain stability in the power grid by using real-time feedback from the utility to distribute high load over time. Through our evaluations, we show that our algorithms are more effective in preserving privacy with smaller batteries than existing ones and appropriately reduce the peak load by distributing it over time. A mathematical analysis of the proposed algorithms will be explored in future work. Our work can also proceed along several directions. First, we will investigate the problems of privacy and peak load reduction in microgrid systems that have local power generation capabilities. Second, we will investigate the interplay between the battery life and the random charging patterns. Last, we will explore machine learning based techniques to derive the battery sizes suitable for each home.

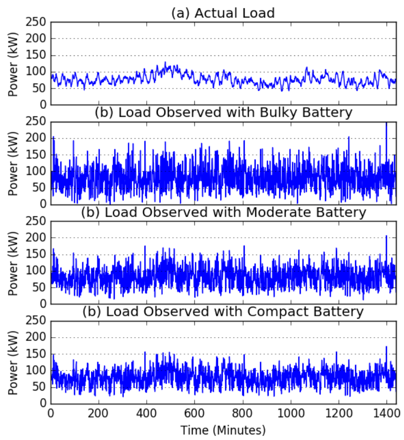

(a) Random Charging across Battery Sizes

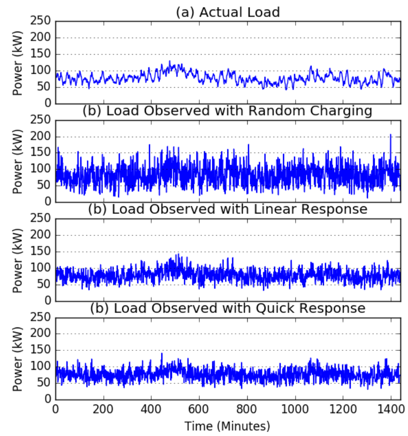

(b) Proposed Solutions using Moderate Battery

Fig. 10: Load Profile at Utility Node

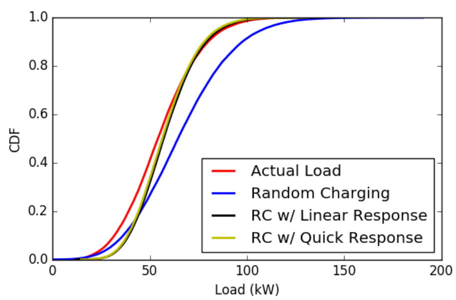

(a) Compact Battery Sizes

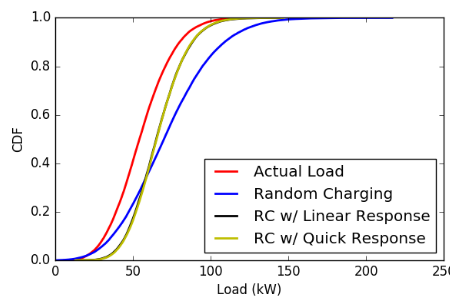

(b) Moderate Battery Sizes

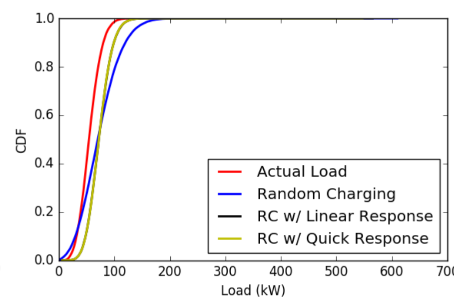

(c) Bulky Battery Sizes

Fig. 11: Load Distribution at Utility Node with Different Battery Sizes.

References

[1] Mikhail A. Lisovich, Deirdre K. Mulligan, and Stephen B. Wicker, "Inferring Personal Information from Demand-Response Systems," IEEE Security and Privacy, 2010.

[2] G. W. Hart, "Nonintrusive appliance load monitoring," Proceedings of the IEEE, 1992.

[3] Sherali Zeadally, Al-Sakib Khan Pathan, Cristina Alcaraz, and Mohamad Badra, "Towards Privacy Protection in Smart Grid," Wireless Personal Communications, 2012

[4] SmartGrid.gov, "Synchrophasor Applications in Transmission Systems," https://www.smartgrid.gov/recovery_act/program_impacts/ applications_synchrophasor_technology.html.

[5] Gaurav Lahoti, Daisuke Mashima and Wei-Peng Chen, "Customer-centric energy usage data management and sharing in smart grid systems," Proceedings of the first ACM workshop on Smart energy grid security, 2013.

[6] Georgios Kalogridis, Costas Efthymiou, Stojan Z. Denic, Tim A. Lewis and Rafael Cepeda, "Privacy for Smart Meters: Towards Undetectable Appliance Load Signatures," Smart Grid Communications (SmartGridComm), 2010.

[7] Stephen McLaughlin, Patrick McDaniel, and William Aiello, "Protecting Consumer Privacy From Electric Load Monitoring," ACN CCS, 2011.

[8] Weining Yang, Ninghui Li, Yuan Qi, Wahbeh Qardaji, Stephen McLaughlin and Patrick McDaniel, "Minimizing Private Data Disclosures in the Smart Grid," ACM CCS, 2012.

[9] Georgios Kalogridis, Zhong Fan, and Sagar Basutkar, "Affordable privacy for home smart meters," Parallel and Distributed Processing with Applications Workshops (ISPAW), 2011

[10] Department of Energy, "GridLAB-D," "https://www.gridlabd.org/".

[11] Patrick McDaniel and Stephen McLaughlin, "Security and Privacy Challenges in the Smart Grid," IEEE Security & Privacy, 2009.

[12] Stephen McLaughlin, Dmitry Podkuiko, Sergei Miadzvezhanka, Adam Delozier and Patrick McDaniel, "Multi-vendor Penetration Testing in the Advanced Metering Infrastructure," Proceedings of the 26th Annual Computer Security Applications Conference, 2010.

[13] Barry Hayes, Igor Melatti, Toni Mancini, Milan Prodanovic, and Enrico Tronci, "Residential Demand Management Using Individualized Demand Aware Price Policies," IEEE Transactions on Power Systems, 2017.

[14] SmartGrid.gov, "Smart Grid Investment Grant Program," "https://www.smartgrid.gov/recovery_act/program_impacts/peak_demand_and_electricity_consumption.html".

[15] Department of Energy, "Automated Demand Response Benefits California Utilities and Commercial & Industrial Customers," "https://www.smartgrid.gov/files/C6-Honeywell-final-draft-091814.pdf".

[16] Nerea Ruiz, Inigo Cobelo, and Jose Oyarzabal, "A Direct Load Control Model for Virtual Power Plant Management," IEEE Transactions on Power Systems, 2009.

[17] Stephen McLaughlin, Dmitry Podkuiko, Adam Delozier, Sergei Miadzvezhanka and Patrick McDaniel, "Embedded Firmware Diversity for Smart Electric Meters," USENIX HotSec, 2010.

[18] Indiana Office of Utility Consumer Counselor, "A Small Direct Load Control Switch," "www.in.gov/oucc/2391.htm".

[19] Clearly Energy, "Residential Demand Response Programs," "https://www.clearlyenergy.com/residential-demand-response-programs".

[20] Leehter Yao, and Hau-Ren Lu, "A Two-Way Direct Control of Central Air-Conditioning Load Via the Internet," IEEE Transactions on Power Delivery, 2009.

[21] BPL Global, "Connected Energy," "http://www.bplglobal.net/".

[22] Aditya Mishra, David Irwin, Prashant Shenoy, and Ting Zhu, "Scaling distributed energy storage for grid peak reduction," Proceedings of the fourth international conference on Future energy systems, 2013

[23] Aditya Mishra, David Irwin, Prashant Shenoy, Jim Kurose, and Ting Zhu, "SmartCharge: cutting the electricity bill in smart homes with energy storage," Proceedings of the 3rd International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet, 2012

[24] Energy Informative, "Grid Systems," "http://energyinformative.org/grid-tied-off-grid-and-hybrid-solar-systems/".

[25] UMass Trace Repository, "Smart* Data Set for Sustainability," "http://traces.cs.umass.edu/index.php/Smart/Smart".