Privacy Disclosure Through Smart Meters: Reactive Power Based Attack and Defense

Jingyao Fan
The Pennsylvania State University
University Park, PA, USA
Email: jfan@cse.psu.edu

Qinghua Li
University of Arkansas
Fayetteville, AR, USA
Email: qinghual@uark.edu

Guohong Cao
The Pennsylvania State University
University Park, PA, USA
Email: gcao@cse.psu.edu

***Abstract*—Smart meters can record fine-grained power consumption data and provide such data to the power supplier through realtime communications. Although smart meters can make power management more efficient and fault-tolerant, they also pose bigger threats to user privacy. Data from smart meters contain fine-grained power consumption information of home appliances and thus can be used to infer the ON/OFF states of home appliances. This problem has received some attention in the literature; however, most of them focus on active power based attacks. This paper focuses on reactive power and demonstrates how attackers can exploit reactive power data to infer appliance usage information. Experiments on real residential smart meter data show that our proposed attack can identify the ON/OFF events of home appliance with high accuracy. To protect users against such attacks, a novel defense technique called Reactive Power Obfuscation (RPO) is proposed. RPO can mask the true reactive power demand from the smart meter by using a capacitor to store and provide reactive power in a controlled manner. We evaluate the performance of RPO based on real household power consumption data. Evaluation results show that the ON/OFF events of home appliances can hardly be revealed from reactive power data when RPO is applied.**

***Keywords*-Smart Meter, Privacy Leakage, Reactive Power**

I. Introduction

Smart meters are being deployed in more and more households [1]. With low-cost microprocessors, solid-state circuits and realtime communication ports, smart meters can measure, store and automatically upload richer and more fined-grained power consumption data [2]. With such data, the power supplier will have better situation-awareness and make timely responses, resulting in better energy efficiency and better fault tolerance [3], [4].

The proliferation of smart meters also raises many security and privacy concerns [5], [6], [7], [8], [9]. Home appliances may have unique power consumption signatures. From the fine-grained power consumption data provided by smart meters, one may be able to infer if a certain home appliance is on or off based on its unique signature [10], [11]. From the home appliance usage information, user behavior such as doing laundry, cooking, or sleeping can be inferred [12].

Some researchers [10], [13], [11] proposed techniques to identify the ON/OFF events of home appliance from the data generated by smart meters, and exploit such information to infer user behavior. Among them, some works [10], [11] analyzed the household power usage data, and extracted some user behaviors and habits through activity recognition and user profiling. Greveler *et al.* [13] analyzed the power consumption of televisions. With detailed power demand information of different programs, they identified the program being played by a television. The aforementioned attacks all rely on active power to infer user behavior. Several defense techniques against active power-based attacks have been proposed [14], [15], [16]. They use load-controlled rechargeable batteries to mask the active power consumption data from smart meters. With these techniques deployed, the above active power-based attacks can be mitigated.

Despite the large amount of work on active power-based attacks and defenses, an open question that has received little attention is whether the *reactive power data alone* measured by smart meters can cause privacy leakage, and if so, how to prevent such privacy leakage? This is an important problem to study since most of today's smart meters measure not only active power but also reactive power.

In this paper, we identify a new attack which only leverages the reactive power data of smart meters and show that reactive power alone can still be exploited to breach user privacy. The attack can extract reactive power-based appliance signatures and then, based on each appliance's unique signature, identify the ON/OFF events of appliances from the load profiles measured by a smart meter. To mitigate privacy disclosure caused by the new attack, we propose a novel technique to obfuscate the reactive power measured by smart meters. The *Reactive Power Obfuscation (RPO)* technique uses capacitors installed within a user's household to store and provide reactive power in a controlled manner, to smooth power fluctuations, so that the reactive power changes caused by the ON/OFF of appliances are hidden from the smart meter.

The contributions of this paper are summarized as follows:

- We are the first to propose a reactive power based attack and demonstrate that reactive power data alone from smart meters can be exploited to infer appliance usage information and pose bigger threats to user privacy. The attack is demonstrated with real residential power consumption data, and is shown to be able to identify appliances with high accuracy.
- We propose a novel RPO technique to address the privacy concerns caused by reactive power. RPO uses capacitors to store and provide reactive power through a novel con-

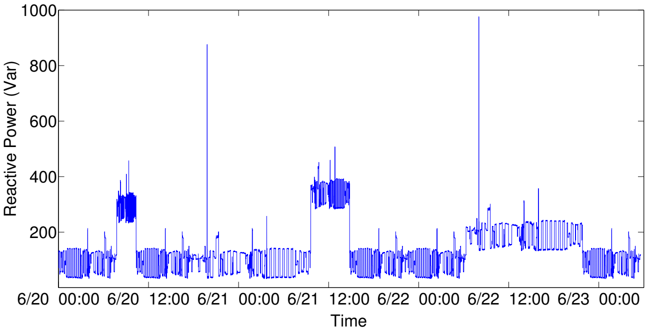

Fig. 1: A three-day reactive power load profile for a two-bedroom apartment

trol system consisting of feedback circuits, amplification circuits and controllable charging circuits.

- We build a circuit simulator and evaluate the performance of RPO using multiple privacy metrics. Evaluation results based on real residential power consumption data show that RPO can effectively prevent privacy leakage from reactive power data.

The rest of the paper is organized as follows. Section II introduces related work. Section III describes the system model and the threat model. In Section IV, we present the reactive power-based attack. Section V describes the proposed RPO technique. Section VI evaluates the performance of RPO, and Section VII concludes the paper.

II. Related Work

Since load monitoring techniques can be used to identify the appliances based on their unique power consumption characteristics, we first present related work in nonintrusive load monitoring before introducing related work in smart grid privacy.

A. Nonintrusive Load Monitoring

Nonintrusive load monitoring techniques are algorithms that can disaggregate a household's electric load profile to infer the ON/OFF states of home appliance. For traditional electric grid, many non-intrusive load monitoring techniques [17], [18], [19] have been proposed to process power consumption data collected from specialized sensors and identify appliances. They basically rely on unique power consumption characteristics of each appliance and use different signal processing techniques and classification algorithms to identify the appliances.

B. Smart Grid Privacy

With fine-grained data provided by smart meters, traditional attacks may pose bigger threats to user privacy. Privacy issues caused by smart meters have been studied extensively in the literature [20], [13], [10], [11]. Wang and Lu [20] gives an overview of potential privacy threats of smart meter data and gives a general idea of how such data can be used to answer questions about users' life. In [13], [10], [11], various techniques have been proposed to identify appliance and extract user activities and habits through activity recognition. However, all these works rely on active power data. The attack proposed in this paper only relies on reactive power data.

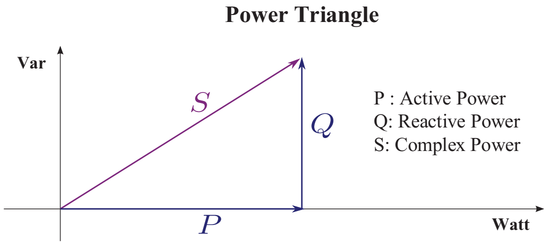

Fig. 2: Relationship between active power, reactive power and complex power

To protect user privacy, many privacy preserving techniques have been proposed [21], [22], [23]. In [21], [22], smart meter readings sent to the server are anonymized to protect user privacy. Privacy-preserving techniques [23], [24], [25] can also be applied to aggregate the meter data of nearby users so that no individual user's power consumption can be learned by the server. However, all these works require interactions with either a third party or other users and thus they bring additional threats. Moreover, they require changes to smart meters and the power control system and thus introduce extra cost.

To avoid interactions with any third party while preserving user privacy, battery-based techniques have been proposed [14], [15], [26]. They mask the true power demand of a household by directing rechargeable batteries to charge and discharge so that the net demand does not change too much. However, their solutions can only be used to mask the active power consumption of a household, and they cannot be used to address reactive power-based privacy attacks. Different from these existing works, the defense technique proposed in this paper can protect users from reactive power-based attacks. Active power and reactive power are orthogonal in the vector space and they do not interfere with each other. Thus, our solution can be combined with battery-based approaches to deal with privacy leakage from smart metering data; i.e., capacitor is used to store and provide reactive power data to mask reactive power usage while battery is used to mask the active power usage.

III. Preliminaries

A. Load Profiles

Load profiles are time series of electric demand. Most recent smart meters can measure and store very fine-grained load profiles with reactive power, such as GE I-210 meters [27] and Landis+Gyr meters [28]. The load profiles used in this paper are collected with a GE I-210 meter at three different types of residences: two bedroom apartment (Apt 1), one bedroom apartment (Apt 2) and a two-story house (H). Apt 1 consists of two bedrooms, 1.5 bathrooms with two adults living in. Apt 2 consists of one bedroom and one bathroom with two adults living in. H consists of four bedrooms and 2.5 bathrooms, with two adults and two kids living in.

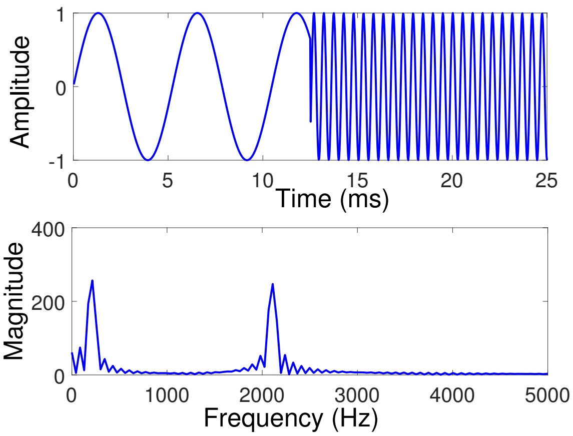

Fig. 3: An example of Fourier transform

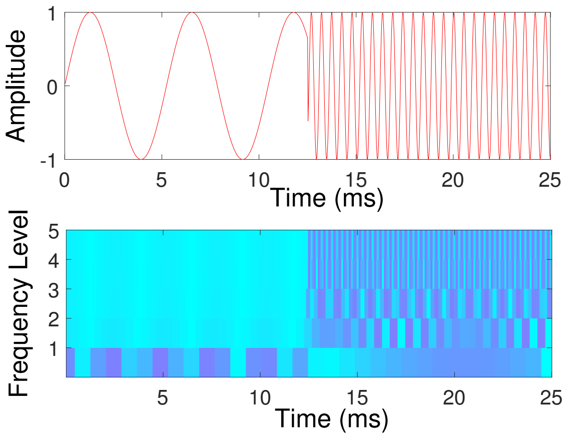

Fig. 4: An example of wavelet transform

The load profiles for Apt 1 were collected from June 1^{st} to 30^{th}, 2016. The load profiles for Apt 2 were collected from February 8^{th} to 22^{th}, 2016. The load profiles for H were collected from March 10^{th} to 17^{th}, 2016. All load profiles contain reactive power at the granularity of one second. There are $53 \times 24 \times 3600 = 4,579,200$ records in our collected load profiles. A three-day load profile with reactive power from June 20^{th} to 22^{th}, 2016 is shown in Fig. 1.

B. Reactive Power and Capacitor

In an Alternating Current (AC) circuit, both voltage and current are sinusoidal. For resistive elements, the voltage and current are always in the same phase, resulting in net transfer of energy. This portion of energy is called active power (or real power). For energy storage elements such as inductors and capacitors, their voltage and current are always 90 degrees out of phase. The energy flowing to these elements is stored and will flow back to the source periodically, without any net energy transfer. This portion of energy is called *reactive power*. Most appliances consist of both resistive elements and energy storage elements, and thus both active power and reactive power flow to loads. The vector sum of active power and reactive power is called *complex power*, as shown in Fig. 2.

As mentioned above, capacitors are energy storage elements and do not really consume net power. They only cause energy transfer between the power source and themselves. We will explain how the capacitors cycle between providing and storing power.

Let C denote the capacitance of the capacitor. In an AC circuit with voltage source $v(t) = V_0 \cos(\omega t)$, based on the power and current law, the current flows through the capacitor is:

$$i(t) = C\frac{dv(t)}{dt}\omega C V_0 \sin(\omega t) = \omega C V_0 \cos(\omega t + 90^\circ)$$

According to the power definition, instantaneous power of the capacitor can be calculated as:

$$p(t) = v(t)i(t) = \omega C V_0^2 \cos(\omega t)\cos(\omega t + 90^\circ)$$

When $p(t)$ is positive, the capacitor charges (storing power) and when $p(t)$ is negative, the capacitor discharges (providing power).

C. Appliance Specifications

The targeted appliances in our three households are: lamp, refrigerator, air conditioner, microwave oven, dishwasher, kettle, laptop, and television. Their power consumption specifications are listed in Table I. The above appliances are chosen in our experiments because they are commonly used in households. Since the signature extraction process is the same for all appliances, the proposed attack is not limited to the above appliances.

D. Fourier Transform and Wavelet Transform

Fourier transform is a very important signal processing technique [29]. It decomposes a function of time into frequency components. The Fourier transform $\hat{f}$ of an integrable function f $(\mathbb{R} \to \mathbb{C})$ is defined as:

$$\hat{f}(\xi) = \int_{-\infty}^{\infty} f(x)\ e^{-2\pi i x \xi}\, dx, \forall\, \xi \in \mathbb{R}.$$

An example of Fourier Transform on a time-domain signal is shown in Fig. 3. We can see from the lower figure that the original signal mainly consists of 200 Hz and 2100 Hz components.

Different from Fourier transform which gives only frequency information (what frequency components are in the signal), wavelet transform can give both time and frequency information: what frequency components are in the signal and when these particular spectral components occur [30]. Wavelet transform of an integrable function f is defined as:

$$\left[W_\psi f\right](a,b) = \frac{1}{\sqrt{|a|}}\int_{-\infty}^{\infty}\overline{\psi\left(\frac{x-b}{a}\right)}f(x)dx.$$

Here, $a = 2^{-j}$ is the binary dilation (scaling), and $b = k \cdot 2^{-j}$ is the binary position (time). The wavelet coefficients c_{jk} are calculated as:

$$c_{jk} = \left[W_\psi f\right]\left(2^{-j}, k \cdot 2^{-j}\right).$$

TABLE I: Appliance Specifications

Appliances	Apt 1		Apt 2		H	
	Active Power	Reactive Power	Active Power	Reactive Power	Active Power	Reactive Power
Lamp	15 W	18 Var	15 W	12 Var	25 W	20 Var
Refrigerator	110 W	170 Var	150 W	215 Var	265 W	305 Var
Air Conditioner	900 W	120 Var	1200 W	180 Var	2000 W	220 Var
Microwave Oven	1000 W	450 Var	800 W	310 Var	1350 W	530 Var
Dishwasher	1100 W	300 Var	700 W	190 Var	1200 W	320 Var
Kettle	1200 W	10 Var	1400 W	12 Var	1800 W	15 Var
Laptop	28 W	14 Var	30 W	15 Var	35 W	15 Var
TV	48 W	25 Var	75 W	35 Var	105 W	44 Var

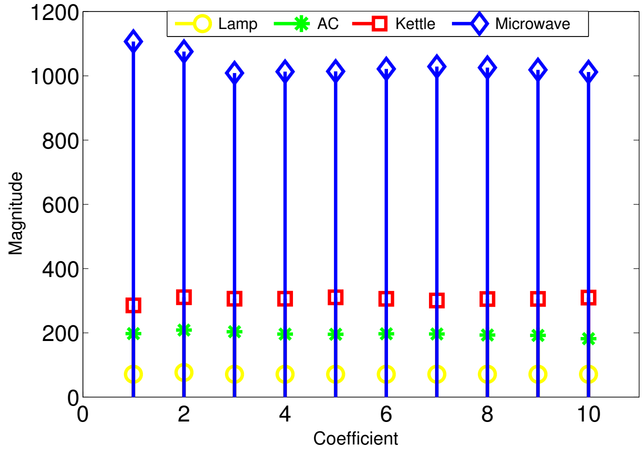

Fig. 5: Wavelet coefficients: using the first 10 wavelet coefficients as appliance signatures

An example of wavelet transform on a time-domain signal is shown in Fig. 4. The shade of color in the figure corresponds to the magnitude of the wavelet coefficients. The darker the color, the bigger the coefficient is. As shown in the lower figure, the first half of the original signal consists of low frequency components and the second half of the signal mainly consists of high frequency components.

E. Threat Model

We assume that the attacker can only obtain reactive power data from the smart meter. Active power is not considered here since the focus of this paper is reactive power-based attack as stated in Section I, and active power-based attack has been addressed in the literature. The attacker can obtain the power consumption signatures of the appliances used in the target user's residence. This assumption is valid since there are many ways for the attacker to obtain such information. Existing research has shown that such information can be leaked to malicious third parties [14]. Attackers may also run these appliances at home to obtain signatures from meter readings or get them from manufactures provided information. Even if the attacker cannot obtain such information directly, he can still learn appliance signatures through training techniques [31].

IV. Reactive Power-Based Attack

In this section, we present the privacy attack which can be used to infer appliance usage information from only reactive power data.

A. Attack Approach

1) Appliance Signature Extraction: For each appliance, we extract an one-minute window from the beginning of its reactive power waveform. One minute is enough to capture the characteristics of appliances because most appliances reach a stable state in less than one minute. Then wavelet transform is performed on the extracted piece of waveform to get the wavelet coefficients. The wavelet coefficients form a time-frequency representation of the original reactive power waveform. They represent the frequency distribution of each appliance when they are turned on, and thus these signatures can capture the characteristics of each appliance. During the experiments, we find that the first 10 coefficients are enough to distinguish appliances and thus the first 10 wavelet coefficients are used as the appliance signature. A larger number can be chosen, but it will increase the computation overhead. Examples of appliance signatures are given in Fig. 5.

2) Appliance Identification: The appliance identification process consists of three steps: filtering, event detection and identification.

Filtering In the load profiles, there are events that appear to be turning ON/OFF appliances but are actually not. These events are called *deceptive events*. Deceptive events are usually caused by the periodic wake-up of appliances such as refrigerator and air conditioner. Since the period of these events is usually relatively long (at the hour level), they compose the low frequency components in the waveform. There is also some noise in the load profiles, i.e., impulse that usually occurs at the beginning of a wave. These noise composes the high frequency components in the waveform. To filter out deceptive events and noise, a Fourier transform is first performed on the waveform to get the frequency distribution. From the frequency distribution along with the original waveform, we identify the estimated frequency of deceptive events and noise. Then a band-pass filter [32] is applied on the waveform to filter out low frequency deceptive events and high frequency noise. The new waveform obtained only contains components that are relevant to appliances. This new waveform is called the *post-filtering waveform*. This filtering process is shown in Fig. 6.

Event Detection After deceptive events and noise are filtered out, we need to detect real events that are caused by

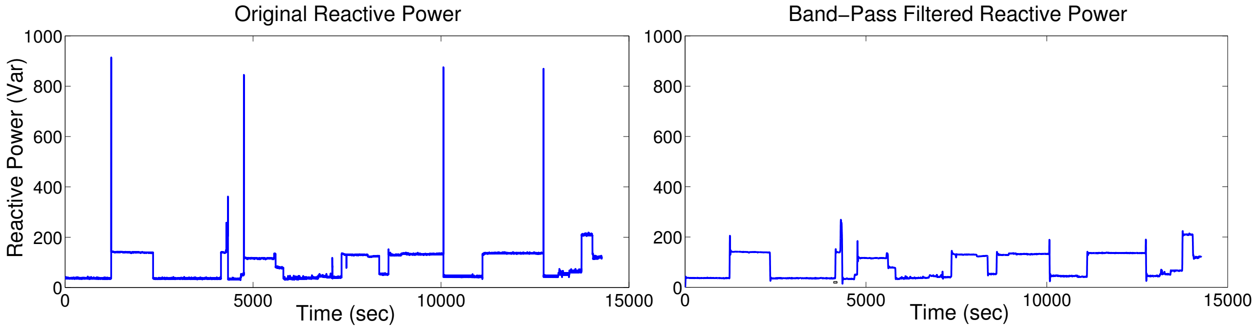

Fig. 6: Filtering out deceptive events and noise with band-pass filter

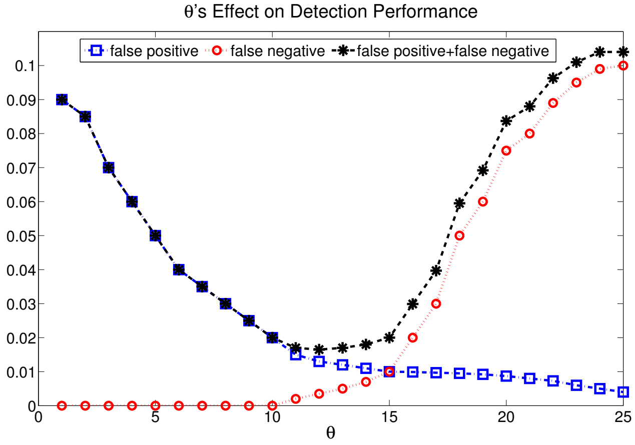

Fig. 7: The effects of θ on the detection accuracy

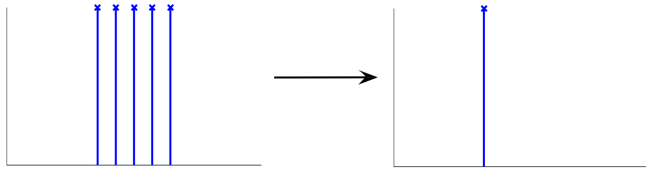

Fig. 8: Redundant Edges where only the first UpEdge is kept

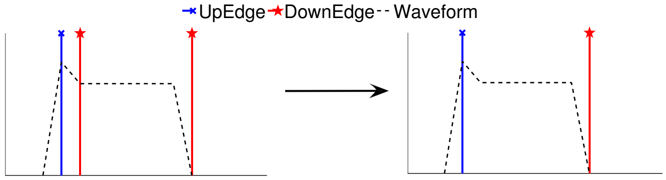

Fig. 9: Noisy Edge caused by a jitter

appliances. When an event happens, an edge always appears in the waveform. Therefore, the key is to detect edges in the post-filtering waveform. To reduce false detection, the detection process distinguishes between the edges going up (UpEdge) and the edges going down (DownEdge). For any two consecutive points in the waveform W, we compute: $\Delta_i = W[i+1] - W[i]$. If Δ_i is larger than the chosen threshold θ, an UpEdge occurs at time i in W. If Δ_i is smaller than $-\theta$, a DownEdge occurs at time i. Obviously, the choice of θ is very important for the performance of event detection. If θ is too small, noise may be mistreated as events. If θ is too big, events triggered by low-power appliances may be ignored. Fig. 7 shows how the detection accuracy changes with θ. As can be seen, with the increase of θ, the false positive rate decreases and the false negative rate increases. In the experiments, θ is set to be 15 since at this point a balance can be achieved between the false positive rate and the false negative rate and the total false detection rate is relatively low. Our waveform is in the granularity of second. Sometimes an edge in the waveform lasts for several seconds, and each second will be recognized as an individual edge. Thus, there will be redundant edges which should be removed. There are two cases of redundancy:

- Consecutive UpEdges: When there are a series of consecutive UpEdges, only the first UpEdge is kept, as shown in Fig. 8.
- Consecutive DownEdges: When there are a series of consecutive DownEdges, only the last DownEdge is kept. This DownEdge captures more information than others so that no useful frequency component is lost.

Here by consecutive we mean that the edges are next to each other in time, i.e. with time interval of less than 5 seconds.

The reactive power is unstable compared with active power. There are always jitters in the waveform when the reactive power changes. Most jitters will be removed during the filtering process. However, some small jitters still exist. The jitters will result in noisy edges that should be ignored. Fig. 9 shows that a jitter occurs when the appliance is turned ON. The waveform (the black dashed line) rises when the appliance is turned ON and drops a little bit. It then stays stable for a while and falls to zero when the appliance is turned OFF. As shown in Fig. 9, three edges are detected: an UpEdge (blue line) when the waveform rises, a DownEdge (red line) right after the UpEdge and a second DownEdge when the waveform falls to zero. The first DownEdge does not represent a real change in appliance state and thus is a noisy edge. We remove the noisy DownEdge and only keep the UpEdge and the second DownEdge. A jitter can also occur when the appliance is turned OFF. Then, similar to the case in Fig. 9, the noisy

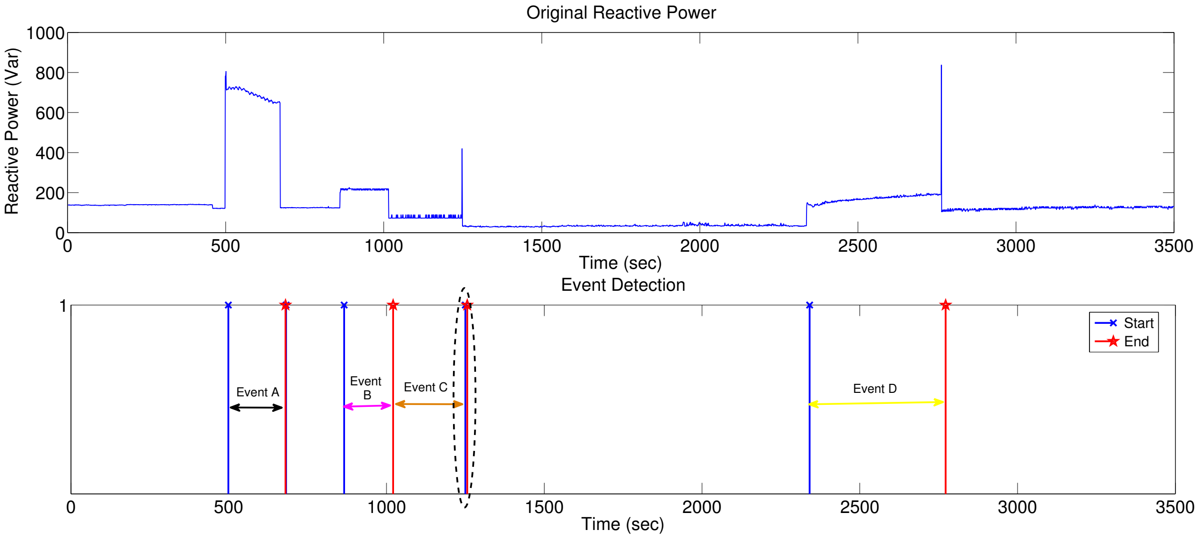

Fig. 10: Extracting waveforms between matching UpEdges/DownEdges

UpEdge right before the DownEdge is removed.

Identification In this step, the appliance that triggers each UpEdge and DownEdge is identified. First, the reactive power waveforms between each pair of UpEdge and DownEdge are extracted. For each piece of extracted waveform, wavelet transform is performed to get the wavelet coefficients. Then the obtained coefficients are compared with the appliance signatures (such as those shown in Fig. 5) to determine which appliance corresponds to this piece of waveform. This identification process is demonstrated in Fig. 10 and Fig. 11. In Fig. 10, we identify the pieces of waveforms needed to be extracted based on the matching UpEdges and DownEdges. As shown in the figure, there is a pair of UpEdge and DownEdge that are consecutive around 1250^{th} second (marked by the black dashed line). It is unlikely that an event only lasts for a second and thus this pair of UpEdge and DownEdge cannot be processed alone. For pairs like this, there are two choices: combining this pair with the previous pair (event C marked by the orange arrow) or ignoring it (event B marked by the magenta arrow). Waveform pieces for both choices are extracted.

In Fig. 11, we match the coefficients of each extracted waveform piece (events A,B,C,D marked by arrows in Fig. 10) with the appliance signatures shown in Fig. 5 using least square method. Event A in Fig. 10 is matched with the signature of microwave oven. Event D in Fig. 10 is matched with the signature of air conditioner. Event B and event C in Fig. 10 both match with the signature of kettle.

B. Experiments

We launched reactive power based attacks against the load profiles collected at the three households. The identification results based on the attacks are compared with the real usage events logged by users to get the false positive rate and false negative rate for each appliance. The results are shown in Table II. We can see that most appliances can be identified with low false positive/negative rates. For example, in Apt 1, the false positive rate can reach 4.2% for lamp, and 1.3% for microwave oven.

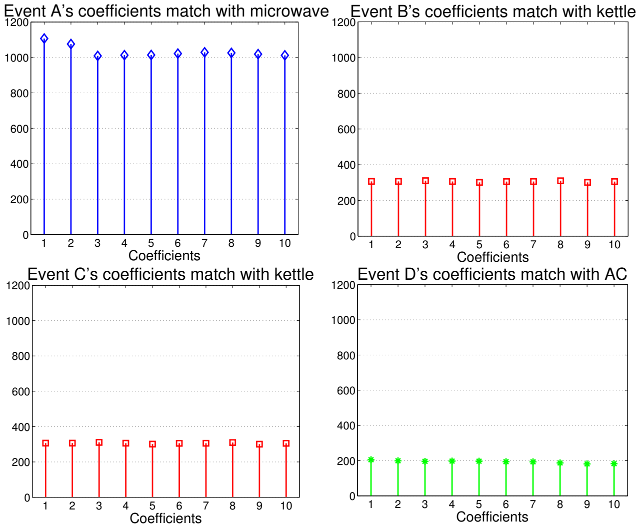

Fig. 11: Matching coefficients of waveforms in Figure. 10 with the appliance signatures in Figure 5

The attack has relatively higher false positive/negative rates for refrigerator and air conditioner, around 15%, due to the following reasons. Appliances such as air conditioner and refrigerator have multiple working states. Different working states consume different amounts of power. For example, the air conditioner consumes more power when it is cooling down the room and consumes less power after the room reaches the preset temperature. The switches of working states may not be matched to signatures sometimes and will result in misdetections.

From the table, we can see that the false positive/negative rates for air conditioner and refrigerator are still pretty low,

TABLE II: The Performance of Reactive Power Based Attack

Appliances	Apt 1		Apt 2		H	
	False Positive	False Negative	False Positive	False Negative	False Positive	False Negative
Lamp	4.2%	5.8%	3.7%	5%	11%	13%
Refrigerator	11.7%	13.5%	9.8%	12.6%	13.5%	14.7%
Air Conditioner	14.5%	16.5%	13.3%	15.7%	17.4%	18.3%
Microwave Oven	1.3%	2.7%	1%	2%	5.7%	7.5%
Dishwasher	∼0	2.6%	0.7%	2.5%	4.3%	5.8%
Kettle	1%	1.7%	1.3%	2%	6.5%	8.1%
Laptop	3.9%	5.6%	3.4%	5.3%	9.4%	8.8%
TV	3.3%	6.7%	3%	5.2%	8.4%	9.3%
Overall	5.8%	7.5%	4.2%	5.9%	7.3%	9.7%

TABLE III: Notations

R_c	The target reactive power
t	Time variable in the power demand function
$r(t)$	The net reactive power demand of appliances
$u(t)$	The obfuscated reactive power demand read by the smart meter
$c(t)$	The amount of reactive power coming in/out of the capacitor
V_{bd}	Breakdown voltage of capacitor
C	Capacitance of the capacitor

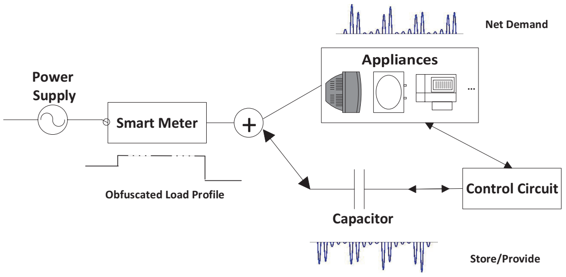

Fig. 12: RPO system model

around 15%. Moreover, the attackers can use context information to improve the detection accuracy because the usage patterns of such appliances are related to the context. For example, there is a high probability that the air conditioner is working on a hot summer day. The refrigerator is more likely to be working at high power level on weekends because people usually do grocery shopping on weekends. Based on information like this (obtained from weather history), probability-based attacks can be applied to further improve the detection accuracy.

Even without considering these context information, as shown in Table II, the good performance of our attack demonstrates that reactive power alone can reveal user privacy and new defense techniques are needed.

V. Reactive Power Obfuscation

In this section, we first present the design of RPO and then evaluate its performance.

A. System Model and Basic Idea

Fig. 12 shows the system model of RPO, which mainly consists of a smart meter, home appliances, a controllable capacitor with relatively large capacitance, and a control circuit. The general idea is as follows. The control circuit controls the capacitor to provide or store reactive power so that the reactive power fluctuations caused by appliances (net demand) will be smoothed out, and will not be shown in the smart meter. By reducing the power fluctuations measured by the smart meter, less appliance usage information can be inferred by the attackers. Thus, RPO aims to reduce the power fluctuations by maintaining the reactive power at a target value. Since the power usage of households may change dramatically and there are some physical limitations in the electric systems, RPO has to address many technical challenges.

B. Design of RPO

To prevent the leakage of appliance ON/OFF information from reactive power data, RPO aims to keep the reactive power around a target value. To achieve this goal, RPO has three processes: initialization, maintaining, and adjusting, which are introduced below using notations in Table III.

- Initialization Process: RPO initializes the target reactive power R_c based on the household's power usage history and the capacitance of the capacitor.
- Maintaining Process: For any given time t, RPO controls the capacitor to store and provide reactive power based on the relationship between the net demand $r(t)$ and the target demand R_c. When $r(t)$ rises beyond R_c, the capacitor should start providing reactive power to compensate for the extra load. When $r(t)$ drops below R_c, the capacitor should start storing reactive power to consume extra power supply.
- Adjusting Process: When the capacitor is not able to maintain the current target demand, RPO adjusts R_c to a new value.

1) Initialization Process: The target reactive power R_c affects the performance. If R_c is not properly chosen, it will have to be frequently adjusted. When selecting the initial value of R_c, the capacitor should have enough capacitance to keep the reactive power demand maintained at R_c based on the given usage history. Also, the capacitor should not exceed its breakdown voltage V_{bd}. According to the principles of electric

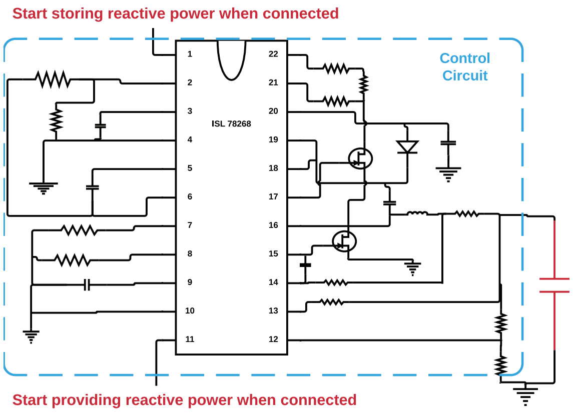
Start storing reactive power when connected
Control Circuit
ISL 78268
1
2
3
4
5
6
7
8
9
10
11
22
21
20
19
18
17
16
15
14
13
12
Start providing reactive power when connected

Fig. 13: Storing/providing module of the control circuit

circuits, the above conditions can be written as follows:

$$r(t) + c(t) = R_c \; and \; |c(t)| < \frac{1}{2}CV_{bd}^2$$

Thus, we have the range of R_c:

$$[\max r(t) - \frac{1}{2}CV_{bd}^2, \frac{1}{2}CV_{bd}^2 + \min r(t)]$$

For simplicity, we initialize R_c to be the midpoint of the range.

2) Maintaining Process: The maintaining process is the main part of RPO. There are two main modules at the circuit level: Decision Making and Storing/Providing. The decision making module decides whether the capacitor should store or provide reactive power. The storing/providing module controls the capacitor to store/provide the right amount of reactive power.

Decision Making Module Whether the capacitor should store or provide reactive power depends on the relationship between $r(t)$ and R_c.

$$r(t) - R_c \begin{cases} > 0: & \text{start providing reactive power;} \\ \leq 0: & \text{start storing reactive power.} \end{cases}$$

$r(t) - R_c$ serves as the control signal for the Storing/ Providing module. When $r(t) - R_c$ is positive, it signals the Storing/Providing module to start providing reactive power; when $r(t) - R_c$ is negative, it signals the Storing/Providing module to start storing reactive power.

Storing/Providing Module The storing and providing module directs the flow of the reactive power from/to the capacitor. There are various methods to charge/discharge the capacitor. Since our goal is to maintain the reactive power of the household at a stable level, we choose the constant current/constant voltage (CICV) method. At the beginning of the charge cycle, our storing and providing module should operate in constant current mode. In this way, a constant current is provided to the capacitor such that its voltage increases linearly. When the capacitor is charged to a target voltage, the our module should enter a constant voltage loop and accurately controls the capacitor charge level to be constant to avoid over charging. To control the capacitor in the above way, the control circuit we design must satisfy the following:

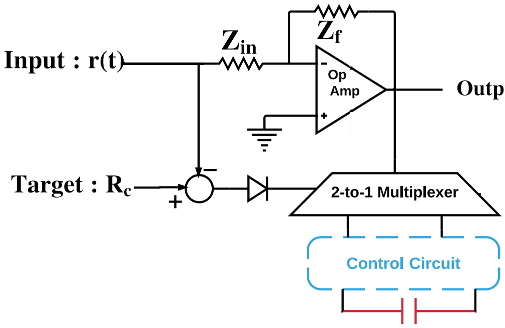

Fig. 14: Combining the three modules with a negative feedback amplifier

- The control circuit can operate at two regulation modes: constant current and constant voltage.
- Transitions between different regulation modes should be automatic.
- The control circuit should have an accurate current sense amplifier that can operate at supply voltage of the system.

To achieve this goal, we use the Intersil's ISL78268 controller [33], which consist of MOSFET driver, amplifiers, didoes and flip-flops, as shown in Fig. 13. ISL78268 is used as the controller for the large capacitance capacitor (red in the figure). Other components in the figure are used to connect the controller with the capacitor so that the capacitor can store reactive power or provide reactive power as needed. The maximum current allowed in ISL78268 is 3 A, which means that there will be an upper-bound of the reactive power stored in the capacitor. If more reactive power needs to be reserved, a group of ISL78268 controllers can be combined in parallel to achieve larger capacity.

In order for RPO to work as designed, the two modules should be combined functionally so that the reactive power from the capacitor can offset the reactive power of the home appliances. In other words, we need to add the providing amount to the net demand or subtract the storing amount from the net demand. To achieve this goal, an amplifier is used as shown in Fig. 14. In the figure, the two resistance R_{in} and R_f and the amplifier form a negative feedback circuit so that reactive power can be added to or subtracted from the net demand according to the relationship between $r(t)$ and R_c. According to the principles of electric circuits:

$$r(t) - R_c \begin{cases} > 0: & u(t) = r(t) - R_{in}/R_f \cdot |c(t)| \\ \leq 0: & u(t) = r(t) + R_{in}/R_f \cdot |c(t)| \end{cases}$$

If R_c and the capacitor are chosen properly, the obfuscated reactive load $u(t)$ will not change much when the states of appliances are changed.

3) Adjusting Process: The initialization of R_c is based on the usage history. It is possible that some time later, the power consumption of the appliances $r(t)$ changes dramatically compared to historic data and the capacitor fails to maintain $u(t)$ at R_c. There are two failure cases:

- High Demand: When the current reactive power demand of appliances is very high, even if the capacitor releases all the stored power, the obfuscated reactive power demand still rises beyond the target level.

- Low Demand: When the current reactive power demand of appliances is very low, even when the capacitor is full, the obfuscated reactive power demand still falls below the target level.

When such failures occur, the reactive power load captured by the smart meter will change as the appliance usage pattern changes. In order to preserve privacy after such failures, R_c needs to be adjusted.

If a failure is caused by low demand, the capacitor is already fully charged but the overall reactive power demand $u(t)$ is still not enough (below R_c). R_c should be adjusted to a lower value. If a failure is caused by high demand, the capacitor has already released all its stored reactive power but $u(t)$ is still too high (above R_c). R_c should be adjusted to a higher value. In RPO, R_c is adjusted according to Algorithm 1.

Algorithm 1 Adjustment of R_c

```
1: Suppose at time $t_f$, a failure occurs and the maximum
   reactive power can be stored by the capacitor is $r_{max}$
2: if $u(t_f) > R_c$ then
3:     $R_c = u(t_f) + CR \cdot r_{max}$
4: end if
5: if $u(t_f) < R_c$ then
6:     $R_c = u(t_f) - CR \cdot r_{max}$
7: end if
```

Here, CR is a parameter. In our experiments, it is randomly chosen in the range [0.2,0.8]. Note that in line 6, $R_c = u(t_f) - CR \cdot r_{max} > u(t_f) - r_{max}$. The high demand failure only occurs when the capacitor is providing all its reactive power. Thus, $u(t_f) > r_{max}$ and R_c is assured to be positive.

4) Cost and Usability: The cost of RPO is a very important factor that should be considered since it will affect whether RPO can be deployed in practice. We evaluate the cost of RPO in two aspects: the power consumption of RPO and the monetary cost of its components.

The resistors used in RPO will consume energy (active power). According to the concept of active power and Ohm's law, we have:

$$P = \frac{1}{T}\int_0^T \frac{UI}{t}(1 - \cos 2\omega t)dt = UI = I^2R = \frac{U^2}{R}$$

where U is the average voltage over the resistor, I is the average current through the resistor, and R is the resistance. For those resistors that are connected to the circuit in series, small resistance (e.g., 0.05 ohm) should be chosen; for those resistors that are connected to the circuit in parallel, large resistance (e.g., 100k ohm) should be chosen. In this way, the active power consumed by each resistor is very small.

In US, the average voltage in residential power supply is $110\,V$. The maximum current allowed by ISL78268 is $3\,A$. Thus, the upper bound for U is $110\,V$ and the upper bound for I is $3\,A$. We can calculate the maximum possible power consumed by the resistors:

$$\text{For parallel resistors: } \frac{110^2}{10^5} = 0.121W;$$
$$\text{For serial resistors: } 3^2 \cdot 0.05 = 0.45W.$$

In our RPO design, we use 10 resistors connected in parallel and 5 resistors connected in serial. Thus, the maximum power consumption of RPO should be $10{*}0.121 + 5{*}0.45 = 3.46\,W$, which is less than $5\,W$. Since the power consumption of RPO is small, it will not introduce noticeable changes to the smart meter readings. With RPO deployed, the smart meter data can still perform its functions as desired.

RPO uses an ISL78268 controller, a supercapacitor with large capacitance and other small components such as resistors, amplifiers, inductors and diodes. The unit price for the ISL78268 controller is less than \$10. The super-capacitor needed for a household is about \$50-\$100. The unit price for other small components is at cent level (they usually sell in bulks) and the total cost should be less that \$5. Thus, the overall cost of RPO is less that \$150, which is acceptable, and can be largely deployed.

C. Compatibility of RPO

RPO focuses on addressing the privacy leakage caused by reactive power. Although it is not designed to address active power-based privacy attacks, it is compatible with existing battery-based approaches [34], [14], [15] which can mask active power usage from the smart meter.

As mentioned in Section III-B, active power and reactive power are orthogonal in the vector space and they do not interfere with each other. Thus, RPO can be combined with battery-based approaches to deal with privacy leakage from smart metering data; i.e., capacitor is used to store and provide reactive power data to mask reactive power usage while battery is used to mask the active power usage.

VI. Performance Evaluations

In this section, we first present the metrics used in the evaluation, and then present the evaluation results using these metrics.

A. Metrics

1) Vulnerability Factor: The attacker infers appliance usage information from the changes in the waveform, and better privacy can be achieved if fewer changes exist in the obfuscated waveform. To measure the maximum information leakage before and after RPO is implemented, we define *vulnerability factor* as the entropy of a reactive power waveform. In information theory, the entropy for signal X with probability mass function P is defined as:

$$\mathrm{H}(X) = -\sum_{i=1}^{n} \mathrm{P}(x_i)\log_b \mathrm{P}(x_i),$$

In this paper, $b = 2$ is used as the standard base. The probability mass function P for the load profile is defined

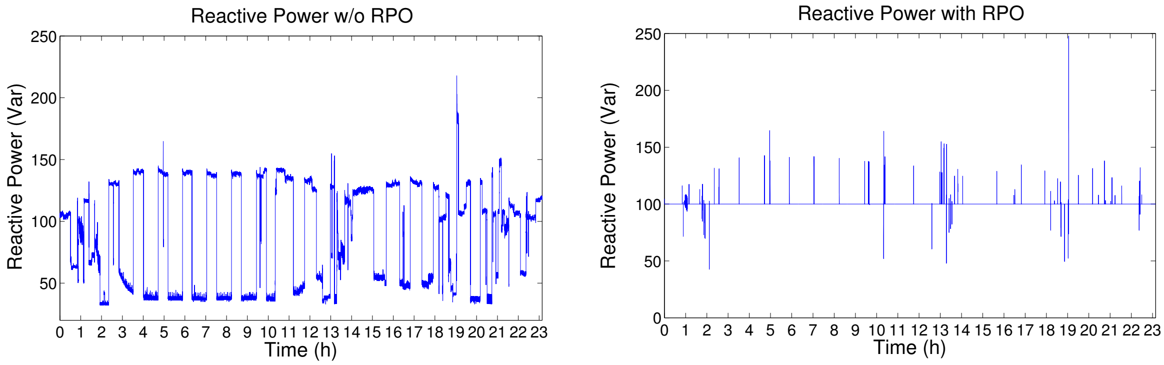

Fig. 15: Reactive Power Load Profiles Before/After RPO is adapted

TABLE IV: Vulnerability Factor and Obfuscation Factor

Apartment	Vulnerability Factor		Obfuscation Factor
	w/o RPO	w/ RPO	
Apt 1	2.307	0.0316	0.9544
Apt 2	5.334	0.0661	0.9985
H	9.043	0.1614	0.9365

as:

$$\mathrm{P}(X) = \frac{1}{n}\sum_{i=1}^{n}\mathrm{I}(x_i),$$

where $I()$ is an indicator function calculated as

$$I(x_i) = \begin{cases} 1 : & x = x_i \\ 0 : & \text{otherwise.} \end{cases}$$

Here, X is the reactive power load waveform and x_i is the reactive power load at time slot i. The entropy $H(X)$ can be interpreted as an upper bound on the information (the number of bits) that the attacker can extract from the reactive power waveform. A smaller vulnerability factor means less information leakage and better privacy.

2) Obfuscation Factor: The vulnerability factor measures to what degree the load profile can leak user privacy (i.e., user's appliance usage information). It shows the upper bound of the amount of information that can be inferred from the load profile. To further quantify how much privacy RPO can provide, it is desirable to show how different the obfuscated waveform and the original waveform are. Based on cosine similarity, a metric *obfuscation factor* (denoted by g) is defined as follows:

$$\begin{aligned} g =& 1 - \text{similarity} = 1 - \cos(\theta) \\ =& 1 - \frac{\mathbf{u(t)} \cdot \mathbf{r(t)}}{\|\mathbf{u(t)}\|\|\mathbf{r(t)}\|} = \frac{\sum\limits_{i=1}^{n} u(t_i)r(t_i)}{\sqrt{\sum\limits_{i=1}^{n} u(t_i)^2}\sqrt{\sum\limits_{i=1}^{n} r(t_i)^2}} \end{aligned}$$

where r denotes the original reactive power and u denotes the obfuscated reactive power given by RPO. It is easy to see that g measures the difference between the original waveform and the obfuscated waveform. The goal of RPO is to obfuscate the waveform as much as possible. Thus, we want g to be as large as possible. In other words, our obfuscation factor g can describe the information loss during obfuscation, and a larger g means better privacy.

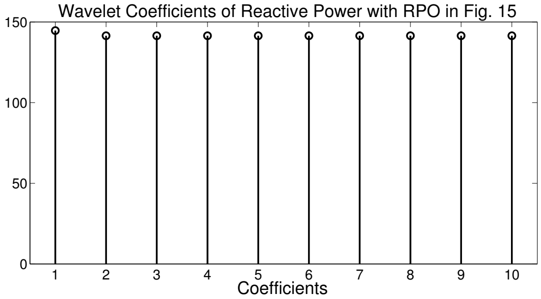

Fig. 16: Coefficients of Obfuscated Reactive Power Data

3) Detection Rate: To evaluate the effectiveness of RPO on mitigating reactive power based attacks, we define another metric *detection rate*, which measures how well the attacker can identify appliances:

$$\text{detection rate} = \frac{\text{number of accurate detection}}{\text{number of appliance usage event detected}}$$

Since an appliance can be turned on or off, for each appliance, there are two detection rates, corresponding to the ON and OFF events.

B. Simulation Setup

To evaluate the performance of RPO, we simulate the complete system using Cadence's OrCAD PSpice software [35]. OrCAD PSpice enables complex circuit design and high-performance circuit simulation. It supports both digital and analog circuits. In PSpice, we choose the mixed A/D circuit mode to set up our designed RPO circuit (as shown in Fig. 13 and Fig. 14). We create waveforms using the load profile data sets described in Section III-A, and add the waveforms as signal sources to the circuit (the input shown in Fig. 14). To capture the obfuscated reactive power waveform, we add a Power Dissipation Marker at the output line (shown in Fig. 14). In this way, the waveform we get from the marker will be the obfuscated reactive power. The results reported below were collected based on a simulated 10 F capacitor.

C. Evaluation Results

1) Original and Obfuscated Load Profiles: In this set of experiments, the initial target power R_c is set to $100\,Var$. As

TABLE V: The effectiveness of appliance recognition attacks with/without RPO

Appliances	detection rate of Apt 1				detection rate of Apt 2				detection rate of H			
	ON		OFF		ON		OFF		ON		OFF	
	No RPO	RPO	No RPO	RPO	No RPO	RPO	No RPO	RPO	No RPO	RPO	No RPO	RPO
Lamp	95.8%	1.2%	94.2%	1.4%	96.3%	1.5%	95%	1.7%	89.1%	1.6%	87.9%	1.5%
Refrigerator	88.3%	1%	86.5%	0.9%	91.2%	1.3%	87.4%	1%	86.5%	1.4%	85.7%	1.2%
AC	85.5%	0.8%	83.5%	0.5%	86.7%	1%	84.3%	0.6%	81.4%	1.2%	80.3%	0.8%
Microwave	98.7%	1.2%	97.3%	1.6%	99.1%	1.4%	98%	1.7%	93.7%	1.8%	92.5%	2%
Dishwasher	∼1	2%	98.4%	1.6%	99.3%	2.6%	97.4%	1.5%	95.3%	2.5%	93.8%	1.8%
Kettle	99%	1.3%	98.3%	1.7%	98.7%	1.1%	98.1%	0.9%	93.2%	1.3%	91.6%	1%
Laptop	96.1%	0.9%	94.6%	1.3%	96.6%	0.8%	94.7%	0.6%	89.6%	1.1%	88.2%	1.1%
TV	96.7%	0.7%	93.3%	1%	97%	0.5%	94.8%	1.1%	90.4%	1%	89.7%	0.9%
Overall	94.2%	0.9%	92.5%	1.1%	95.8%	1%	94.1%	0.8%	91.7%	1.2%	89.2%	1%

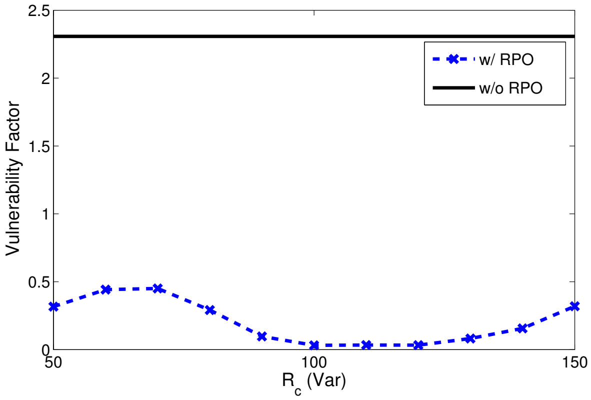

Fig. 17: R_c's effects on the vulnerability factor

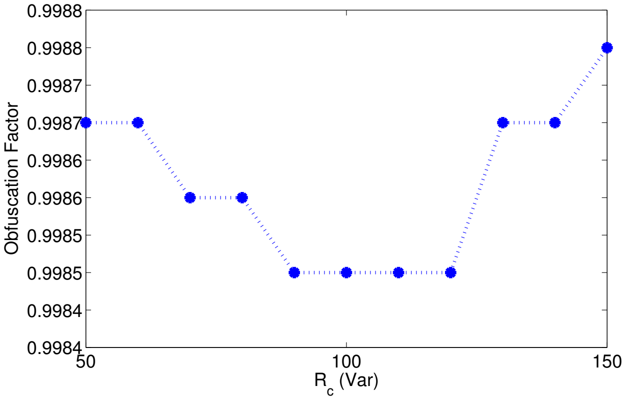

Fig. 18: R_c's effects on the obfuscation factor

an example, Fig. 15 shows a 24 hours-long load profile. The left figure is the original reactive power, and the right figure is the obfuscated reactive power. The figures show that after using RPO, the waveform appears to be more stable. Most of the time, the obfuscated reactive power stays around $100\,Var$ and no appliance usage can be inferred, which means that RPO is able to mask the reactive power changes caused by turning ON/OFF appliances.

As shown in the figure, some fluctuations still occur in the obfuscated data due to the delay of the storing/providing process of the capacitor. When the appliance changes state, it takes some time (about one second) for the capacitor to fully compensate the change and there will be a very short fluctuation before the reactive power changes back to $100\,Var$. However, these short fluctuations will not reveal detailed appliance usage information such as which appliance is turned ON/OFF. This is because the reactive power before the fluctuation is not the true reactive power consumption of the home appliance, and it is set to be $100\,Var$ by RPO. Since the starting reactive power is different, the reactive power change is not related to the power consumption of the appliance. Moreover, this fluctuation is too short (only about one second) to extract the wavelet coefficient. Thus, the attacker cannot link the fluctuation with the appliance correctly. To further validate this, we perform a wavelet transform on the obfuscated data and the coefficients are shown in Fig. 16. We can see from the figure that the coefficients of the obfuscated data do not have any pattern and can not be mapped to any appliance.

2) Vulnerability Factor and Obfuscation Factor: The vulnerability factor and the obfuscation factor with or without RPO are shown in Table IV. After applying RPO, the vulnerability factor for all three households drop significantly (more than 98%). It means that RPO can effectively reduce the amount of useful information contained in the reactive power data and makes it much harder for the attacker to make inferences. It shows that the obfuscated data and the original data are significantly different and RPO can obfuscate the original reactive power data very well.

3) Effectiveness of Mitigating Reactive Power-based Attacks: In this subsection, we run the attacks described in Section IV, based on the original reactive power data and the obfuscated reactive power data. The results are shown in Table V. Before applying RPO, the attacker can identify appliance ON/OFF events with high accuracy. After applying RPO, the correctness of identifying ON/OFF events is very low (less than 2%). For example, as shown in the Table, for Apt 1, without RPO, the detection rate for the attacker to identify Lamp "ON" can reach 95.8%, but this detection rate drops to 1.2% when RPO is applied. Thus, the results demonstrate that RPO can prevent attackers from inferring appliance usage information and hence can preserve user privacy.

4) Effects of initial R_c: The initial target load R_c is an important factor of RPO. We evaluated the vulnerability factor and obfuscation factor with different initial R_c and the results are shown in Fig. 17 and Fig. 18. As shown in Fig. 17, the vulnerability factor changes between 0.05 and 0.45 when the initial R_c changes. The best result appears around 100 to 120 Var. Note that despite the changes, the vulnerability factor is always much better compared to 2.307 (the value when RPO is not applied).

Fig. 18 shows that the obfuscation factor does not change much (ranged from 0.9985 and 0.9988) when the initial R_c changes. Its high value indicates that RPO performs well even when R_c changes.

VII. Conclusions

This paper presented a new attack against user privacy which uses only reactive power data measured by smart meters. The attack demonstrates that reactive power can be exploited by attackers to infer appliance usage information and thus poses bigger threats to user privacy. To address this problem, we proposed a novel technique called RPO to obfuscate the reactive power measured by smart meters. RPO uses capacitors to store and provide reactive power in a controlled manner, to smooth power fluctuations, so that the reactive power changes caused by the ON/OFF of appliances are hidden from the smart meter. Evaluation results show that RPO can prevent adversary from exploiting reactive power data to infer appliance usage information and hence preserve user privacy.

Acknowledgments

This work was supported in part by the National Science Foundation (NSF) under grant No. CNS-1526425, and by the Defense Threat Reduction Agency under grant HDTRA1-10-1-0085.

References

[1] U.S. Department of Energy, “The Smart Grid: An Introduction,” http://energy.gov/oe/downloads/smart-grid-introduction, 2008.

[2] B. Cook, J. Gazzano, Z. Gunay, L. Hiller, S. Mahajan, A. Taskan, and S. Vilogorac, “The Smart Meter and a Smarter Consumer: Quantifying the Benefits of Smart Meter Implementation in the United States,” *Chemistry Central Journal*, vol. 6, no. 1, 2012.

[3] X. Dong, H. Lin, R. Tan, R. K. Iyer, and Z. Kalbarczyk, “Software-Defined Networking for Smart Grid Resilience: Opportunities and Challenges,” in *ACM Workshop on Cyber-Physical System Security (CPSS)*, 2015.

[4] M. Hossain, N. Madlool, N. Rahim, J. Selvaraj, A. Pandey, and A. F. Khan, “Role of Smart Grid in Renewable Energy: An Overview,” *Renewable and Sustainable Energy Reviews*, vol. 60, 2016.

[5] Q. Li and G. Cao, “Multicast authentication in the smart grid with one-time signature,” *IEEE Transactions on Smart Grid*, vol. 2, no. 4, pp. 686–696, 2011.

[6] Y. Bachy, F. Basse, V. Nicomette, E. Alata, M. Kaâniche, J.-C. Courrège, and P. Lukjanenko, “Smart-TV Security Analysis: Practical Experiments,” in *IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, 2015.

[7] V. B. Krishna, K. Lee, G. A. Weaver, R. K. Iyer, and W. H. Sanders, “F-DETA: A Framework for Detecting Electricity Theft Attacks in Smart Grids,” in *IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, 2016.

[8] H. Lin, H. Alemzadeh, D. Chen, Z. Kalbarczyk, and R. K. Iyer, “Safety-Critical Cyber-Physical Attacks: Analysis, Detection, and Mitigation,” in *Symposium and Bootcamp on the Science of Security*, 2016.

[9] E. Fernandes, J. Jung, and A. Prakash, “Security Analysis of Emerging Smart Home Applications,” in *IEEE Symposium on Security and Privacy (S&P)*, 2016.

[10] T. A. Nguyen and M. Aiello, “Energy Intelligent Buildings Based on User Activity: A Survey,” *Energy and buildings*, vol. 56, 2013.

[11] C. Dinesh, B. W. Nettasinghe, R. I. Godaliyadda, M. P. B. Ekanayake, J. Ekanayake, and J. V. Wijayakulasooriya, “Residential Appliance Identification Based on Spectral Information of Low Frequency Smart Meter Measurements,” *IEEE Transactions on Smart Grid*, 2015.

[12] G. Eibl and D. Engel, “Influence of Data Granularity on Smart Meter Privacy,” *IEEE Transactions on Smart Grid*, vol. 6, no. 2, 2015.

[13] U. Greveler, B. Justus, and D. Loehr, “Multimedia Content Identification Through Smart Meter Power Usage Profiles,” *Computers, Privacy and Data Protection*, vol. 1, p. 10, 2012.

[14] W. Yang, N. Li, Y. Qi, W. Qardaji, S. McLaughlin, and P. McDaniel, “Minimizing Private Data Disclosures in the Smart Grid,” in *ACM Conference on Computer and Communications Security (CCS)*, 2012.

[15] D. Egarter, C. Prokop, and W. Elmenreich, “Load Hiding of Household’s Power Demand,” in *IEEE International Conference on Smart Grid Communications (SmartGridComm)*, 2014.

[16] L. Yang, X. Chen, J. Zhang, and H. V. Poor, “Cost-Effective and Privacy-Preserving Energy Management for Smart Meters,” *IEEE Transactions on Smart Grid*, vol. 6, no. 1, pp. 486–495, 2015.

[17] G. W. Hart, “Nonintrusive Appliance Load Monitoring,” *Proceedings of the IEEE*, vol. 80, no. 12, pp. 1870–1891, 1992.

[18] S. N. Patel, T. Robertson, J. A. Kientz, M. S. Reynolds, and G. D. Abowd, “At the Flick of a Switch: Detecting and Classifying Unique Electrical Events on the Residential Power Line,” in *ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*, 2007.

[19] M. Zeifman and K. Roth, “Nonintrusive Appliance Load Monitoring: Review and Outlook,” *IEEE Transactions on Consumer Electronics*, vol. 57, no. 1, 2011.

[20] W. Wang and Z. Lu, “Cyber Security in the Smart Grid: Survey and Challenges,” *Computer Networks*, vol. 57, no. 5, pp. 1344–1371, 2013.

[21] F. Diao, F. Zhang, and X. Cheng, “A Privacy-Preserving Smart Metering Scheme Using Linkable Anonymous Credential,” *IEEE Transactions on Smart Grid*, vol. 6, no. 1, pp. 461–467, 2015.

[22] Y. Gong, Y. Cai, Y. Guo, and Y. Fang, “A Privacy-Preserving Scheme for Incentive-Based Demand Response in the Smart Grid,” *IEEE Transactions on Smart Grid*, vol. 7, no. 3, pp. 1304–1313, 2016.

[23] L. Chen, R. Lu, and Z. Cao, “PDAFT: A Privacy-Preserving Data Aggregation Scheme With Fault Tolerance for Smart Grid Communications,” *Peer-to-peer networking and applications*, vol. 8, no. 6, 2015.

[24] Q. Li, G. Cao, and T. La Porta, “Efficient and privacy-aware data aggregation in mobile sensing,” *IEEE Transactions on Dependable and Secure Computing*, vol. 11, no. 2, 2014.

[25] J. Fan, Q. Li, and G. Cao, “Privacy-Aware and Trustworthy Data Aggregation in Mobile Sensing,” in *IEEE Conference on Communications and Network Security (CNS)*, 2015.

[26] L. Yang, X. Chen, J. Zhang, and H. V. Poor, “Optimal Privacy-Preserving Energy Management for Smart Meters,” in *IEEE International Conference on Computer Communications (INFOCOM)*, 2014.

[27] “GE I-210 meters,” http://www.gedigitalenergy.com/SmartMetering/catalog/i210_family.htm.

[28] “Landis Gyr Meters,” http://www.landisgyr.com/products/.

[29] R. Bracewell, “The Fourier Transform and Its Applications,” 1965.

[30] Y. Meyer, “Wavelets and Operators,” 1995.

[31] O. Parson, S. Ghosh, M. Weal, and A. Rogers, “An Unsupervised Training Method for Non-Intrusive Appliance Load Monitoring,” *Artificial Intelligence*, vol. 217, pp. 1–19, 2014.

[32] A. V. Oppenheim, R. W. Schafer, and J. R. Buck, *Discrete-Time Signal Processing*. Prentice-hall Englewood Cliffs, 1989.

[33] “ISL78268,” https://www.intersil.com/products/ISL78268.

[34] S. McLaughlin, P. McDaniel, and W. Aiello, “Protecting Consumer Privacy from Electric Load Monitoring,” in *ACM Conference on Computer and Communications Security (CCS)*, 2011.

[35] “OrCAD: Cadence PCB Solutions,” http://www.orcad.com/.