A Privacy-Preserving Scheme for Incentive-Based Demand Response in the Smart Grid

Yanmin Gong, *Student Member, IEEE,* Ying Cai, *Member, IEEE,* Yuanxiong Guo, *Member, IEEE,*
and Yuguang Fang, *Fellow, IEEE*

***Abstract*—The advanced metering infrastructure (AMI) in the smart grid provides real-time information to both grid operators and customers, exploiting the full potential of demand response. However, it introduces new privacy threats to customers. Prior works have proposed privacy-preserving methods in the AMI such as temporal or spatial aggregation. A main assumption in these works is that fine-grained data do not need to be attributable to individuals. However, this assumption does not hold in incentive-based demand response (IDR) programs where fine-grained metering data are required to analyze individual demand curtailments and hence need to be attributable. In this paper, we propose a privacy-preserving scheme for IDR programs in the smart grid, which enables the demand response provider (DRP) to compute individual demand curtailments and demand response rewards while preserving customer privacy. Moreover, a customer can reveal his/her identity and prove ownership of his/her power usage profile in certain situations such as legal disputes. We achieve both privacy and efficiency in our scheme through a combination of several cryptographic primitives such as identity-committable signatures (ICS) and partially blind signatures. As far as we know, we are the first to identify and address privacy issues for IDR programs in the smart grid.**

I. Introduction

The smart grid is a modernized power grid that uses information and communication technologies to improve the efficiency, reliability, economics, and sustainability of the generation, transmission, distribution, and consumption of electricity. In the smart grid, a full measurement and collection system called the advanced metering infrastructure (AMI) replaces traditional electromechanical meters. The AMI collects fine-grained, time-based information and transmits them to various parties through a communication network, enabling the integration of demand-side resources into the wholesale market and hence the demand response (DR).

According to the U.S. Department of Energy, DR refers to "changes in electric use by demand-side resources from their normal consumption patterns in response to the varying electricity price, or to incentive payments designed to reduce electricity use when wholesale market prices are high or when system reliability is jeopardized" [1]. In the power grid, generation and consumption should be balanced instantaneously. The load-following strategy, where a power plant adjusts its power supply to match the fluctuating demand, has been dominant in the traditional power grid operations. However, this strategy incurs a high cost in terms of environment, grid reliability and operational efficiency. On the contrary, the smart grid places great emphasis on the DR strategy where consumers shape their power demand to match the supply [2]. DR supports high penetration of renewable energy generation by shaping the demand to match the intermittent and unpredictable power output of renewable generation, and it also brings other benefits such as peak shaving, reliability enhancement, and generation cost reduction.

Generally speaking, there are two types of DR programs: price-based demand response (PDR) programs that motivate customers to change their consumption patterns according to time-varying electricity prices and incentive-based demand response (IDR) programs that reward participating customers for reducing their electricity usage in response to DR requests. Although more utilities offer some types of PDR programs to customers than IDR programs, PDR accounts for just a small part of the total DR resource base [3]. Since IDR programs can be tailored to specific operational goals such as localized load reduction during transmission congestion, they diversify the ways in which demand-side management contributes to reliable and efficient grid operations. In IDR programs, the time interval of measurements varies from hours to seconds based on different trigger conditions [4], which poses a serious threat to customer privacy [5], [6]. It has been shown that power usage profiles at a granularity of 15 minutes may reveal whether a child is left alone at home and at a finer granularity may reveal the daily routines of customers [7]. Despite its importance, the privacy issues in IDR programs have never been addressed before. The unique challenge of IDR programs lies in the fact that the meter measurements should be both attributable and fine-grained, excluding some popular privacy-preserving approaches that address privacy issues in PDR programs. In IDR programs, there is a new party called the demand response provider (DRP), who aggregates demand-side resources of customers and rewards customers based on their demand curtailments in DR events. The DRP can either be the electric utility company or a third party, and it collects fine-grained metering measurements in order to calculate the customer baseline (CBL) and hence the demand curtailments.

In this paper, we aim at preserving customer privacy for IDR programs in the smart grid. We propose a scheme that enables

Y. Gong and Y. Fang are with the Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL 32611 USA (e-mail: ymgong@ufl.edu; fang@ece.ufl.edu).

Y. Cai is with Department of Computer Science and Technology, Beijing Information Science and Technology University, Beijing 100101, China and with State Key Laboratory of Information Security, Institute of Information Engineering, Chinese Academy of Sciences, Beijing 100093, China. (e-mail: ycai@bistu.edu.cn).

Y. Guo is with the School of Electrical and Computer Engineering, Oklahoma State University, Stillwater, OK 74078 USA (e-mail: richard.guo@okstate.edu).

the DRP to profile, reward, and provide feedback to customers in IDR programs without violating customer privacy. The DRP is able to analyze fine-grained metering data to calculate CBLs, schedule demand curtailments, and correctly reward customers, but it cannot link the real identity of a customer to the fine-grained metering data. Our scheme is constructed by cryptographic primitives. Individual metering data are signed with a special technique such that the authenticity can be verified without revealing the real identity of the signer. When customers want to inquire their metering data or claim their DR rewards, they prove their eligibility to the DRP but reveal no additional information about themselves. With these techniques combined, the anonymity of customers is guaranteed throughout the IDR processes. *As far as we know, we are the first to address the privacy issues in IDR programs.*

The rest of the paper is organized as follows. Related work is provided in Section II. Section III presents the cryptographic primitives used in our scheme. We provide some background on IDR programs and describe the components, system flow, and design goals of our scheme in Section IV. Section V elaborates on the proposed scheme, where we design privacy-preserving protocols for different processes in IDR programs. Practical considerations and useful extensions are presented in Section VI. Section VIII and Section IX analyze the security and the efficiency of the proposed scheme, respectively. Finally, Section X concludes the paper.

II. Related Work

While instrumental to the implementation of DR, fine-grained metering data collected by the AMI can be used to determine occupant activities, raising serious privacy concerns [8]. Research studies on non-intrusive load monitoring (NILM) [6], [9], [10] have shown the possibility of deducing appliance usage patterns from fine-grained metering data. The appliance usage patterns can be further analyzed to learn the health status, daily routines or unusual behaviors such as "you slept late at night" and "your child is left alone at home" [7]. Hence, a growing number of research activities have been carried out to address privacy issues in the AMI.

The approaches to addressing privacy issues in the AMI can be divided into three categories. The first category proposes to aggregate individual metering data before sending them out to utility companies since most benefits of the smart grid can be achieved with the aggregate data. Aggregation can be either performed at a central point [11]–[13] or distributed in the network [14]. The second category uses cryptographic tools to hide sensitive information, mainly adopted for private billing purposes in PDR programs [7], [12], [15]. These works intend to calculate bills at the customer side and ask customers to prove the correctness of their bills to utilities. The third category uses anonymity to protect user privacy [16]. The aforementioned approaches share a common assumption: *metering data for operational purposes do not need to be attributable to a specific customer, and metering data for billing purposes do not require to be in fine granularity.* There is also a line of research that preserve privacy by adding noise to the power usage profiles using local rechargeable batteries [9], [17]. This approach is complementary to previous privacy-preserving approaches. However, challenges such as battery maintenance costs and additional capital costs still exist in the large-scale integration of local storage devices. In this paper, we address privacy issues in IDR programs, for which the previous assumption no longer holds. In IDR programs, fine-grained metering data are required when the DRP schedules demand curtailments, calculates CBLs, and allocates DR rewards for individual customers. Hence both fine granularity and attributability are required for IDR programs. In this case, aforementioned privacy protection mechanisms such as aggregation are not applicable, and a new approach is needed.

III. Cryptographic Primitives

This section gives an introduction to the cryptographic primitives used as the building blocks in our scheme.

A. Identity-Committable Signature (ICS)

The identity-based signature scheme [18] avoids the use of certificates in conventional public key infrastructure by deriving the public key of a signer from his public identity information such as email address and telephone number. The scheme designed in [19] makes use of bilinear pairings on elliptic curves, a popular technique in identity-based public key cryptography. Let $\mathbb{G}$ be an additive group with generator P, and $\mathbb{G}_T$ be a multiplicative group. A mapping $\hat{e} : \mathbb{G} \times \mathbb{G} = \mathbb{G}_T$ is called a bilinear pairing if it satisfies the following:

- Bilinearity: $\hat{e}(aP, bP) = \hat{e}(P, Q)^{ab}$ for all $a, b \in Z_p$ and $P \in \mathbb{G}$.
- Non-degeneracy: If P is a generator of G, then $\hat{e}(P, P) \neq 1$.
- Computability: There exists an efficient algorithm to compute $\hat{e} = (P, Q)$ for all $P, Q \in \mathbb{G}$.

Chu and Tzeng [20] construct an ICS scheme which allows a signer to sign a message on behalf of an organization or a group. The scheme is setup as follows. The private key generator (PKG) chooses a master secret key $(x, y) : x, y \in_R \mathbb{Z}_p$ and three hash functions $H_1 : \{0,1\}^* \rightarrow \mathbb{G}$, $H_2 : \{0,1\}^* \times \mathbb{G} \rightarrow \mathbb{Z}_p$, and $H_2' : \{0,1\}^* \times \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{Z}_p$. Then it computes $P_X = xP$, $P_Y = yP$ and publishes $(\mathbb{G}, \mathbb{G}_T, \hat{e}, P, P_X, P_Y, H_1, H_2, H_2')$ as the public parameters. For identity I, the DRP calculates $Q_I = H_1(I)$, $Q_I' = xQ_I$, and $S_I = xyQ_I$. The public and private key pairs for the user are Q_I and (Q_I', S_I), respectively. To generate an ICS on message m, the signer randomly selects a value $r \in Z_p$, computes $h = H_2(m, U)$, and generates $U_I = rQ_I'$, $V_I = (r+h)S_I$. The signer then chooses a secret $\mu \in Z_p^* \setminus \{1\}$ and computes $Q = \mu Q_I, Q' = \mu Q_I', U = \mu U_I, V = \mu V_I$. The ICS on message m is $\delta_{IC} = (Q, Q', U, V)$. To verify the signature, the verifier calculates $h = H_2'(m, Q, U)$ and accepts the signature if and only if $\hat{e}(Q, P_X) = \hat{e}(Q', P)$ and $\hat{e}(U, P_Y) = \hat{e}(V, P)\hat{e}(Q', -P_Y)^h$ hold.

B. Zero-Knowledge Proof(ZKP)

The notion of ZKP is introduced by Goldwasser et al. [21] in which the prover takes interactive input from the verifier and

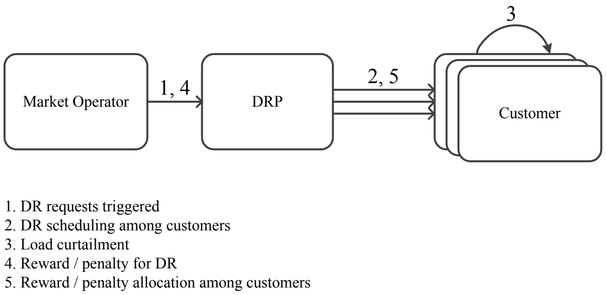
3
Market Operator
1, 4
DRP
2, 5
Customer
1. DR requests triggered
2. DR scheduling among customers
3. Load curtailment
4. Reward / penalty for DR
5. Reward / penalty allocation among customers

Fig. 1. Electricity Market for Incentive-Based Demand Response (IDR)

responds based on this input. With Fiat-Shamir heuristic [22], the ZKP can be transformed into the non-interactive form where interaction is not needed between the verifier and the prover. The non-interactive form has been proved to be secure under the random oracle model. We follow the notions introduced by Camenisch and Stadler [23] to describe the ZKP protocols, where $PK\{\cdot\}$ denotes the zero-knowledge proof of a statement. For instance, $PK\{\alpha : C = g^{\alpha}\}$ is used to prove the knowledge of the discrete logarithm of C with base g, and $PK\{(a, b) : C = g_0^a \wedge C' = g_1^b\}$ is used to prove the knowledge of both a and b which satisfy the expression on the right side of the colon. With this notion, we can describe a ZKP without involving details.

C. Partially Blind Signature

Partially Blind Signature. Commitment schemes enable one to commit a chosen value without revealing it. A well known commitment scheme is the Pedersen Commitment [24]. Let $\mathbb{G}$ be a group of prime order p, and g and h be generators of $\mathbb{G}$. To commit a value $x \in \mathbb{Z}_p$, the committer randomly chooses $r \in \mathbb{Z}_p$, computes $C = g^x h^r$, and outputs C as the commitment. To reveal x, the committer discloses x, r. The verifier can verify if $C = g^x h^r$. Multiple values can be committed in a single commitment. For example, the commitment for x_1, x_2 is $C = g_1^{x_1} g_2^{x_2} h^r$, where g_1, g_2 are generators of $\mathbb{G}$. We denote the Pedersen Commitment on message x as $\mathrm{CM}(x)$.

An application of commitment schemes is the BBS+ signature designed in [25] and [26]. The construction of BBS+ signature is partially blinded: the signer can sign messages in a commitment without knowing their values. Let $\mathbb{G}$, $\mathbb{G}_T$ be two cyclic groups of prime order p, and $\hat{e} : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$ be a bilinear pairing function. Let $g, g_0, g_1, g_2 \in \mathbb{G}$ be generators of $\mathbb{G}$, which are public parameters. The signer randomly chooses $\gamma \in \mathbb{Z}_p$ as his secret key and computes $\omega = g^{\gamma}$ as his public key. To sign messages m_1, m_2, the signer randomly chooses $c, z \in Z_p$, computes $A = (g g_0^z g_1^{m_1} g_2^{m_2})^{1/(c+\gamma)}$, and outputs (A, c, z) as the signature. One can verify a BBS+ signature by checking if $\hat{e}(A, \omega g^c) = \hat{e}(g g_0^z g_1^{m_1} g_2^{m_2}, g)$ holds.

IV. System Model

We describe the system model of the proposed privacy-preserving scheme in this section.

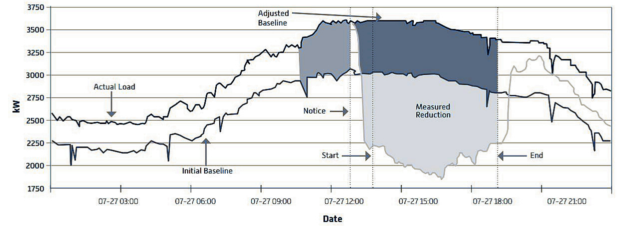

Fig. 2. Example Baseline and Performance Measurement for Demand Response Asset [4]

A. Background

As shown in Fig. 1, the electricity market for IDR involves three entities: the market operator, the DRP, and the customers. The market operator manages the electricity market and triggers DR events based on the status of the power grid. When a DR event is triggered, the DRP schedules load curtailment among customers and aggregates these demand side resources. Participating customers reduce their load during a DR event as scheduled. The market operator pays the DRP for its aggregate curtailment, and the DRP then allocates the reward among participating customers. Since customers have unequal contributions to the aggregate load curtailment, they should be rewarded based on their contributions so that active customers are encouraged.

Individual customer contribution is calculated as the difference between his real-time power consumption and his baseline consumption level (CBL), which represents the "behave-as-usual" usage pattern of a customer. Calculation of CBL is among the most important factors in an IDR program because it should neither reward nor penalize a customer for his natural load variances. Fig. 2 (depicted by [4]) gives an example of CBL, where the initial baseline is adjusted according to the actual load on that day so that the effort of demand reduction of the customer can be fairly estimated. In order to mimic the dynamic shape of the customer load, CBL calculation algorithms of the DRP take as input an extensive data set including both fine-grained historical meter measurements and peripheral data (e.g., weather and time of the day) [4], [27]. However, these fine-grained metering data as required by the DRP in the CBL and curtailment calculation raise serious customer privacy concerns.

B. Components

To address these privacy concerns, we propose a scheme which enables the DRP to perform all the required operations without linking customer identity and fine-grained metering data. The scheme involves four components, i.e., *smart meters*, *customer devices*, *a proxy*, and *a DRP*.

Smart Meters. The utility company installs smart meters at customer premises, one for each customer. Smart meters are assumed to be tamper-resistant and able to perform elementary cryptographic operations, but they cannot store long-term metering data or perform CBL calculation due to limited storage and computation capabilities.

The Proxy. The proxy plays the role of an anonymizer which hides the static IP address of smart meters. It can be either the gateway or a trusted third party. The proxy is semi-trusted, meaning curious but not malicious, in the sense that it may attempt to learn the customer privacy, but it will faithfully relay the metering data and hide the smart meter IP address from the DRP. From now on, when we refer to “anonymous channel”, we mean an anonymous communication channel established by the proxy.

The DRP. The DRP schedules DR events among customers, records customer performance in DR events, and calculates their corresponding rewards. The DRP is semi-trusted, meaning that it may attempt to learn the customer privacy, but it will faithfully follow protocol specifications.

Customer Devices. Customers query the DRP to learn their own metering data and claim DR rewards through customer devices (e.g., personal computers or smartphones). Customers are assumed to be curious and potentially malicious. They may impersonate other customers or collude with the DRP to learn power usage profiles of other customers, or cheat to gain undeserved rewards.

There may be external attackers who launch denial-of-service attack or man-in-the-middle attack, or eavesdrop. However, addressing these attacks is beyond the scope of this paper.

C. System Flow

The scheme includes the following processes. In the **registration process**, the DRP creates two accounts for a customer, one associated with his real identity and the other associated with his pseudonym. The real identity can be any information that uniquely identifies the customer, such as the account number or telephone number. Since a customer can only enroll in a single IDR program at a time, the DRP needs to make sure that a customer does not register multiple pseudonyms. This is achieved with the anonymous ticket: the customer obtains a ticket when he registers the real identity and presents it to the DRP when he anonymously registers the pseudonym. In the **metering process**, the smart meter collects metering data, constructs signatures on them, and sends them together with its pseudonym to the DRP through the anonymous channel. The DRP stores the data by pseudonym in the database and analyzes the data for operational and settlement purposes. The ICS signature ensures the authenticity of the metering data, while the ZKP ensures that adversaries cannot change the pseudonym in the message. The ZKP also enables customers to prove ownership of their pseudonyms when making personal inquiries for CBLs or metering data in the **querying process**. Customers claim rewards with a partially blinded signature (BBS+) which hides the real identity but ensures the integrity in the **settlement process**. The pseudo accounts of customers are revoked in the **revocation process** when customers leave the DRP programs.

D. Design Goals

We intend to design a scheme that guarantees privacy, integrity, and availability.

Privacy. Customers need to register their real identities for security reasons. However, they want to remain anonymous when querying their metering data or claiming their rewards. We guarantee this by allowing no other party except the customer himself to know his fine-grained power usage profile.

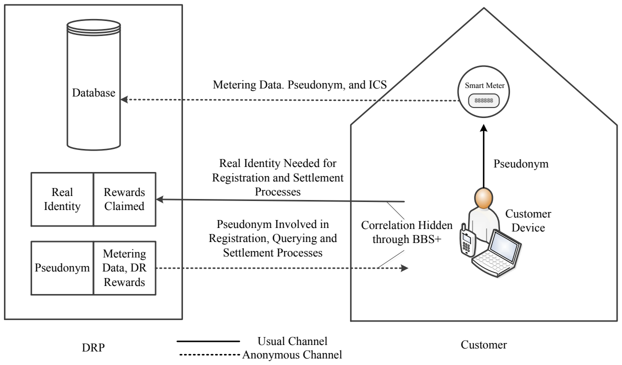
Database
Metering Data. Pseudonym, and ICS
Smart Meter
888888
Pseudonym
Real Identity Needed for Registration and Settlement Processes
Real Identity
Rewards Claimed
Pseudonym Involved in Registration, Querying and Settlement Processes
Correlation Hidden through BBS+
Customer Device
Pseudonym
Metering Data, DR Rewards
DRP
Usual Channel
Anonymous Channel
Customer

Fig. 3. System Diagram

Integrity. We also need to ensure the integrity of the scheme. Misbehaviors such as falsifying the metering data or double spending should be detected immediately.

Availability. Guaranteeing the availability of IDR programs means that all the features required by IDR programs are fulfilled and the efficiency is guaranteed. Specifically, the DRP can gather information to profile, reward and provide feedback to customers while customers can learn their DR performance and claim their rewards. Moreover, since the metering data should be transmitted and processed with low latency, the metering process should have low computation and communication overhead.

V. Basic Protocol Design

We describe the basic protocols in this section. The system diagram is given in Fig. 3. We use the ICS scheme to authenticate metering data in the metering process, which enables the DRP to authenticate the data without knowing the real identity of the signer. The BBS+ scheme is used to hide the relationship between the identifiable account and the pseudo account. For ease of presentation, we use $PK\{\cdot\}$ to denote the ZKP of a statement. The detailed construction of the ZKPs will be given in Sec. VII.

A. Setup Process

The DRP plays the role of the private key generator (PKG) and sets up the master key and public parameters for the ICS scheme and the BBS+ scheme.

ICS Scheme. Let $\mathbb{G}_1$ be an additive group with generator P_1, and $\mathbb{G}_{1T}$ be a multiplicative group. The private key generator (PKG) chooses a master secret key (x, y) : $x, y \in_R \mathbb{Z}_p$ and three hash functions $H_1 : \{0,1\}^* \to \mathbb{G}_1$, $H_2 : \{0,1\}^* \times \mathbb{G}_1 \to \mathbb{Z}_p$, and $H_2' : \{0,1\}^* \times \mathbb{G}_1 \times \mathbb{G}_1 \to \mathbb{Z}_p$. Then it computes $P_X = xP_1$, $P_Y = yP_1$ and publishes $(\mathbb{G}_1, \mathbb{G}_{1T}, \hat{e}, P_1, P_X, P_Y, H_1, H_2, H_2')$ as the public parameters for the ICS scheme.

BBS+ Scheme. Let $\mathbb{G}_2$ be an additive group with generator P_2. To generate the signing key of the DRP for BBS+ signature construction, the PKG selects $g, g_0, g_1, g_2, g_3, g_4 \in \mathbb{G}_2$ and $\gamma \in \mathbb{Z}_p$. It further computes $\omega = g^\gamma$ as the private signing key. Finally, the PKG publishes the public parameters of the system as $\mu = (\mathbb{G}_2, P_2, g, g_0, g_1, g_2, g_3, g_4, \omega)$.

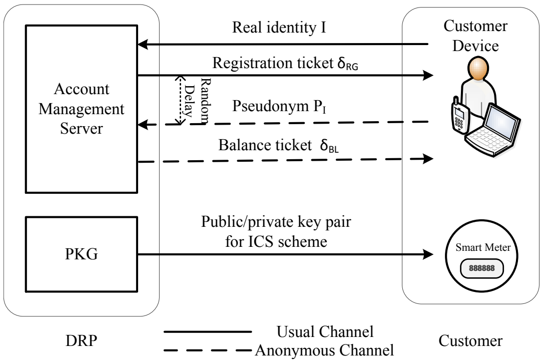

Fig. 4. Registration Process

B. Registration Process

Fig. 4 describes the registration process. The customer reveals his identity I to the DRP for registration. After verifying his eligibility, the DRP computes and sends the public/private key pair (Q_I', S_I) to the smart meter of the customer. Moreover, the DRP creates and returns δ_s^{RG}, a BBS+ signature on s tagged with "RG", to the customer. As described in Sec. III, the construction process of the signature involves several interactions between the DRP and the customer. The customer first commits a random secret s in the commitment $\mathrm{CM}(s) = g_0^{z'} g_3^s$ and sends the commitment $\mathrm{CM}(s)$ to the DRP. He also constructs and sends a ZKP PK_1 to show the correctness of the commitment, where

$$PK_1\{(z', s) : C = g_0^{z'} g_3^s\}. \tag{1}$$

The DRP checks the correctness of PK_1 and aborts the ongoing process if it is incorrect. Otherwise, the DRP picks $z'', c \in_R \mathbb{Z}_p$, computes $A = (g g_0^{z''} g_1^I)^{1/(c+\gamma)}$, and returns A, z'', c to the customer. Next, the customer computes $z = z' + z''$ and checks if $\hat{e}(A, \omega g^c) = \hat{e}(g g_0^z g_1^I g_3^s, g)$ holds to ensure its correctness. If the verification fails, he refuses to accept the ticket and restarts the whole process. Otherwise, he parses and stores the signature $\delta_s^{RG} = (A, c, z, s)$ as the registration ticket for his pseudonym. The value of s remains hidden during the process.

The customer also registers a pseudonym. To this end, the customer selects a random number λ_I as his secret and computes his pseudonym P_I as $P_I = g_4^{\lambda_I}$. After a random delay, the customer sends P_I and (δ_s^{RG}, s) to the DRP through an anonymous channel. He proves to the DRP that (1) δ_s^{RG} is a valid signature on s, and (2) $P_I = g_4^{\lambda_I}$ with a ZKP PK_2:

$$\begin{aligned} &PK_2\{(\lambda_I, A, c, z, I, z') : \\ &\quad P_I = g_4^{\lambda_I} \wedge \hat{e}(A, \omega g^c) = \hat{e}(g g_0^z g_1^I g_3^s, g)\}. \end{aligned} \tag{2}$$

If the DRP verifies the validity of the ZKP, it establishes a pseudo account associated with P_I. To initiate the balance in the pseudo account, the customer randomly selects a new value of s, sends commitment $\mathrm{CM}(I, B, s)$ to the DRP via the anonymous channel, provides a ZKP PK_3, and obtains δ_s^{BL}, a BBS+ signature on (I, B, s). Here, B denotes the balance and is initialized to 0, and the ZKP PK_3 is defined as follows:

$$PK_3\{(I, s, z') : C = g_0^{z'} g_1^I g_3^s\}. \tag{3}$$

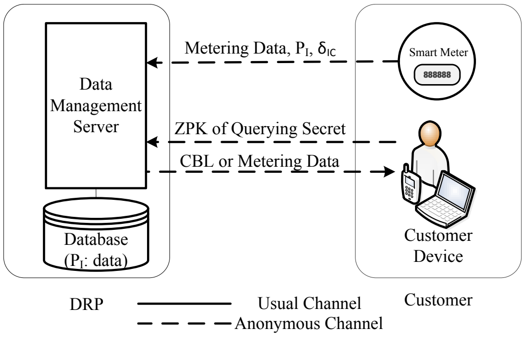

Fig. 5. Metering and Querying Process

The customer stores (δ_s^{BL}, I, B, s) as the balance ticket. Note that s is updated every time, and hence a customer cannot use the same ticket twice.

After the registration of pseudonym P_I, the customer inputs the pseudonym into the smart meter. The smart meter stores the pseudonym locally and only uses the pseudonym for metering purposes.

The whole process can be described with the following protocols, where CD, SM, and anon. are used to denote the customer device, the smart meter, and the anonymous channel, respectively:

1) CD → DRP : I;
2) DRP → CD : δ_s^{RG};
3) DRP → SM : $Q_I, (Q_I', S_I)$;
4) CD $\xrightarrow{anon.}$: DRP : δ_s^{RG}, P_I;
5) DRP $\xrightarrow{anon.}$: δ_s^{BL};
6) CD → SM : P_I.

C. Metering and Querying Processes

Fig. 5 describes the metering and querying processes. At each reporting cycle t, the smart meter collects metering data m_t and generates an ICS signature δ_{IC} on the metering data. The construction process has been described in III. It then attaches the pseudonym of the customer to the message and sends the entire message $(m_t, t, P_I, \delta_{IC})$ to the DRP through the anonymous channel. Upon receiving the message, the DRP checks the validity of δ_{IC}. If δ_{IC} passes the verification, the DRP stores (m_t, δ_{IC}) as the metering record at time t for the pseudo account P_I. Otherwise, the DRP discards the message. The metering records associated with P_I can be used to calculate individual CBL and allocate DR rewards.

Specifically, the CBL associated with P_I is calculated as $b_t = f(\{m_\tau\}_{\tau \in \Gamma})$, where m_τ represents the historical metering data, Γ is a baseline window over which demand data are collected, and $f(\cdot)$ is a mapping from historical measurements to the CBL.

We describe the metering process with the following protocol:

$$\text{SM} \xrightarrow{anon.} \text{DRP} : m_t, t, \delta^{IC}, P_I. \tag{4}$$

In the querying process, the customer proves his knowledge about the secret key λ_I of pseudo account P_I with a ZKP PK_4:

$$PK_4\{\lambda_I : P_I = g_4^{\lambda_I}\}. \tag{5}$$

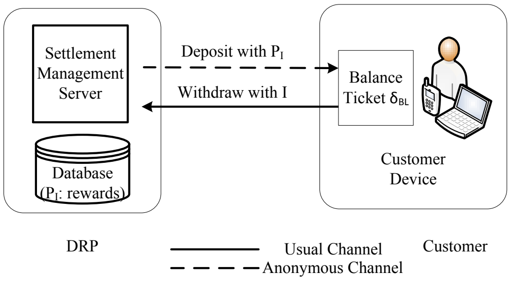

Fig. 6. Settlement Process

The customer sends a querying request together with PK_4 to the DRP via the anonymous channel. If PK_4 is correctly constructed, the DRP locates the requested data in the database and sends it back over the established anonymous channel. Otherwise the request is rejected and the corresponding request is ignored. We describe the process with the following protocols:

1) CD $\xrightarrow{anon.}$ DRP : PK_4, P_I;
2) DRP $\xrightarrow{anon.}$ CD : CBL, m_t.

D. Settlement Process

The DR rewards are allocated to pseudo accounts by the DRP based on individual curtailments. Customers may claim their rewards in two steps, as described in Fig. 6. First, a customer adds the reward into his balance ticket through the anonymous channel. Suppose his old balance ticket is $(\delta_{\tilde{s}}^{BL}, I, \tilde{B}, \tilde{s})$. To transfer reward d from the pseudo account, the customer first checks if the reward in his pseudo account is larger than d. If yes, he selects a random secret s and sends commitment CM$(s, I, B, \tilde{B})$ to the DRP, together with the following ZKP PK_5:

$$PK_5\{(\lambda_I, \tilde{A}, \tilde{c}, I, \tilde{z}, \tilde{B}, \tilde{s}, z') : P_I = g_4^{\lambda_I} \wedge B - d > 0$$
$$\wedge\, C = g_0^{z'} g_1^I g_2^{\tilde{B}} g_3^s \wedge \hat{e}(\tilde{A}, \omega g^{\tilde{c}}) = \hat{e}(g g_0^{\tilde{z}} g_1^I g_3^{\tilde{s}}, g)\},$$

which shows that his pseudonym is P_I, the new balance is positive, and the balance ticket is correctly formed. Now the DRP verifies if both $\tilde{s}$ is never shown before and PK_5 is true. If yes, it replies with a new BBS+ signature δ_s^{BL} on the tuple (I, B, s). The customer stores (δ_s^{BL}, I, B, s) as the new balance ticket. The process is described as follows:

1) CD $\xrightarrow{anon.}$ DRP : $\delta_{\tilde{s}}^{BL}, PK_5, s$;
2) DRP $\xrightarrow{anon.}$ CD : δ_s^{BL}.

Second, the customer redeems reward from the balance ticket with his real identity. The customer selects a new s and sends the balance ticket $\delta_{\tilde{s}}^{BL}$, the withdrawal amount d, and a ZKP to the DRP, which is constructed in the same way as PK_5. The DRP then verifies the validity of the ZKP and checks if $\tilde{s}$ is never used before. If both are true, it returns a new BBS+ signature δ_s^{BL} on (I, B, s), and the customer stores (δ_s^{BL}, I, B, s) as the new balance ticket. The process of the second step is described as follows:

1) CD → DRP : $\delta_{\tilde{s}}^{BL}, PK_5, s$;
2) DRP → DRP : δ_s^{BL}.

E. Revocation Process

When the customer quits from an IDR program, the DRP needs to ensure that both the identifiable and the pseudo accounts of the customer are closed. This is guaranteed through a revocation ticket. When the customer revokes the pseudo account through the anonymous channel, he obtains a revocation ticket δ_s^{RV} from the DRP. The revocation ticket contains a BBS+ signature on (I, s) with s being the random secret selected by the customer. After a random period, the customer presents his real identity, the revocation ticket, and a ZKP PK_6 together to the DRP, where

$$PK_6\{(A, c, z, I, z') : \hat{e}(A, \omega g^c) = \hat{e}(g g_0^z g_1^I g_3^s, g)\}.$$

This ticket proves the revocation of the pseudo account associated with customer I. Then the DRP can continue to complete the rest of the revocation process. The whole process is described as follows:

1) DRP $\xrightarrow{anon.}$ CD : δ_s^{RV};
2) CD → DRP : $\delta_{\tilde{s}}^{RV}, PK_6$.

VI. Practical Considerations and Extensions

In this section, we discuss some practical issues and provide useful extensions to solve them.

A. Cloaking Mechanism

In the metering process, all the metering records of a customer are associated with the same pseudonym, which enables the DRP to link the metering data and perform basic operations. In theory, the DRP only knows the pseudonym of the power usage profiles, and thus the real identity of the customer is hidden. In practice, however, the DRP may still infer the real identity of the customer by mining the relationships between rewards and withdrawals. For example, if a customer withdraws all the available rewards in his account every settlement cycle, the withdrawals will equal the rewards; the DRP can then use the rewards as a quasi-identifier to find the real identity associated with the pseudo account. To avoid such a linkage, customers can use cloaking mechanisms when they withdraw from the balance tickets.

In general, the cloaking rules hide the relationship between withdrawals and rewards by reducing the withdrawal amounts and frequency. Ideally, if a customer withdraws once per year and leaves some balance unredeemed, the DRP can only learn an estimate of his total reward through the year. This information does not reveal the relationship between the real identity and the pseudonym since it applies to many customers. However, customers usually want to use rewards whenever they are available, and redeeming rewards motivates them to be more active in future DR events. Hence, we need to balance privacy and timeliness.

In the following, we propose two cloaking mechanisms, i.e., floor function withdrawal (FFW) mechanism and partition and random selection (PRS) mechanism. Without loss of generality, we assume that withdrawal decisions are made once per settlement cycle.

The FFW divides the range of rewards into non-overlapping intervals. Each customer falls into one interval based on their remaining balance in the balance ticket. At the end of a settlement cycle, a customer withdraws the floor value of his interval. In this way, customers in the same interval are indistinguishable because they withdraw the same amount of rewards. The achieved anonymity is determined by interval size: intervals of larger size may include more customers, and therefore provide stronger anonymity guarantee. Intervals do not need to be of the same size. For intervals which contain a dense population, the size can be chosen smaller, while for other intervals, the size should be larger.

The PRS defines a set of cells, say $\{5, 10, 20, 40\}$. Customers first partition the rewards into cells. Then they select each cell with a probability p and withdraw an amount equal to the sum of the selected cells. For example, if the reward is 45, the customer may divide it into $\{5, 10, 10, 20\}$, choose a subset of it with selection probability 0.8, and finally select cells $\{5, 10, 10\}$. The withdrawal is the sum of these cells, which is 25.

B. Pseudonym Update

Cloaking schemes can reduce information leaked to the DRP. However, in the long run, the DRP can still gain enough information for de-anonymization. Suppose that Alice receives rewards $R_1, R_2, \dots, R_N$ and withdraws $W_1, W_2, \dots, W_N$ in the first N settlement cycles. The DRP learns that

$$\sum_{k=1}^{n} W_k \le \sum_{k=1}^{n} R_k, n = 1, 2, \dots, N, \tag{6}$$

where W_k is the k-th withdrawal, and R_k is the k-th reward of Alice. If the withdrawals of customer Bob also satisfy (6), that is,

$$\sum_{k=1}^{n} W_k \le \sum_{k=1}^{n} R'_k, n = 1, 2, \dots, N, \tag{7}$$

where R'_k is the k-th reward of Bob, then Alice and Bob are indistinguishable from the DRP side. However, as N becomes large, it becomes harder to find a "Bob" who is indistinguishable from Alice. With cloaking mechanisms customers can slow down this process but not stop it. Hence, Alice needs to update the pseudonym after a few settlement cycles. Updating the pseudonym includes revocation of current pseudonym (Sec. V-E) and registration of the new one (Sec. V-B). To avoid linkage of the two pseudonyms, after revocating the old one, the customer waits a certain period before registering the new pseudonym.

C. Re-identification

There are scenarios where the customer needs to provide his power usage profiles to others. For example, a customer may use his power usage profile to justify in a legal dispute. This is especially important when the DRP is not a third party, but the utility company itself. In this case, the DRP should enable the customer to prove ownership of his profile. In other words, the customer should be able to prove to the DRP or other parties that the power usage profile is linked to his real identity. This feature can be provided through the ICS scheme.

To prove his ownership of a metering record, the customer presents the secret λ_I together with the metering data, corresponding ICS signatures, and the real identity I to the verifier. The verifier parses the ICS of the metering data in the record as $\delta_{IC} = (Q, Q', U, V)$, computes the public key for the customer as $Q_I = H_1(I)$, and checks if $Q_I = \lambda_I^{-1} Q$ holds. If the result is yes, then the verifier is convinced that the signed metering data is generated by the customer with identity I. Hence, the customer with identity I is re-identified to be the owner of the power usage profile.

Since the DRP already knows the linkage between the pseudonym P_I and the metering data, it can now readily link the real identity to the pseudonym. If a customer wants to keep his future power usage profile hidden after the re-identification process, he needs to update his pseudonym following the update protocol in Section VI-B.

VII. Detailed Constructions of $PK_1 - PK_6$

In this section, we show the detailed constructions of $PK_1 - PK_6$. To expedite the process of proof generation, we add the following system parameters : $h, h_1, h_2 \in \mathbb{G}$, $\pi = h^\sigma$, $\zeta_i = h^{1/\sigma+i}, i = 1, 2, \dots, M$, where σ is a secret that should not be leaked, and M is the maximum amount of an account. We also add $\hat{E} = \hat{e}(g, g)$, $\hat{E}_i = \hat{e}(g_i, g), i = o, \dots, 3$, $\hat{H} = \hat{e}(h, h)$, $\hat{H}_0 = \hat{e}(h_1, \omega)$, $\hat{H}_1 = \hat{e}(h_1, g)$, $\hat{H}_2 = \hat{e}(h_1, h)$, $\hat{H}_3 = \hat{e}(h_1, \pi)$ for better efficiency. All these parameters are included in the public parameters of the system. Following the Fiat-Shamir transformation, we model $\mathcal{H}$ as a random oracle.

A. $PK_1\{(z', s) : C = g_0^{z'} g_3^s\}$

First, the DRP sends a challenge F to the customer. The customer randomly picks $\rho_{z'}, \rho_s \in \mathbb{Z}_p$, and computes $T = g_0^{\rho_{z'}} g_3^{\rho_s}$. Then the customer sends T, F to the random oracle, and obtains the output $m = \mathcal{H}(T, F)$. Based on this output, the customer computes $k_{z'} = \rho_{z'} - mz'$, $k_s = \rho_s - ms'$, and sends m, $k_{z'}$, k_s to the DRP. Finally, the DRP computes $T' = C^m g_0^{k_{z'}} g_3^{k_s}$ and accepts the proof if and only if the equation $m = \mathcal{H}(T', F)$ holds.

B. $PK_2\{(\lambda_I, A, c, z, I, z')$: $P_I = g_4^{\lambda_I} \wedge \hat{e}(A, \omega g^c) = \hat{e}(g g_0^z g_1^I g_3^s, g)\}$

First, the DRP sends a challenge F. The customer randomly picks $\theta_1, \theta_2 \in \mathbb{Z}_p$, and computes $\Theta_1 = h_1^{\theta_1} h_2^{\theta_2}$, $\Theta_2 = A h_1^{\theta_2}$. He also chooses $\rho_w, \rho_{\theta_1}, \rho_{\theta_2}, \rho_{z'}, \rho_I, \rho_c, \rho_{\beta_1}, \rho_{\beta_2} \in \mathbb{Z}_p$, computes

$$\begin{aligned}
T_1 &= g_4^{\rho_w}, \\
T_2 &= h_1^{\rho_{\theta_1}} h_2^{\rho_{\theta_2}}, \\
T_3 &= \Theta_1^{-\rho_{\theta_2}} h_1^{\rho_{\beta_1}} h_2^{\rho_{\beta_2}}, \\
T_4 &= \hat{H}_0^{\rho_{\theta_2}} \hat{H}_1^{\rho_{\beta_2}} \hat{E}_0^{\rho_{z'}} \hat{E}_1^{\rho_I} \hat{e}(\theta_2, g)^{-\rho_c}.
\end{aligned}$$

Next, the customer sends them to a random oracle, and obtains $m = \mathcal{H}(\Theta_1, \Theta_2, T_1, T_2, T_3, T_4, F)$. The customer computes $k_w = \rho_w - m\lambda_I$, $k_{\theta_1} = \rho_{\theta_1} - m\theta_1$, $k_{\theta_2} = \rho_{\theta_2} - m\theta_2$, $k_{z'} = \rho_{z'} - mz'$, $k_I = \rho_I - m \cdot ID$, $k_s = \rho_s - ms$, $k_c = \rho_c - mc$,

$k_{\beta_1} = \rho_{\beta_1} - m\theta_1 c$, $k_{\beta_2} = \rho_{\beta_2} - m\theta_2 c$ and sends them together with m, Θ_1, Θ_2 to the DRP. Finally, the DRP computes

$$\begin{aligned} T'_1 &= P_I^m g_4^{k_w}, \\ T'_2 &= \Theta_1^m h_1^{k_{\theta_1}} h_2^{k_{\theta_2}}, \\ T'_3 &= \Theta_1^{-k_c} h_1^{k_{\beta_1}} h_2^{k_{\beta_2}}, \\ T'_4 &= (\hat{e}(\Theta_2, \omega)\hat{E}^{-1}\hat{E}_3^{-s})^m \hat{H}_0^{k_{\theta_2}} \hat{H}_1^{k_{\beta_2}} \hat{E}_0^{k_I} \hat{E}_1^{k_B} \hat{e}(\Theta_2, g)^{-k_c}. \end{aligned}$$

If the DRP verifies that $m = \mathcal{H}(\Theta_1, \Theta_2, T'_1, T'_2, T'_3, T'_4, F)$, it will accept the proof.

C. $PK_3\{(I, s, z') : C = g_0^{z'} g_1^I g_3^s\}$

The construction of PK_3 is part of PK_5.

D. $PK_4\{\lambda_I : P_I = g_4^{\lambda_I}\}$

The construction of PK_4 is contained in PK_2.

E. $PK_5\{(\lambda_I, \tilde{A}, \tilde{c}, I, \tilde{z}, \tilde{B}, \tilde{s}, z') : P_I = g_4^{\lambda_I} \wedge B - d > 0 \wedge C = g_0^{z'} g_1^I g_2^{\tilde{B}} g_3^s \wedge \hat{e}(\tilde{A}, \omega g^{\tilde{c}}) = \hat{e}(g g_0^{\tilde{z}} g_1^I g_3^{\tilde{s}}, g)\}$

The first statement $P_I = g_4^{\lambda_I}$ is the same as the one in PK_2, and thus we only describe how to construct the ZKP for the rest of the statements. First, the DRP sends a challenge F. The customer randomly picks $\theta_1, \theta_2 \in \mathbb{Z}_p$, and computes $\Theta_1 = h_1^{\theta_1} h_2^{\theta_2}$, $\Theta_2 = \tilde{A} h_1^{\theta_2}$. The customer chooses $\rho_{\theta_1}, \rho_{\theta_2}, \rho_{z'}, \rho_I, \rho_{\tilde{B}}, \rho_s, \rho_{\tilde{z}}, \rho_{\tilde{c}}, \rho_{\beta_1}, \rho_{\beta_2} \in \mathbb{Z}_p$, and computes

$$\begin{aligned} T_1 &= g_0^{\rho_{z'}} g_1^{\rho_I} g_2^{\rho_{\tilde{B}}} g_3^{\rho_s}, \\ T_2 &= h_1^{\rho_{\theta_1}} h_2^{\rho_{\theta_2}}, \\ T_3 &= \Theta_1^{-\rho_{\tilde{c}}} h_1^{\rho_{\beta_1}} h_2^{\rho_{\beta_2}}, \\ T_4 &= \hat{H}_0^{\rho_{\theta_2}} \hat{H}_1^{\rho_{\beta_2}} \hat{E}_0^{\rho_{\tilde{z}}} \hat{E}_1^{\rho_I} \hat{E}_2^{\rho_{\tilde{B}}} \hat{e}(\Theta_2, g)^{-\rho_{\tilde{c}}}. \end{aligned}$$

Then the customer sends $\Theta_1, \Theta_2, T_1, T_2, T_3, T_4, F$ to the random oracle, and obtains $m = \mathcal{H}(\Theta_1, \Theta_2, T_1, T_2, T_3, T_4, F)$. The customer computes $k_{\theta_1} = \rho_{\theta_1} - m\theta_1$, $k_{\theta_2} = \rho_{\theta_2} - m\theta_2$, $k_{z'} = \rho_{z'} - mz'$, $k_s = \rho_s - ms$, $k_{\tilde{c}} = \rho_{\tilde{c}} - m\tilde{c}$, $k_I = \rho_I - m \cdot ID$, $k_{\tilde{B}} = \rho_{\tilde{B}} - m\tilde{B}$, $k_{\beta_1} = \rho_{\beta_1} - m\theta_1 c$, $k_{\beta_2} = \rho_{\beta_2} - m\theta_2 c$, and sends them together with m, Θ_1, Θ_2 to the DRP. Finally, the DRP computes

$$\begin{aligned} T'_1 &= C^m g_0^{k_{z'}} g_1^{k_I} g_2^{k_{\tilde{B}}} g_3^{k_s}, \\ T'_2 &= \Theta_1^m h_1^{k_{\theta_1}} h_2^{k_{\theta_2}}, \\ T'_3 &= \Theta_1^{-k_{\tilde{c}}} h_1^{k_{\beta_1}} h_2^{k_{\beta_2}}, \\ T'_4 &= (\hat{e}(\Theta_2, \omega)\hat{E}^{-1}\hat{E}_1^{-ID}\hat{E}_2^{-\tilde{B}}\hat{E}_3^{-s})^m \cdot \\ &\quad \hat{H}_0^{k_{\theta_2}} \hat{H}_1^{k_{\beta_2}} \hat{E}_0^{k_{\tilde{z}}} \hat{E}_1^{k_I} \hat{E}_2^{k_{\tilde{B}}} \cdot \hat{e}(\Theta_2, g)^{-k_{\tilde{c}}}. \end{aligned}$$

The DRP accepts the zero knowledge proof if and only if $m = \mathcal{H}(\Theta_1, \Theta_2, T'_1, T'_2, T'_3, T'_4, F)$.

F. $PK_6\{(A, c, z, I, z') : \hat{e}(A, \omega g^c) = \hat{e}(g g_0^z g_1^I g_3^s, g)\}$

The construction of PK_6 is part of PK_2.

VIII. Privacy and Security Analysis

In this section, we show that the proposed scheme achieves the design goals of privacy and integrity. We also give formal security analysis in the appendix A.

A. Privacy

We preserve customer privacy by ensuring the anonymity of fine-grained metering data. Each customer registers both a real identity and a pseudonym, and only pseudonyms are attached to metering data. The DRP can neither infer the real identities from the metering data in the metering process nor link real identities and pseudonyms in other processes.

1) Anonymity of Metering Data: During the metering process, smart meters only attach pseudonyms to metering data. To verify the authenticity of the metering data, they sign the metering data with ICS. The secret value μ of the ICS is stored locally at the smart meter, and the DRP does not know it. Hence, due to the anonymity property of the ICS, the DRP can verify the data source, but it cannot identify the customer [20].

In the registration process, BBS+ signatures are used to hide the relationship between the real identity and the pseudonym. A customer obtains a BBS+ signature (i.e., the registration ticket in Sec. V) after he registers his real identity. In a separate communication session, the customer uses this signature to prove his eligibility of enrollment and to register his pseudonym. Since the BBS+ signature hides the value of the real identity, the DRP does not know his real identity when registering the pseudonym and thus cannot link these two identities. The same conclusion can be given for the settlement process and the revocation process.

In the querying process, customers inquire their data through pseudonyms and no information on real identity is involved, and thus the DRP learns nothing about their real identities. Besides, in order to obtain their data, customers need to enclose a ZKP in the querying request which proves knowledge of the secret of the pseudonym. As a result, nobody except the customer himself could learn his fine-grained metering data through the querying process. This provides an additional layer of protection to customer privacy.

Since information involving pseudonyms is sent through a proxy who hides the static physical address of a smart meter from the DRP, anonymity is also ensured in the physical layer.

2) Unlinkability Between Pseudo Accounts and Identifiable Accounts: In Sec. VI, we show that the relationship between rewards and withdrawals may compromise anonymity and propose two cloaking mechanisms to mitigate the attack. The cloaking mechanisms divide customers into several sets and customers in the same set are indistinguishable. Denote the set as S. A set with larger size provides stronger anonymity.

Suppose the DRP has 500 subscribed customers and customers withdraw money once per settlement cycle (e.g., a month). We assume that DR rewards follow a Gaussian distribution with mean 50 and variance 20. We simulate the withdrawal behaviors of customers with the FFW mechanism, and the PRS mechanism under two parameter settings, i.e., PRS with cells $\{1, 5, 10, 50\}$ and probability $p = 0.2$, and PRS with cells $\{1, 5, 10, 50\}$ and probability $p = 0.5$. We demonstrate the ratio of customers with different sizes of S in Fig. 7. Overall, most of the customers are indistinguishable at least from 9 others. However, as the DRP gradually gains more information, the sets are becoming smaller, and customers need to update their pseudonyms. We compare their average undrawn amounts on a monthly basis in Fig. 8. We can

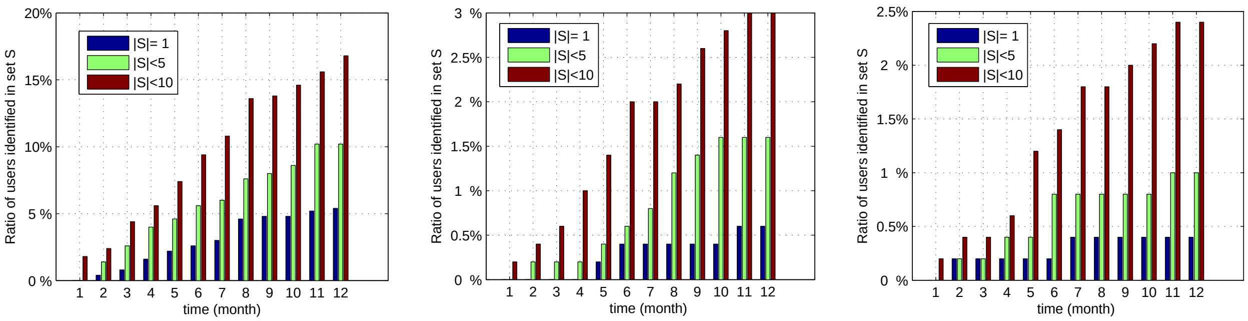

(a)

(b)

(c)

Fig. 7. Anonymity for FFW, PRS with $p = 0.2$ and PRS with $p = 0.5$.

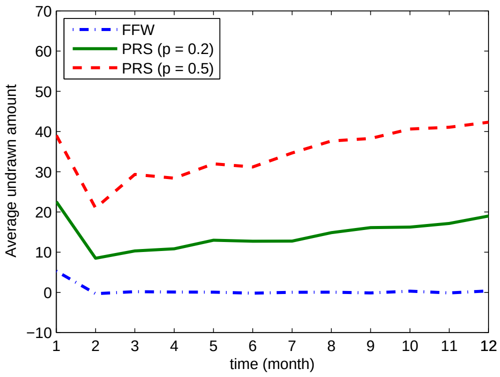

Fig. 8. Evaluation of Cloaking Mechanisms

see that the FFW mechanism requires the fewest amount of undrawn rewards among all the three approaches, while the PRS mechanism with selection probability $p = 0.5$ requires the most. This illustrates the trade-off between privacy and timeliness: If you want better privacy, you should withdraw less frequently.

B. Data Integrity

Data integrity of the proposed scheme is guaranteed in the following aspects.

Authenticity. The authenticity of the metering data, ensured by the ICS scheme, allows the DRP to verify that the data are generated by a genuine and registered smart meter. No attackers can forge or tamper the metering data since they cannot forge the ICS under adaptively chosen message attack [20].

Tickets are used in our scheme for registration process, revocation process, and settlement process. The tickets, or BBS+ signatures, have been proven to be secure against existential attack [26]. Hence attackers cannot forge tickets to gain monetary benefits. An important feature with the tickets is that each ticket can only be used once to avoid double spending. This is accomplished by the random secret s committed in the ticket. When a ticket is shown to the DRP, the DRP parses the ticket and extracts the secret s in the ticket. The DRP maintains a list of previously shown secrets and determines whether the newly parsed secret is in the list. Hence, a malicious user who replays used tickets will be discovered by the DRP.

Confidentiality. In addition to protocols we described in the paper, standard asymmetric and symmetric encryptions are used to provide confidentiality. For normal communication, either encryption scheme is good. For anonymous communication, asymmetric encryption schemes are required. For example, in the metering process, smart meters need asymmetric encryption to ensure confidentiality. To this end, they encrypt messages with the public key of the DRP, who then decrypts them with its private key. This ensures end-to-end confidentiality of the metering data.

Binding between the Metering Data and the Real Identity. Malicious customers can also cheat in re-identification process by presenting usage profiles owned by other customers. However, the success of this attack requires proof of ownership on power usage profiles. In our scheme, the binding property of ICS ensures that the real identity I is bound with the metering records, and thus a probabilistic polynomial-time algorithm adversary cannot output a different identity I' from the identity I that has been bound in the metering records [20].

Consistency of the Pseudonym. In the registration process, the pseudonym is chosen by the customer and sent to the smart meter. An adversary may send a different pseudonym from what he registered. However, he could neither endanger the availability of the IDR programs, nor gain economical benefits with this move. The reason is obvious. The DRP knows all the pseudonyms that have been registered. If the adversary uses an unregistered pseudonym or a registered pseudonym of other customers, the DRP can easily detect it.

IX. Performance Analysis

In this section, we analyze the efficiency and cost of the proposed scheme. The computation cost comes mainly from pairings and exponentiations in signature schemes (ICS, BBS+) and ZKP ($PK_1 - PK_6$). We summarize number of these two operations in basic protocols for smart meters, customers, and the DRP in Table I. From this table, we can see that smart meters do not need to perform any of the two operations. The most time consuming processes performed by smart meters are ICS generation in the metering process, which involves no paring or exponentiation operations and can be handled by existing smart meters efficiently.

The customer devices or DRP servers are assumed to be powerful enough to conduct computation-intensive operations.

TABLE I
Number of Pairing and Exponentiation Operations

	Registration		Metering			Querying		Settlement	
	Customer	DRP	Customer	Smart Meter	DRP	Customer	DRP	Customer	DRP
Group $\mathbb{G}$ exponentiation (pre-processed)	22	14	0	0	0	2	1	48	21
Group $\mathbb{G}$ exponentiation (direct)	1	6	0	0	0	0	1	3	9
Group $\mathbb{G}_T$ exponentiation (pre-processed)	4	6	0	0	0	0	5	13	16
Group $\mathbb{G}_T$ exponentiation (direct)	1	2	0	0	0	0	1	2	3
Pairing (one parameter is constant)	3	2	0	0	5	0	1	6	2
Pairing (both parameters are not constant)	2	0	0	0	0	0	0	2	0

However, some of the operations, such as the exponentiations that have constant base (e.g., $g_0^{z'}$) and the pairings with one of the parameter being constant (e.g., $\hat{e}(gg_0^z g_1^I g_3^s, g)$), can be preprocessed, which expedites the calculation greatly.

Based on simulation results of [28] that uses similar cryptographic tools, we can estimate the computation time of our algorithm. If we use a smart phone HTC Desire HD with QSD8255 1GHz CPU and 1.5G ROM to simulate the customer device, the registration time for the customer is less than 3 seconds, and the settlement time is less than 6 seconds. If we use a desktop with Q6600 2.4GHz CPU and 3GB RAM as the server of DRP, the registration time for the DRP is 0.2 seconds, the metering time is 0.05 seconds, and the settlement time is 0.3 seconds. Since registration and billing processes happen at low frequency, the processing time in the order of a few seconds is insignificant. Since smart meters are not involved in any computation intensive operations such as exponentiation or pairing operations, our proposed protocols can be implemented efficiently.

X. Conclusion

In this paper, we have identified and addressed the unique privacy issues in incentive-based demand response (IDR) programs. We have proposed a scheme which provides fine-grained metering data to the demand response provider (DRP) for basic operations, ensuring data integrity throughout all the processes. The scheme protects customer privacy by separating the real identity and the fine-grained metering data, i.e., the DRP can only learn either the real identity or the fine-grained metering data at a time but cannot link them together. In the case when re-identification is required, the linkage between real identity and metering data can be easily restored. Hence, our scheme provides an integrated solution for privacy-aware IDR programs, which promotes the acceptance of IDR programs.

Appendix

We analyze the security of our scheme with a game-based approach. Our security goals include privacy and integrity. Each security goal is modeled as a game played between a probabilistic polynomial time adversary $\mathcal{A}$ and a challenger $\mathcal{C}$. The game is defined so that it captures the capabilities and behavior of an adversary. The adversary winning the game implies that it may defeat a security goal. Using reduction argument, we would then show any adversary winning the game could be used to forge a BBS+ signature, which have been proven to be unforgeable in [26]. We also provide metering authenticity with ICS signature in our scheme, the authenticity of which has been proved in [20]. Some of the notations we use in the appendix have been introduced in VII. The security proof is adapted based on [29] and [28].

A. Integrity

We model the interaction between a cheating user $\mathcal{A}$ and an honest DRP $\mathcal{C}$ as the following game. The DRP $\mathcal{C}$ keeps a running balance B possessed by user $\mathcal{A}$. User $\mathcal{A}$ wins the game if it can make B to be negative. In this game we allow $\mathcal{A}$ to register multiple times, which models the situation when several users collude together. Of course, the interaction only involves registration an settlement because the balance tickets are only created and updated in these two processes.

System Parameter. $\mathcal{C}$ creates and publishes the system parameter and keeps the secret key γ private.

Interactions. $\mathcal{A}$ can make the following two types of interaction freely with $\mathcal{C}$.

- *Registration.* $\mathcal{A}$ interacts with $\mathcal{C}$ in the registration process. Upon successful completion of the process, B is initialized by $\mathcal{C}$ as 0.
- *Settlement.* $\mathcal{A}$ interacts with $\mathcal{C}$ in the settlement process to claim reward of value d.

Winning. $\mathcal{A}$ wins the game if there exists a sequence of interaction query so that B becomes negative.

Proof. We prove our security by reduction. Specifically, assume there exists $\mathcal{A}$, we show how to construct a forgery attack against the underlying BBS+ signature. Sine BBS+ signature is known to be unforgeable, this means no PPT adversary $\mathcal{A}$ can win in the above game. That is, our system prevents users from cheating.

Before stating our proof, let us assume the zero-knowledge proof PK_3 and PK_5 are both sound. That is, given black-box access to the prover that makes these ZKPs, there exists extractor algorithm EX_{BL} and EX_4 which output the witnesses used by the prover. Indeed, the ZKP protocols described in our paper are sound in the random oracle model.

Next, we describe an algorithm, called simulator, $\mathcal{S}$, which provides the view to $\mathcal{A}$ as the honest DRP and at that same time forges a BBS+ signature. $\mathcal{S}$ is given the public key of the BBS+ signature in the form of $(\hat{e}, \mathbb{G}, \mathbb{G}_T, g, g_0, g_1, g_2, \omega)$, together with a black-box SO, normally referred to as signing oracle. SO outputs a BBS+ signature (A, c, z) on any input message (m_1, m_2). $\mathcal{S}$ successfully forges a BBS+ signature if it can output a valid signature (A^*, c^*, z^*) on message (m_1^*, m_2^*) such that the former is not the output of SO.

Now we describe the behavior of $\mathcal{S}$. It sets public parameter as $(\hat{e}, \mathbb{G}, \mathbb{G}_T, g, g_0, g_1, g_2, \omega)$ and gives it to $\mathcal{A}$. Note that $\mathcal{S}$ does not know the secret key of the DRP $\mathcal{C}$ but the public parameter is distributed correctly. Below we show how $\mathcal{S}$ interacts with $\mathcal{A}$ in each of the possible interactions which involve the balance tickets.

- Registration. Upon executing PK_3 with $\mathcal{A}$, $\mathcal{S}$ uses EX_{BL} to extract the witness (I, z', s). $\mathcal{S}$ issues a signature query with input (I, B, s) to SO, where $B = 0$. $\mathcal{S}$ receives (A, c, z), computes $z'' = z - z'$, and returns (A, z'', c) to $\mathcal{A}$.
- Settlement. Upon executing PK_5 with $\mathcal{A}$, $\mathcal{S}$ uses EX_4 to extracts the witness $(\tilde{A}, \tilde{c}, \tilde{z}, I, \tilde{B}, z', s)$. If $\tilde{A}, \tilde{c}, \tilde{z}$ is not the output of SO, $\mathcal{S}$ outputs them as the forgery on $I, \tilde{B}, \tilde{s}$. Otherwise, it checks if $\tilde{s}$ is fresh and issues a signature query with input $(I, \tilde{B} - d, s)$ to SO. S receives (A, c, z) from $\tilde{A}$ and computes $z'' = z - z'$. It returns (A, z'', c) to $\mathcal{A}$. $\mathcal{S}$ sets $B = B - d$.

Due to the setting of the game, the value B remains positive if $\mathcal{S}$ never aborts. This is because in order to reduce the value of B, $\mathcal{A}$ has to interact with $\mathcal{S}$ in the settlement process and the number of signatures given to $\mathcal{A}$ via $\mathcal{S}$ is limited and that PK_5 assures $\mathcal{S}$ will not accept on message of the form $(\cdot, B, \cdot)$ with $B < d$. Thus, in order for $\mathcal{A}$ to win the game, $\mathcal{S}$ will abort and obtain a forgery to the underlying BBS+ signature.

B. Privacy

The following game models the user privacy. The rationale is that the curious DRP cannot tell which one of two honest users is responsible for a particular interaction under the extreme condition that all other interaction sequences have been specified by the curious DRP. The particular interaction could be during pseudonym registration, settlement, querying, and revocation, but not during real identity registration or metering, because real identity is to be known in real identity registration, and no real identity is revealed by the smart meter in metering. Our definition also guarantee that the settlement interactions are unlinkable.

System Parameter. The malicious adversary $\mathcal{A}$ creates and publishes the system parameter.

Interactions. Adversary $\mathcal{A}$ can make the following four types of interactions freely with $\mathcal{C}$, who acts on behalf of two honest users U_0, U_1.

- *Registration ($b \in \{0,1\}$).* $\mathcal{A}$ interacts with $\mathcal{C}$ who acts on behalf of U_b in the registration protocol. The value b is specified by $\mathcal{A}$.
- *Querying ($b \in \{0,1\}$).* $\mathcal{A}$ interacts with $\mathcal{C}$ who acts on behalf of U_b in the querying process. The value b is specified by $\mathcal{A}$.
- *Settlement ($b \in \{0,1\}$).* $\mathcal{A}$ interacts with $\mathcal{C}$ who acts on behalf of U_b in the settlement process with value d. The value b is specified by $\mathcal{A}$.
- *Revocation ($b \in \{0,1\}$).* $\mathcal{A}$ interacts with $\mathcal{C}$ who acts on behalf of U_b in the revocation process. The value b is specified by $\mathcal{A}$.

Challenge. $\mathcal{A}$ chooses a type of interaction among pseudonym registration, querying, settlement, and revocation. Note that both U_1 and U_0 should have sufficient balance in the settlement process. $\mathcal{C}$ flips a fair coin $\hat{b} \in \{0,1\}$ and interacts with $\mathcal{A}$ on behalf of user $U_{\hat{b}}$.

Winning. $\mathcal{A}$ outputs a guess bit b and wins the game if $b = \hat{b}$.

Proof. Our security proof is to show that probability of $\mathcal{A}$ winning the game is always 0.5. That is, the action of two honest users are completely indistinguishable. The view of $\mathcal{A}$ is provided by a simulator $\mathcal{S}$ who has control over the random oracle. Next we describe the behavior of $\mathcal{S}$:

- *Registration.* $\mathcal{S}$ acts on behalf of the two users honestly.
- *Querying.* $\mathcal{S}$ acts on behalf of the two users honestly.
- *Settlement.* $\mathcal{S}$ acts on behalf of the two users honestly.
- *Revocation.* $\mathcal{S}$ acts on behalf of the two users honestly.

In the *Challenge Phase*, $\mathcal{S}$ first checks if both users are eligible to participate in the interaction. That is, they are eligible for registration in pseudonym registration process, or they are having sufficient balance if the interaction is settlement. Then $\mathcal{S}$ flips a fair coin $\hat{b} \in \{0,1\}$. The following protocols are slightly different for different interactions such as pseudonym registration and querying. For example, for PK_2 in pseudonym registration process, $\mathcal{S}$ randomly picks $P_I \in \mathbb{Z}_p, \tilde{s} \in \mathbb{Z}_p$ and sends them to $\mathcal{A}$. Then $\mathcal{A}$ chooses a random challenge F and send it to $\mathcal{S}$. Upon receiving the random challenge F, $\mathcal{S}$ randomly chooses $\Theta_1, \Theta_2, m, k_w, k_{\theta_1}, k_{\theta_2}, k_{z'}, k_I, k_s, k_c, k_{\beta_1}, k_{\beta_2} \in \mathbb{Z}_p$ and computes

$$\begin{aligned}
T_1 &= P_I^m g_4^{k_w}, \\
T_2 &= \Theta_1^m h_1^{k_{\theta_1}} h_2^{k_{\theta_2}}, \\
T_3 &= \Theta_1^{-k_c} h_1^{k_{\beta_1}} h_2^{k_{\beta_2}}, \\
T_4 &= (\hat{e}(\Theta_2, \omega)\hat{E}^{-1}\hat{E}_3^{-s})^m \hat{H}_0^{k_{\theta_2}} \hat{H}_1^{k_{\beta_2}} \hat{E}_0^{k_I} \hat{E}_1^{k_B} \hat{e}(\Theta_2, g)^{-k_c}.
\end{aligned}$$

Finally, $\mathcal{S}$ sets $m = \mathcal{H}(\Theta_1, \Theta_2, T_1, T_2, T_3, T_4, F)$. This is possible since $\mathcal{S}$ is in control of the random oracle and can decide what value to return as the output of the random oracle. $\mathcal{S}$ sends $\Theta_1, \Theta_2, m, k_w, k_{\theta_1}, k_{\theta_2}, k_{z'}, k_I, k_s, k_c, k_{\beta_1}, k_{\beta_2} \in \mathbb{Z}_p$ to $\mathcal{A}$ as PK_2 in the interaction. We can get interactions between $\mathcal{S}$ and $\mathcal{A}$ for $PK_3 - PK_6$ in a similar way. Note that the values are correctly distributed and can be based on the storage of U_0 or U_1. For any valid storage such as $(\tilde{\delta}_s, I, \tilde{s})$ or $(\tilde{\delta}_s, I, \tilde{B}, \tilde{s})$, there exists a set of randomness that maps it to the view of the above protocol and that the value $\tilde{s}$ is completely hidden from $\mathcal{A}$. Thus the value $\check{b}$ is completely hidden from the view of $\mathcal{A}$ and the probability that $\mathcal{A}$ guess correctly is always 0.5.

References

[1] U. DOE, "Benefits of demand response in electricity markets and recommendations for achieving them," 2006.

[2] A. Ipakchi and F. Albuyeh, "Grid of the future," *Power and Energy Magazine, IEEE*, vol. 7, no. 2, pp. 52–62, 2009.

[3] C. Goldman, M. Reid, R. Levy, and A. Silverstein, "Coordination of energy efficiency and demand response," Lawrence Berkeley National Laboratory, Tech. Rep., 2010.

[4] "The demand response baseline," EnerNoc, 2008. [Online]. Available: http://www.naesb.org/pdf4/dsmee_group2_022609w2.pdf

[5] G. W. Hart, "Non-intrusive appliance load monitoring," *Proc. of the IEEE*, vol. 80, no. 12, pp. 1870–1891, 1992.

[6] D. C. Bergman, D. Jin, J. P. Juen, N. Tanaka, C. A. Gunter, and A. K. Wright, "Distributed non-intrusive load monitoring," in *Proc. IEEE PES ISGT*, 2011, pp. 1–8.

[7] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin, “Private memoirs of a smart meter,” in *ACM Workshop on Embedded Sensing Systems for Energy-efficiency in Building*, 2010, pp. 61–66.

[8] E. Quinn, “Smart metering and privacy: Existing laws and competing policies,” *Available at SSRN 1462285*, 2009.

[9] S. McLaughlin, P. McDaniel, and W. Aiello, “Protecting consumer privacy from electric load monitoring,” in *Proc. ACM CCS*, 2011, pp. 87–98.

[10] M. A. Lisovich, D. K. Mulligan, and S. B. Wicker, “Inferring personal information from demand-response systems,” *IEEE Security & Privacy*, vol. 8, no. 1, pp. 11–20, 2010.

[11] F. D. Garcia and B. Jacobs, “Privacy-friendly energy-metering via homomorphic encryption,” in *Security and Trust Management*. Springer, 2011, pp. 226–238.

[12] K. Kursawe, G. Danezis, and M. Kohlweiss, “Privacy-friendly aggregation for the smart-grid,” in *Proc. 11th Int. Conf. on Privacy Enhancing Technologies*, 2011, pp. 175–191.

[13] Z. Erkin and G. Tsudik, “Private computation of spatial and temporal power consumption with smart meters,” in *Proc. 10th Int. Conf. on Applied Cryptography and Network Security*, 2012, pp. 561–577.

[14] F. Li, B. Luo, and P. Liu, “Secure information aggregation for smart grids using homomorphic encryption,” in *Proc. IEEE SmartGridComm*, 2010, pp. 327–332.

[15] A. Rial and G. Danezis, “Privacy-preserving smart metering,” in *Proc. 10th ACM Workshop on Privacy in the Electronic Society*, 2011, pp. 49–60.

[16] C. Efthymiou and G. Kalogridis, “Smart grid privacy via anonymization of smart metering data,” in *Proc. IEEE SmartGridComm*.

[17] G. Kalogridis, C. Efthymiou, S. Z. Denic, T. A. Lewis, and R. Cepeda, “Privacy for smart meters: Towards undetectable appliance load signatures,” in *Proc. IEEE SmartGridComm*.

[18] A. Shamir, “Identity-based cryptosystems and signature schemes,” in *Advances in cryptology*. Springer, 1985, pp. 47–53.

[19] J. C. Cha and J. H. Cheon, “An identity-based signature from gap diffie-hellman groups,” in *PKC 2003*, 2002, pp. 18–30.

[20] C.-K. Chu and W.-G. Tzeng, “Identity-committable signatures and their extension to group-oriented ring signatures,” in *Information Security and Privacy*, 2007, pp. 323–337.

[21] S. Goldwasser, S. Micali, and C. Rackoff, “The knowledge complexity of interactive proof systems,” *SIAM Journal on computing*, vol. 18, no. 1, pp. 186–208, 1989.

[22] A. Fiat and A. Shamir, “How to prove yourself: practical solutions to identification and signature problems,” in *CRYPTO*, 1987, pp. 186–194.

[23] J. Camenisch and M. Stadler, “Efficient group signature schemes for large groups,” in *CRYPTO*. Springer, 1997, pp. 410–424.

[24] T. P. Pedersen, “Non-interactive and information-theoretic secure verifiable secret sharing,” in *CRYPTO*, 1992, pp. 129–140.

[25] D. Boneh, X. Boyen, and H. Shacham, “Short group signatures,” in *CRYPTO*, 2004, pp. 41–55.

[26] M. H. Au, W. Susilo, and Y. Mu, “Constant-size dynamic k-taa,” in *Proc. 5th Int. Conf. on Security and Cryptography for Networks*, 2006, pp. 111–125.

[27] K. Coughlin, M. A. Piette, C. Goldman, and S. Kiliccote, “Estimating demand response load impacts: evaluation of baseline load models for non-residential buildings in california,” *Lawrence Berkeley National Laboratory*, 2008.

[28] J. K. Liu, M. H. Au, W. Susilo, and J. Zhou, “Enhancing location privacy for electric vehicles (at the right time),” in *ESORICS 2012*.

[29] J. Camenisch and A. Lysyanskaya, “An efficient system for non-transferable anonymous credentials with optional anonymity revocation,” in *Advances in Cryptology*. Springer, 2001, pp. 93–118.