Privacy Enhanced Demand Response with Reputation-based Incentive Distribution

Matteo Vasirani, Tri Kurniawan Wijaya, Georgios Liassas, Karl Aberer
School of Computer and Communication Sciences
LSIR, EPFL, Switzerland
{matteo.vasirani, tri-kurniawan.wijaya, georgios.liassas, karl.aberer}@epfl.ch

ABSTRACT

Demand response (DR) is a Smart Grid application that aims at managing consumption of electricity in response to supply-side signals. The exchange of sensitive information (i.e., electricity consumption data) between the consumers and the DR provider is necessary for the functioning of DR, but at the same time it is an obstacle to its wide-spread adoption, due to the related privacy concerns. This paper proposes the use of homomorphic encrypted user aggregation and reputation-based incentive distribution to address the trade-off between enhancing user privacy and correctly assessing the contribution of each user to the demand reduction.

Categories and Subject Descriptors

E.3 [**Data Encryption**]: Public key cryptosystems

Keywords

homomorphic encryption, demand response

1. INTRODUCTION

Demand Response (DR) is an application that is considered a fundamental building block of the Smart Grid vision. The goal of DR is influencing the energy consumption pattern of users in response to supply-side signals. When a DR event is issued by a DR provider, a DR signal is sent to the users, containing the start and end time of the event, and optionally the amount of consumption to be reduced. DR is based on the exchange of privacy invasive energy consumption data between users and DR providers [6]. For example, in incentive-based DR, where an agreement between the DR provider and the energy consumers is established, it is necessary to have access to the user consumption data in order to assess how the user responded to the DR signal, and therefore the amount of incentives (e.g., money, discount vouchers redeemable at local shops, bill rebates) that have to be granted to the user. The collection of this type of data might entail many privacy related issues [7]. For instance, non-intrusive load monitoring techniques (NILM) can be used to infer the type and number of appliances, as well as their state (*on* or *off*), using only the data collected by the house-level meter. Also the activity of the user can be inferred from energy consumption data. For instance, Greveler *et al.* [4] employed a method to identify the displayed TV channels. They exploit smart meter measurements with sampling rate of 0.5 Hz to develop a function that predicts the power consumption of a LCD monitor lighting system. The power consumption of the monitor is correlated to the brightness of the content that is displayed, so that by inferring the energy consumption of the TV it is possible to infer the program or movie that is being displayed.

To alleviate the privacy-related problems, several Privacy Enhancing Techniques (PETs) such as anonymisation [5] or perturbation [2] are available to protect users against activity and behavioural analysis [1],[3]. On the other hand, these techniques usually imply an obfuscation phase that might deprive the DR provider from important information that impedes an accurate assessment of the DR performance of each user (i.e., demand reduction). As a matter of fact, there exists a trade-off between enhancing user privacy and accurately assessing the demand reduction.

In this paper we propose the use of *homomorphic encrypted user aggregation* as a way to enhance user privacy. With homomorphic encryption it is possible to sum up encrypted consumption time series of individual users to generate an aggregated time series of a group of users. Homomorphic encrypted user aggregation is very effective for DR. User aggregation helps to increase the privacy, since the collected information does not refer to a certain individual energy user (i.e., data producer), but to a group of users. Another advantage offered by homomorphic encryption is that exact aggregates can be calculated. Thus, the DR provider is able to accurately assess the demand reduction at the group level, and to allocate the corresponding incentives to each group. On the other hand, the lack of individual consumption data introduces error in the *distribution of incentives* to the group members, because different users in a group might have contributed to the group's demand reduction in a different way. There-

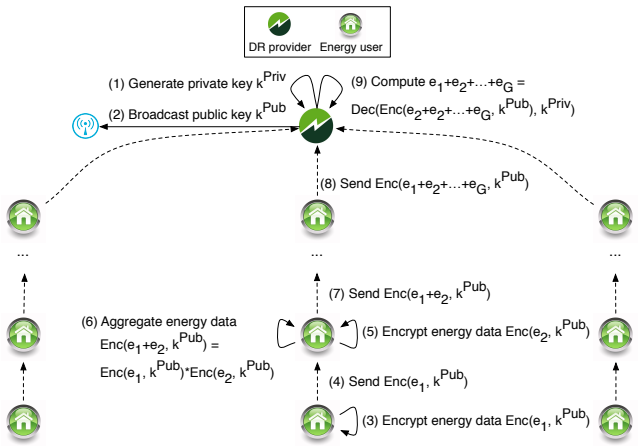

Figure 1: Homomorphic encryption for user aggregation

fore, it is possible that users that contributed the most to the group's demand reduction receive less incentives than those they deserve or, even worse, users that did not contribute at all to the group's demand reduction receive an undeserved share of the group incentives. For this reason, we also define a *reputation-based incentive distribution* mechanism that aims at minimising the misallocation of incentives by creating groups of users with similar response to DR signals.

The paper is structured as follows: Section 2 details the homomorphic encryption protocol for user consumption data aggregation, while Section 3 describes our reputation-based mechanism for group formation; the metrics used for the experimental evaluation are defined in Section 4, while in Section 5 we outline the results of the experimental evaluation; finally, we conclude in Section 6.

2. HOMOMORPHIC ENCRYPTION

In this work we propose an aggregation protocol over a group of data producers based on the Paillier cryptosystem [8]. Figure 1 illustrates how the aggregation protocol works. The DR provider (1) generates the private key for decryption (k^{Priv}) and (2) broadcasts the Paillier public key (k^{Pub}). The public key k^{Pub} is used by the energy users to encrypt their individual energy consumption time series, while k^{Priv} is used by the DR provider to decrypt the aggregate of the individual energy measurements. The energy user that correspond to a leaf in the tree (3) Paillier-encrypts its consumption time series with the provided public key, and (4) sends its encrypted consumption $Enc(e_1, k^{Pub})$ to its direct parent, where e_1 is a time series of energy consumption values. The parent (5) Paillier-encrypts its consumption time series e_2 and then (6) homomorphically adds the result with the encrypted readings of its child in the tree. In the Paillier cryptosystem, the additions in plaintext are translated to multiplication in the ciphertext. Thus, in order to generate the encrypted version $Enc(e_1 + e_2, k^{Pub})$ of the sum of the two time series e_1 and e_2, the parent performs the multiplication operation directly on the encrypted time series $Enc(e_1, k^{Pub})$ and $Enc(e_2, k^{Pub})$. After that, (7) the parent sends the encrypted aggregated consumption of itself and its child to its respective parent. After G aggregations have been performed, where G is the size of the group, (8) the last parent sends the encrypted aggregated time series $Enc(e_1 + e_2 + \cdots + e_G, k^{Pub})$ to the DR provider, which (9) decrypts it with the private key k^{Priv} to obtain the energy consumption time series of the whole group, $e_1 + e_2 + \cdots + e_G$.

In this work we assume an *honest-but-curious* adversary model. Thus, the energy users will not tamper the protocol execution process, but at the same time they might try to read other users' intermediate energy aggregation results. The Paillier encryption used in the protocol ensures that no intermediate node will ever be able to disclose individual data, since none of the intermediate participants in the tree has the private key. In the end, only the DR provider will be able to decrypt the Paillier-encrypted aggregate. Furthermore, given a group size $G > 1$, the DR provider will not be able to distinguish the individual consumption of any of the participants.

We remark that the aggregation protocol proposed here is one of the possible protocols to generate the encrypted aggregated time series $Enc(e_1 + e_2 + \cdots + e_G, k^{Pub})$. Indeed, this time series can be constructed with different protocols and communication paths. For example, all the users in a group could elect a leader h that receives all the individual encrypted time series $Enc(e_i, k^{Pub})$ $\forall i \neq h$ and performs the multiplication operation on behalf of the group members. However, analysing the characteristics of different aggregation protocols (e.g., bandwidth, resiliency to communication failures, etc.) is not the scope of the paper.

3. INCENTIVE DISTRIBUTION

Each user i is characterised by the reduction rate $\gamma_i \geq 0$. We assume two types of users: *legitimate* and *free-riders*. A user is legitimate if $\gamma_i > 0$, while it is considered a free-rider if $\gamma_i = 0$. Given that free riders do not provide any demand reduction but they might still receive some incentives, the primary goal of an incentive distribution mechanism is to group together users with similar performance by separating legitimate users from free riders.

Let N be the total population of users, G the size of the groups created before each DR event by the aggregation protocol, and $N_G = N/G$ the corresponding number of groups. Furthermore, let $N^{\text{FR}} < N$ be the number of free riders in the population, and $\rho^{\text{FR}} = N^{\text{FR}}/N$ the corresponding percentage.

Each user i has an intended consumption ℓ_i (i.e., the

intended energy that would have been consumed if there were no DR event) and a realised consumption r_i (i.e., the actual load during the DR event), defined as

$$\ell_i = \sum_{t=t_{\text{start}}}^{t_{\text{end}}} e_i(t) \qquad r_i = \sum_{t=t_{\text{start}}}^{t_{\text{end}}} (1-\gamma_i) e_i(t) \quad (1)$$

where t_{start} and t_{end} are the start and end time of the DR event respectively, and $e_i(t)$ is the intended energy consumption at time t. Therefore, the contribution of user i to the group demand reduction is

$$\delta_i = \ell_i - r_i = \gamma_i \sum_{t=t_{\text{start}}}^{t_{\text{end}}} e_i(t) \quad (2)$$

For a group j, the intended and realised consumption are defined as

$$L_j = \sum_{i \in j} \ell_i \qquad R_j = \sum_{i \in j} r_i \quad (3)$$

We remark that the DR provider does not have access to the individual terms r_i. In fact, the DR provider can only accurately measure the aggregated value R_j, provided by the individual users through the homomorphic encryption protocol described in Section 2. Furthermore, we assume that, using past aggregated consumption data and possibly other contextual information (e.g., weather forecasts, calendar events, etc.), the DR provider is able to perfectly estimate the aggregated intended consumption L_j.

The demand reduction of group j can be computed as $\Delta_j = L_j - R_j$. After the DR event, the DR provider has to allocate a certain amount of incentives to each group. Let M be the total amount of incentives, the fraction allocated to a group j is defined as

$$M_j = \frac{\Delta_j}{\sum_{j=1}^{N_G} \Delta_j} \cdot M \quad (4)$$

Although the amount of incentives allocated to a group is accurate with respect to the demand reduction of the group, the DR provider has no error-free way of splitting M_j among the group members such that each of them receives its deserved share, given that the individual contribution δ_i is impossible to measure. For this reason, each group member receives an amount of incentives equal to

$$m_i = \frac{M_j}{G} \quad (5)$$

which may differ from the deserved amount, defined as

$$m_i^* = \frac{\delta_i}{\Delta_j} \cdot M_j \quad (6)$$

The difference between m_i and m_i^* determines the incentive distribution error committed by the DR provider, which is defined as

$$E = \frac{1}{N} \sum_{i=1}^{N} (|m_i^* - m_i|) \quad (7)$$

Although the DR provider is not able to measure δ_i, which is determined by the reduction rate γ_i, it might try to create groups of users with similar reduction rate γ_i, in the attempt of minimising the incentive distribution error.

For this reason, we defined a reputation-based mechanism that aims at creating groups composed of users with similar reduction rate γ_i. The reputation score of a user is defined as

$$\pi_i = \frac{1}{|\mathcal{D}|} \sum_{d \in \mathcal{D}} \frac{m_i(d)}{M_{i \in j}(d)} \quad (8)$$

where $\mathcal{D}$ is the set of DR events occurred so far, $m_i(d)$ is the incentive received by user i after DR event d, and $M_{i \in j}(d)$ is the incentive received by the group j that i belonged to during DR event d.

The objective of the reputation-based mechanism is grouping together users with similar reputation values. Prior to a DR event, the DR provider sorts the users in decreasing order based on the reputation score. Then the DR provider selects as the next user to place in the current group that is being formed either the first user of the list with probability $1-\epsilon$, or a random user with probability ϵ. Once the user has been selected, it is removed from the list, so that the first element of the list is always the user with the highest reputation score.

The probabilistic insertion rule is necessary in order to perform some exploration, otherwise after the first DR events always the same groups would be formed. To enforce convergence, we geometrically decrease the value of ϵ after a DR event d using the update scheme $\epsilon(d+1) = a \cdot \epsilon(d)$, where $a \in (0,1)$ is a parameter to control the speed of convergence and $\epsilon(0) \in (0,1)$ is the initial exploration rate. With this scheme, in the long run ϵ tends to 0, which implies that the same groups will be formed, based on the assumption that the reputation score of the users is accurate.

4. EVALUATION METRICS

There are several evaluation criteria to evaluate the user aggregation protocol and the reputation-based incentive distribution. In this section, we describe the following metrics that will be used for the experimental evaluation:

Group homogeneity is defined as:

$$\text{GH} = 1 - \frac{H^{\text{FR}}}{H^*} = 1 - \frac{\sum_{j=1}^{N_G} \frac{n_j^{\text{FR}}}{G} log\left(\frac{n_j^{\text{FR}}}{G}\right)}{N_G \cdot (\ \rho^{\text{FR}} \cdot log(\rho^{\text{FR}})\ )} \quad (9)$$

where H^{FR} is the entropy of percentage of free riders (n_j^{FR}/G) that are present in each group, and H^*

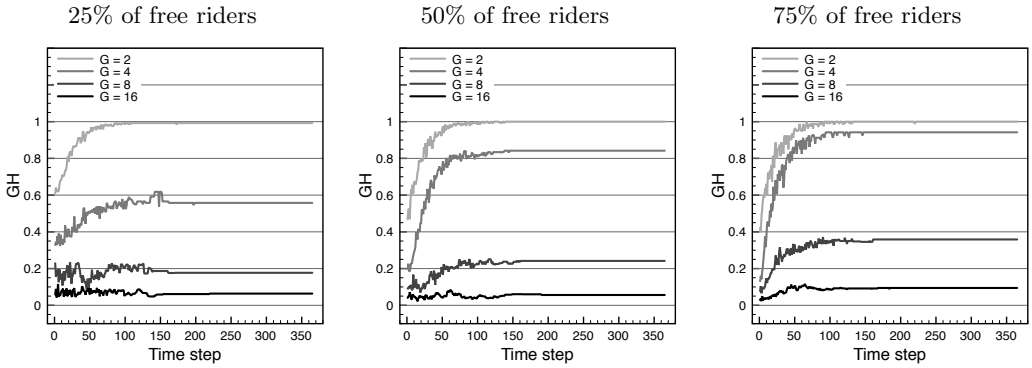

Figure 2: Group homogeneity (the higher GH, the better)

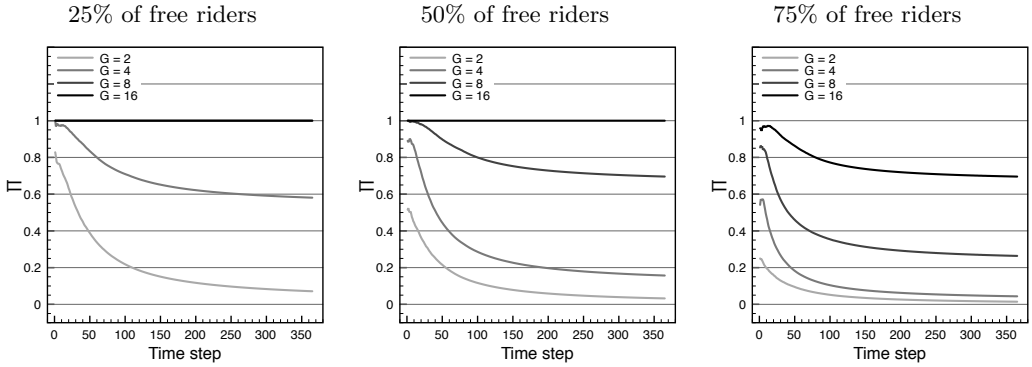

Figure 3: Average reputation of free riders (the lower Π^{FR}, the better)

is the maximal entropy of percentage of free riders, corresponding to a situation where all the created groups have the same percentage of free riders $\rho^{\mathrm{FR}} = N^{\mathrm{FR}}/N$ inside. When GH = 0, the group entropy is maximal, as each group is composed of legitimate users and free riders. GH approaches 1 when each group is composed of a single type of user: either all legitimate users or all free riders.

Average reputation of free riders is defined as:

$$\Pi^{\mathrm{FR}} = \frac{\frac{1}{N^{\mathrm{FR}}} \sum_{k=1}^{N^{\mathrm{FR}}} \pi_k}{\pi^*} \tag{10}$$

where N^{FR} is the total number of free riders, π_k is the reputation score of free rider k, and π^* is the maximum reputation value for a user in a group of size G. When the reputation is computed according to Eq. 8 and the incentive is equally split among group members, we have that $\pi^* = 1/G$. When the free riders have been isolated into homogeneous groups, the average reputation Π^{FR} approaches 0 because they do not provide any demand reduction nor they receive incentives.

Incentive distribution error reduction is defined as:

$$\mathrm{IER} = 100 \frac{E_{\mathrm{RND}} - E_{\mathrm{REP}}}{E_{\mathrm{RND}}} \tag{11}$$

where E_{REP} is the incentive distribution error of the reputation-based mechanism (see Eq. 7), and E_{RND} is the incentive distribution error of a benchmark mechanism for group creation that forms random groups. The metric IER quantifies the gains offered by the reputation-based mechanism.

Privacy is related to the aggregation of the energy consumption values of G users into a single value Y. An adversary is therefore faced with the task of disaggreating the single value Y into G values, one for each user, corresponding to the user original consumption values. To quantify the privacy of the aggregation of G values into a single aggregated value Y, we use the Shannon entropy associated with the disaggregation of Y into G values. In general, the entropy of a system with $\mathcal{S}$ states is expressed as $H = \sum_{s \in \mathcal{S}} -p(s) \cdot log(p(s))$, where $p(s)$ is the probability that the system is in state s. In our case, $\mathcal{S}$ is the set of all the possible disaggregations, i.e., all the possible ways a value Y can be split into G values such that the sum of the G values equals Y. The number of possible disaggregations (i.e., the state space size $|\mathcal{S}|$) is called weak integer composition of Y into G parts, and it is computed as

$$|\mathcal{S}| = \binom{Y+G-1}{G-1} = \frac{(Y+G-1)!}{(G-1)!\ Y!}. \tag{12}$$

Assuming that each disaggregation Y into G values has the same probability, we can rewrite the entropy as

$$H(Y) = log\left(|\mathcal{S}|\right). \tag{13}$$

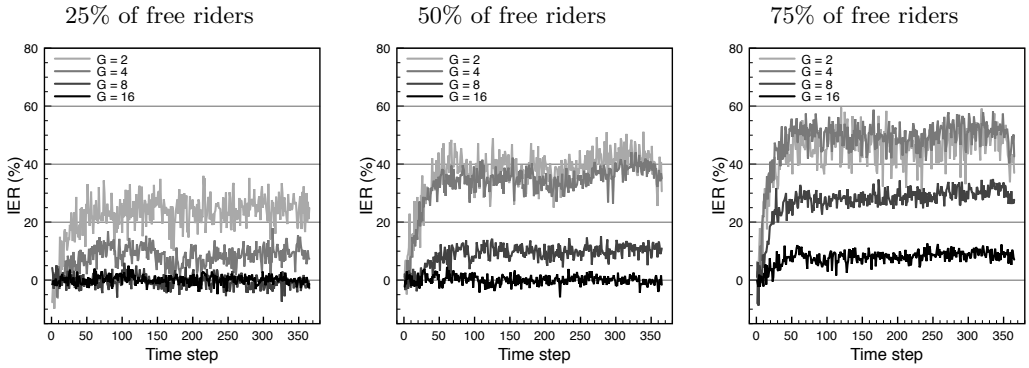

Figure 4: Incentive distribution error reduction (the higher IER, the better)

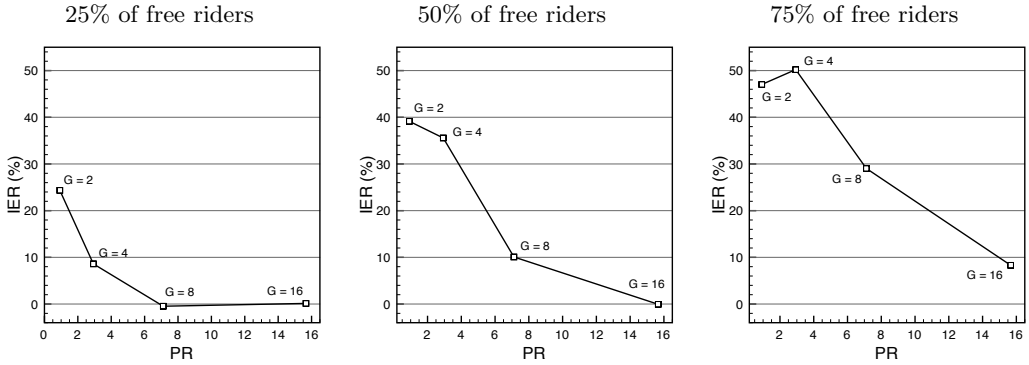

Figure 5: Trade-off between privacy and incentive distribution error reduction

The privacy metric is therefore defined as

$$\mathrm{PR} = \frac{1}{N_G|\mathcal{D}|} \sum_{d\in\mathcal{D}} \sum_{j=1}^{N_G} H(R_j(d)) \qquad (14)$$

where N_G is the number of groups, $\mathcal{D}$ is the set of DR events, and $H(R_j(d))$ is the realised consumption of group j during DR event d, that is, the value Y to be disaggregated by the adversary, as in Eq. 13.

5. EXPERIMENTAL EVALUATION

For the evaluation, we used the electricity usage dataset published by Ireland's Commission for Energy Regulation (CER) in 2012.[1] The project ran in the period 2009-2010 with 5000 residential and business energy consumers participating. The data are cleaned from missing values and filtered out to contain the energy consumption time series of 782 residential users that belong to the control group. From this set, we randomly picked 512 consumers, in order to evaluate groups of increasing size $G \in \{2, 4, 8, 16\}$. Legitimate users are assumed to have a reduction rate $\gamma_i = 0.2$. We performed several experiments with different percentages of free riders in the population of users, $\rho^{\mathrm{FR}} \in \{25\%, 50\%, 75\%\}$.

DR events of 3 hour duration are simulated with the dispatch of a DR signal to the participating users, which are aggregated into several groups of size G. The DR signal requests a demand reduction during 3 hours, from 18:00 in the evening to 21:00 in the night, which corresponds to the energy peak in Ireland. The parameters of the insertion rule $\epsilon(d+1) = a \cdot \epsilon(d)$ were selected by trial-and-error and set to $a = 0.95$ and $\epsilon(0) = 0.9$.

Figure 2 shows the evolution of the group homogeneity metric over time. The lower the value of the group homogeneity metric, the more able the free riders are to infiltrate all the groups. When the group size G is small (2 to 4 users per group), the reputation-based mechanism is able to learn with time to separate the legitimate users from the free riders and create groups with a single type of user. Furthermore, the bigger the percentage of free riders in the population of users, the easier is for the reputation mechanism to spot and isolate them. On the other hand, when the group size G is bigger (8 to 16 users per group), the mechanism is less able to spot the free riders, so that it is easier for them to spread among bigger groups. As a result, the group homogeneity is lower, because all groups tend to be invaded by free riders.

The successful separation of legitimate users and free riders performed by the reputation-based mechanism is also visible in the average reputation value of free riders Π^{FR} (see Figure 3). When free riders are grouped together, their reputation value tends to go to zero, since

[1] `http://www.cer.ie`

the system detects that they do not contribute to any demand reduction. For bigger group sizes, and when the number of free riders is relatively small, their reputation does not decrease substantially with time, since the system is less able to distinguish their behaviour from that of legitimate users. For example, when the group size is 8 or 16 and the percentage of free riders is 25, their reputation is maximal,[2] indicating that they are able to perfectly hide inside a group and receive undeserved incentives, thus increasing their reputation. When the percentage of free riders is very high, there are fewer groups of legitimate users to invade, thus the reputation tend to decrease even when the group size is 16.

The ability of the reputation-based mechanism to separate legitimate users from free riders has a direct effect on how well the incentives are distributed among the users. To quantitatively assess this aspect, we computed the incentive distribution error reduction, plotted in Figure 4. When the group size is small, the reputation-based mechanism provides an error reduction in the range of 20% to 50%, depending on the percentage of free riders in the system. This reduction tends to decrease with bigger groups. When the group size is 16, the reduction vanishes when the percentage of free riders is 25 to 50%, and it is severely reduced when the percentage of free riders is 75%.

Finally, one of the objectives of this paper is to assess the trade-off between privacy and incentive distribution error reduction (Figure 5). When the group size G is 2 or 4, the privacy-level of the user is small, but a great error reduction (20% to 50%) is achieved. When the group size is set to 8, there is a significant increase in privacy, while the error reduction is lowered to 0%, 10% and 30% when the percentage of free riders is 25%, 50% and 70% respectively. Finally, when the group reaches size 16, the error reduction disappears completely, although the privacy of the users is very effectively ensured.

6. CONCLUSIONS

In this paper we proposed homomorphic encryption as a privacy enhancing technique for incentive-based DR. We defined a protocol to aggregate the consumption time series of individual users, combined with a reputation-based mechanism that aims at minimising the error committed by the DR provider in the assessment of the individual contribution to the demand reduction. We analysed the trade-off between the correct incentive distribution among users and the enhancement of user privacy. The experimental results showed that for small groups, or when the percentage of free riders is big, the reputation-based mechanism is able to identify free riders and isolate them from legitimate users. When the users are aggregated in bigger groups, the privacy of the users is greatly improved, at the expense of the correct assessment of the DR performance of each user, which leads to misallocations of incentives.

Future work includes analysing different mechanisms for creating groups of users and different reputation scores. Another line of work is improving the reputation-based mechanism to cope with a heterogeneous population of users (i.e., all legitimate users but with different reduction rate γ_i) and/or dynamic changes in the population (i.e., reduction rates γ_i that change over time). Finally, another possible extension is assessing the privacy by actually developing an adversary model to infer the individual consumption of each user from the aggregated consumption value.

[2]In Figure 3, 25% of free riders, the series for $G = 8$ and $G = 16$ are completely overlapping

Acknowledgment

The research leading to these results has received funding from the EU FP7 under grant agreement no. 288322 (Wattalyst) and 288021 (EINS)

7. REFERENCES

[1] G. Acs and C. Castelluccia. I have a DREAM! (DIffeRentially PrivatE smart Metering). *Information Hiding*, 2011.

[2] C. Dwork. Differential Privacy. *33rd International Colloquium on Automata, Languages and Programming*, 4052:1–12, 2006.

[3] C. Efthymiou and G. Kalogridis. Smart Grid Privacy via Anonymization of Smart Metering Data. *2010 IEEE International Conference on Smart Grid Communications*, pages 238–243, 2010.

[4] U. Greveler, B. Justus, and D. Loehr. Multimedia content identification through smart meter power usage profiles. *Computers, Privacy and Data Protection*, 2012.

[5] M. Jawurek, M. Johns, and K. Rieck. Smart metering de-pseudonymization. *Proceedings of the 27th Annual Computer Security Applications Conference on - ACSAC '11*, page 227, 2011.

[6] H.-y. Lin, W.-g. Tzeng, S.-t. Shen, and B.-s. P. Lin. A Practical Smart Metering System Supporting Privacy Preserving Billing and Load Monitoring. pages 544–560, 2012.

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

[8] P. Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *Advances in Cryptology - EUROCRYPT '99*, pages 223–238. 1999.