EPF Lausanne, Aalto University & Technical University of Denmark

Double Degree Programme in Security and Mobile Computing

Georgios Liassas

Privacy Enhancing Mechanisms in the Smart Grid

Master's Thesis
Lausanne, Switzerland, July 30, 2013

Supervisor: Prof. Tuomas Aura, Aalto University
Prof. Christian W. Probst, Technical University of Denmark

Advisors: PostDoc Matteo Vasirani, LSIR, EPF Lausanne
PhD Tri Kurniawan Wijaya, LSIR, EPF Lausanne

EPF Lausanne, Aalto University & Technical University of Denmark

Double Degree Programme in
Security and Mobile Computing

ABSTRACT OF
MASTER'S THESIS

Author:	Georgios Liassas		
Title:	Privacy Enhancing Mechanisms in the Smart Grid		
Date:	July 30, 2013	**Pages:**	70
Professorship:	NordSecMob	**Code:**	T-110
Supervisor:	Prof. Tuomas Aura, Aalto University Prof. Christian W. Probst, Technical University of Denmark		
Advisors:	PostDoc Matteo Vasirani, LSIR, EPF Lausanne PhD Tri Kurniawan Wijaya, LSIR, EPF Lausanne		

The Smart Grid constitutes a hot research topic, nowadays, due to the potential that is has to further improve and optimize the power generation, delivery and consumption. The set of components it is comprised of, such as the smart meters, as well as the advanced communication technologies it incorporates, renders it capable of bringing significant societal benefits and high reliability in reference with its orderly operation.

An interesting technology in the context of the Smart Grid is the Demand Response. This technology attempts to change the way that the electricity customers used to perceive the power consumption by engaging them in an interaction with the energy producer. Essentially, the customers are asked to adapt their power needs based on the state of the power grid. In that way, the energy capacity or resources could be shared more efficiently and unpleasant incidents, such as power outages, could be prevented. In return, the utility company offers monetary incentives, rewarding in this way the customers' power curtailment efforts.

Nevertheless, the fulfillment of the DR goals requires the exchange of information between the utility company and the customers. From the customers' point of view this interaction might be privacy invasive. Consequently, DR programs could not be widely accepted by the public before the privacy concerns are alleviated. This thesis investigates the trade off between the privacy and the efficiency of a DR mechanism by simulating the stakeholders and their interactions in a mock DR environment.

Keywords:	Smart Grid, Demand Response, Privacy, Dynamic Pricing, Incentives, Differential Privacy, Aggregation, Utility, Paillier, Homomorphic, PET
Language:	English

Acknowledgements

I wish to thank Prof. Karl Aberer, head of the Distributed Information Systems Laboratory LSIR - EPFL for giving me the chance to come at EPFL and work on such an interesting topic. Special thanks goes to my advisors at LSIR EPFL, PostDoc Matteo Vasirani and PhD Tri Kurniawan Wijaya. Without their invaluable support and their impactful suggestions the completion of my thesis would not be possible.

In addition, I would like to thank Prof. Tuomas Aura, Aalto University and Prof. Christian W. Probst, Technical University of Denmark for their constructive criticism and remarks. Their comments helped me to improve the quality of my thesis.

Last but not least, I would like to express my gratitude and affection to my family. Their support throughout the two years of my MSc studies was invaluable.

Espoo, July 30, 2013

Georgios Liassas

Abbreviations and Acronyms

SG	Smart Grid
TOU	Time-of-Use
DR	Demand Response
RTP	Real-Time Pricing
DSB	Demand-Side Bidding
TTP	Trusted Third Party
DLC	Direct Load Control
ZKP	Zero-Knowledge Proof
IOU	Investor Owned Utility
CPP	Critical Peak Pricing
EIP	Emergency Incentive Program
PET	Privacy Enhancing Technology
WAMS	Wide Area Monitoring Systems
RTO	Regional Transmission Operator
AMI	Advanced Metering Infrastructure
EMCS	Energy Management Control System
IIDM	Improved Interfaces and Decision Making

Contents

Chapter 1

Introduction

The Smart Grid (SG) constitutes the electrical grid of the 21^{st} century. Many countries around the world have initiated research projects [12], [33], with the development and deployment of this electric delivery system being the primary objectives. The SG is a modern, computerized electric grid that is comprised of advanced data communication technologies and metering capabilities, among others. By 2020, the European Commission requires that the 80% of consumers should have installed smart meters, in order to achieve its SG implementation goals. The smart meters will drastically change the way we, as users, conceive power consumption by integrating a set of technologies that will allow for instant electricity consumption monitoring and track of cost.

Demand Response (DR) technology is an integral part of the SG and has a key role in it. It mainly concerns the bidirectional communication between the users and the utility company aiming to a stable power load and a total demand that matches the production. In other words, DR paradigm attempts to influence the electric power consumption pattern of the users. As a consequence, it manages to handle disturbances and avoid blackouts. In addition, it boosts the quality of supply and it offers high reliability of the power grid.

Nevertheless, the frequent and highly granular communication of power usage data, from the users to the utility company, is considered privacy invasive [23]. Due to this reason, Privacy Enhancing Technologies (PETs) are used to protect users, primarily, against activity and behavioral analysis. On the other hand, the power consumption data, that flows towards the utility, usually undergo a processing phase due to the application of PET schemes. In fact, this processing might deprive the utility from important

information based on which it designs and parameterizes SG services, such as load forecasting. Therefore, as a matter of fact, the trade-off between privacy and SG efficiency constitutes a critical factor towards the wide adoption of the modern electrical grid.

1.1 Problem statement

There is a quite extended literature, [11], [22], [3], proposing privacy-enhancing technologies for the Smart Grid. However, they do not sufficiently study the implications that those mechanisms impose on the orderly operation of the Demand Response paradigm. More specifically, they do not consider the effect of privacy on the DR programs performance. By setting up a simulated DR environment we investigate the challenges that the application of a certain PET class causes. In particular, we investigate the accuracy of a DR incentive allocation mechanism in the context of a homomorphic aggregation protocol.

1.2 Structure of the Thesis

The rest of the thesis is organized as follows. Chapter 2 constitutes an introduction to the world of Smart Grid. Its components and stakeholders are reviewed whereas significant attention is paid on the DR paradigm and the implications it causes on the users' privacy. As for Chapter 3, we present the requirements of a DR system from the perspective of both the users and the utility company. Furthermore, we review the related work on the field by evaluating the effect of some privacy enhancing mechanisms on the efficient operation of DR programs. As far as the Chapter 4 is concerned, it describes the methodology of the experiments as well as its components. Moreover, in this chapter we also analyze the risks that are present in our system and we finally suggest mitigation techniques. Last but not least, in Chapter 5, we present the results and the evaluation of the experiments.

Chapter 2

Background

2.1 The Smart Grid

The Smart Grid refers to a modernized power delivery and monitoring system that intends to substitute the deprecated electrical grid of the 20th century. Technological breakthroughs during the past decades in the field of information and communication technology have facilitated the development of the SG. Fig.2.1 illustrates[1] the evolution of the electrical grid.

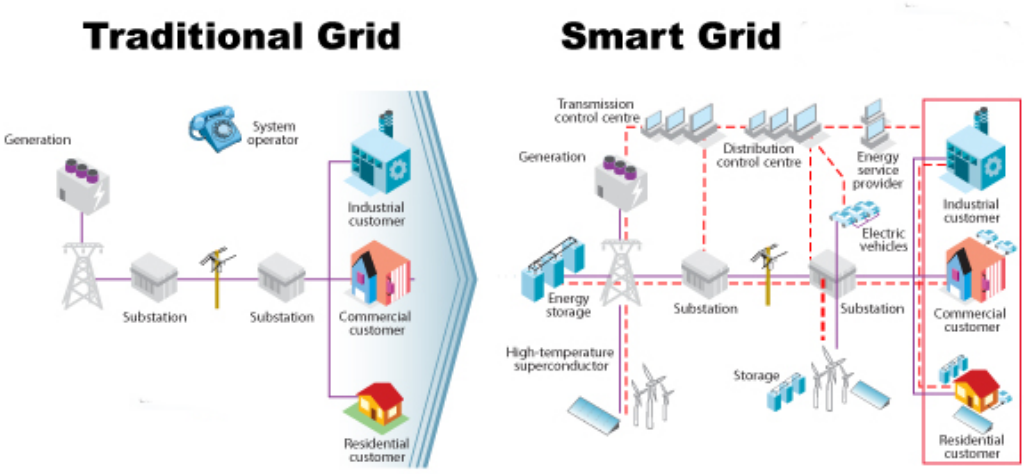

Figure 2.1: The evolution of the power grid.

This sophisticated power grid utilizes computer-based remote control and automation to collect the necessary information that allows serving its purpose more efficiently. The devices that comprise the Smart Grid are equipped

[1] An illustration of the International Energy Agency -IEA

with sensing capabilities and two-way communication technologies. The sensors collect a range of data that is reported to the operational center of the utility. As a consequence, the automation technology that is integrated in the SG fosters the reliability, efficiency and sustainability of the power infrastructure, whereas it offers new benefits to the stakeholders.

2.1.1 Smart Grid Key Components

The SG incorporates a collection of enabling technologies and components that facilitate its orderly operation [2].

Integrated Communications

Integrated communication technologies are a critical component of the SG. All the data produced during the SG operation has to be transmitted from and to several other SG entities. However, different SG components use different communication protocols. In fact, this non-uniformity is an impediment towards an effective, fully-integrated communication infrastructure. Integrated communications is believed to create a dynamic and interactive grid where users and sophisticated devices, such as control centers and smart meters, will interact efficiently.

Sensing and Measurement

Sensing and measurement technologies of the SG primarily focus on the evaluation of the equipment health and the integrity of the infrastructure. They also help mitigate congestions and radically reduce emissions by engaging motivated customers into DR programs. Advanced sensing and measurement technologies include, among others, smart meters, asset condition monitors and wide area monitoring systems(WAMS). In particular, smart meters describe digital meters that record energy usage data and frequently report their measurements both to the users and the utility company. The communication of the power usage information is backed by the advanced metering infrastructure (AMI). In its turn, the AMI is an architecture that automates and facilitates the bidirectional communication between the smart meters and the aforementioned stakeholders.

Advanced Control Methods

Advanced Control Methods aspire to provide the appropriate technologies in terms of hardware and software which will contribute in analyzing, diagnos-

ing and predicting the conditions under which Smart Grid's orderly operation can fail. Moreover, advanced SG devices and algorithms assist in determining the appropriate actions to be taken when alert conditions are identified. The ultimate objective is the avoidance of outages and power quality disturbances. Advanced control functions are supported by distributed intelligent systems (control agents), analytical tools (statistical algorithms) and operational applications such as SCADA and substation automation.

Advanced Electric Components

The modern grid needs advanced electric components to meet the performance requirements in power transmission and distribution. Consequently, they determine the electrical behavior of the grid. Advanced components realization has relied on the significant research and development efforts in the areas of power electronics, superconductivity, materials and chemistry. Examples of such components are the distributed generation and energy storage, the fault current limiters, the advanced switches and conductors, the solid state transformers as well as the microgrids.

Improved Interfaces & Decision Support

Improved interfaces and decision making (IIDM) are essential enabling technologies for the SG. The focus of this Smart Grid component is the transformation of complex power-system data to comprehensible information. It is achieved by virtual reality approaches and other sophisticated data-display techniques. As an immediate consequence, the operators can identify potential problems faster and take the appropriate actions to prevent them. IIDS technologies include, among others, visualization, decision support and system operator training.

2.1.2 Smart Grid Stakeholders

Customers

Customers are the end users of the Smart Grid. They are divided into 3 categories, namely industrial, commercial and residential customers. No matter which category they belong to, they all receive the power supply from the electricity distribution network and based on their energy needs they produce data reflecting their power consumption trends. Thus, hereafter we will also refer to the customers as *data producers*. The energy consumption patterns

of the data producers are used by other SG stakeholders such as utility companies to monitor the state of the grid and ensure its orderly operation.

Utilities

Under the term "utilities" a broad range of bodies is represented. It includes the investor owned utilities (IOU's), the public utilities, the regional transmission operators (RTO's) and the power marketers. These entities supply the Smart Grid with electricity and control the distribution and transportation infrastructure. An important aspect of their operation is the collection and analysis of customers' power consumption data. Hence, hereafter we will also refer to the utilities as *data consumers*. Acting as the Smart Grid facilitators, the utilities take advantage of technological advancements to better control and optimize the grid functions. Ultimately, they aim to provide *advanced quality services* including power load balancing, efficient power generation and distribution, as well as fraud detection.

Policy & Regulation Bodies

Policy makers and regulatory bodies have traditionally exercised their role in the electricity grid by offering supervisory services and enforcing market rules. Their contribution oversight the transparency in the power market and ultimately ensures the public benefit. Smart Grid brings new challenges to this particular class of stakeholders since the way that electricity is traded and distributed has chanced. More stakeholders come into the scene, and the electricity market adapts itself to the new developments. The regulators' role in the SG is twofold. First, they need to update the current regulations to live up to the emerging requirements. Second, they have to act as the intermediary between different stakeholders in order to achieve consensus so that the full potential of the new power grid paradigm can be realized.

Vendors & Technology Providers

Vendors and technology providers refer to a class of companies,organizations and institutions that develop technological solutions and innovative products to support the realization of the Smart Grid. The modern power grid incorporates a number of technological enhancements, among others, in the area of control automation and monitoring, advanced electronics as well as reliable hardware. Essentially, the evolution of technology constitutes the actuator of the SG. With this said, vendors and technology providers have a first class opportunity to actively participate and offer solutions that strengthen the

SG potential and bring financial growth to them.

Other Stakeholders

Advocacy groups, such as environmental organizations as well as governments are, among others, SG stakeholders. The role of advocacy groups is primarily to protect end users' rights in the SG landscape where multiple stakeholders with diverging priorities and interests exist. On the other hand, local or state governments set the initial objectives and financially lead the initiative as part of a national policy.

In conclusion, due to the fundamental need for energy consumption, several stakeholders exist in the modern electrical grid. In spite of their different objectives and expectations the Smart Grid offers benefits to everyone.

2.2 Demand Response Systems

Smart Grid (SG) intends to substitute the traditional electrical grid offering new advancements in favor of both the utilities as well as the customers. Due to the deprecated technology and the lack of appropriate infrastructure, some utilities have no flexibility to deal with severe power incidents other than activating additional power plants or inevitably disconnecting some customers from the distribution network. Such incidents usually include sudden power demand peaks or power plants failure. However, the aforementioned approach to mitigate emergency power conditions is not efficient. Utility companies could face significant financial losses whereas the customers would be frustrated by the low quality of services.

Demand Response paradigm, an integral component of the SG, aspires to provide the utility companies with the appropriate tools to effectively tackle serious power incidents. DR paradigm intends to motivate the customers to play a more substantial role in the operation of the electric grid by shifting or decreasing the power footprint during peak periods. The motivation usually takes the form of time-based tariffs or other monetary rewards. Fig. 2.2 illustrates[2] the operation of a DR system.

Similarly to DR, the Demand Side Management as well as the Price-Responsive Demand are techniques that aim to balance the power demand

[2]An illustration from the Wattalyst project - www.wattalyst.org

Figure 2.2: An overview of a Demand Response system operation

in case of emergencies [32]. In other words, Demand Response actuates the users to adjust their power needs based on the state of the grid, with the goal of achieving fare distribution of available power resources and alleviating peak loads.

2.2.1 Dynamic Pricing

Dynamic pricing is the first type of Demand Response programs. Under this scheme the users are exposed to the ranging cost of electricity production, transmission as well as other auxiliary functions, such as power distribution. The motivation for the customers is clearly the power-bill cost reduction, while at the same time the operator can manage to handle power disturbances, level the peak demand and then mitigate system overloads. This DR approach is realized by prices signaling from the utility to the customers' AMI infrastructure, almost in real time. Various implementations of the Dynamic Pricing scheme provide tariff information to the customers in variable time intervals. Real-Time Pricing (RTP), Time-of-Use (TOU) and Critical Peak Pricing (CPP) are examples of such price-based DR implementations. Nevertheless, dynamic pricing programs are not contract-binding for the customers. Hence, it is not mandatory for them to respond to electricity tariff

fluctuation.

Time-of-Use - TOU

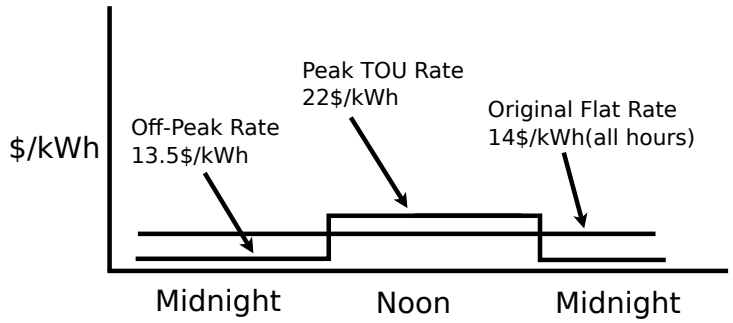

Figure 2.3: Time-of-Use Pricing

TOU schemes divide certain time periods in blocks where the energy cost rates vary, as in Fig[3]. 2.3. A day, a week or even a year can be considered a time period. Usually, a day period is separated in three blocks, in terms of the anticipated power peak load, the low, normal, and peak hours. Each block price rate reflects the average cost of electricity production and delivery. As far as a week time period is concerned, in TOU is typically divided in weekdays and weekends blocks, during which the power consumption trends significantly change. Even in the year time period there are variations. For instance, in the summer time or in public holidays the electricity production cost changes.

The Time-of-Use schemes are considered the second most effective option for realizing a share of the DR potential. This can be achieved with minimum information exchange cost, since the TOU price rates, although time-varying, are fixed [32].

Real-Time Pricing - RTP

In RTP scenarios, as in Fig. 2.4, the DR customers are informed about tariff changes on a-day-ahead or an-hour-ahead basis. In Real-Time Pricing the energy price rates fluctuate hourly reflecting the actual market price. This scheme is the most variable one since the electricity price changes very frequently. It is considered the most appropriate approach for realizing the full

[3]Figures 2.3 to 2.5 are adopted from Fox-Penner (2009), p.41

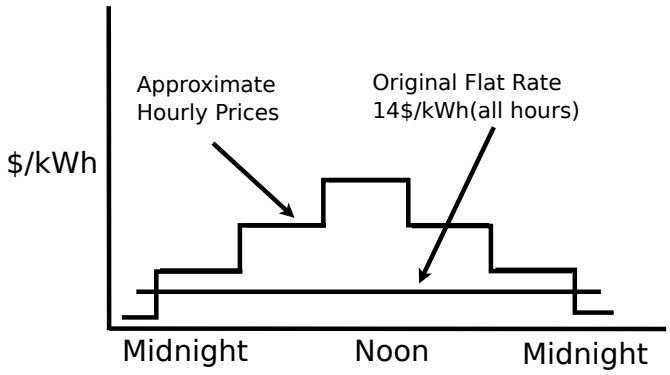

Figure 2.4: Real-Time Pricing

DR potential. Nevertheless, the adoption of RTP programs remains a challenge, especially for residential customers, due to the frequent price updates [32].

Critical Peak Pricing - CPP

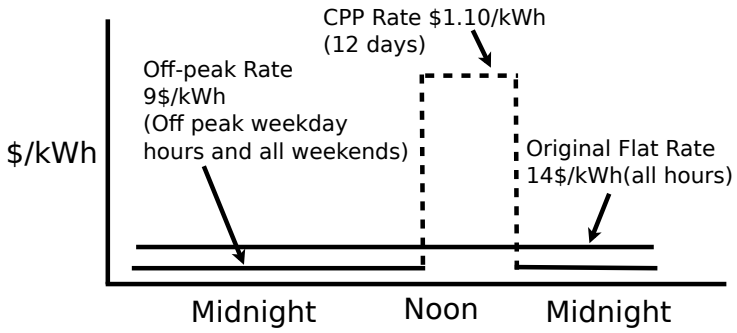

Figure 2.5: Critical Peak Pricing

CPP is a dynamic pricing program where a predefined high rate, as seen in Fig. 2.5, is imposed on the DR participants under critical SG circumstances. Such circumstances include excessive power demand or significant decrease in production capacity. The ultimate objective of a CPP scenario is to drastically reduce the power demand when system reliability is compromised or threatened.

Usually, CPP schemes are used in conjunction with TUO or RTP programs, where utilities inform the customers about the price changes at relatively short intervals. CPP events, typically ,take place few times per year. Moreover, the utility attempts to abet customers' participation by offering energy bill discounts out of the CPP periods. Nevertheless, a basic shortcoming in the broad adoption of such a DR scheme is the way that AMI operates. They are mainly designed to record power on an interval basis where the price is fixed. Sudden electricity rate fluctuation complicates the billing procedure. In addition, another CPP deficiency is the allowed number of peak periods and their duration. If that number has already been reached before a new critical peak incident, the utility will have to activate additional DR programs to alleviate the load [32].

2.2.2 Incentive Programs

In contrast to Dynamic Pricing, Incentive Programs entail an agreement between the service provider and the consumers. Based on the contract that has been signed, the customer is obliged to respond to the utility request for energy reduction. If not, then a penalty is imposed. On the other hand, the power consumer who is engaged in such a power-load shedding scheme enjoys some bill discount benefits. A significant number of incentive-based DR approaches prioritize the power supply safety high in the requirements list. That is the reason why they offer competitive and binding power-consumption packages to the participants. Demand-side bidding (DSB), emergency incentive program (EIP) and direct load control (DLC) are all, among others, different flavors of the incentive programs paradigm.

Direct Load Control - DLC

In DLC programs, the utility directly disconnects electrical appliances or equipment, on short notice, to achieve immediate customers load reduction. All residential, commercial and industrial customers are candidates for such a power curtailment program [1]. The incentive given by the utility, for the possibility of disconnection, is a discount in the electricity bill. Furthermore, the contract between the two parties defines the period of power interruption as well as the groups of appliances whose operation will be suspended. By the deployment of direct load control programs, the utility assures the system balancing and fosters its reliability.

Interruptible/Curtailable (I/C) Services

Under the interruptible service, the contract between the utility and the customer includes curtailment options which are reflected on the power retail rates. These tariffs allow for rate discounts or bill credit for agreeing to lessen the energy demand during critical periods, as in system contingencies. Apart from this, in case a customer does not comply with the contract rules, economic penalties might be imposed by the utility. Last but not least, I/C services have been traditionally offered only to industrial or big commercial customers.

Capacity Market Programs

Capacity market programs are planned on a months-ahead timescale and the utility can activate them on a short notice, usually two or less hours before the program initiation. This DR scheme intends to guarantee the power reserve capacity above the reliability level and to comply with the utility reserve obligations. Incentives include up-front payments proportional to the capacity-market prices.

Ancillary Services Market Programs

Under this program, the customers perform bidding on the load curtailment. They are paid the market price and they agree to immediately reduce electricity consumption when asked. The utility requests load curtailment with less than an hour's notice.

Demand Bidding/Buyback Programs

In this demand side bidding scenario, the utility asks the customers to propose bids along with the amount of power load they intend to curtail. These so called buyback programs are mainly offered to large customers with a significant contribution to the system load [1]. The agreed electricity rate can be either part of the bid or authority-posted, based on the wholesale energy market prices. However, participants who fail to reduce their load are penalized [32].

Emergency Demand Response Programs

Emergency programs are tightly connected to the SG reliability. They are launched when power reserve shortfalls occur. Under such a scenario, the customers are rewarded by incentive payments which are linked to real time wholesale market prices or customer's blackout cost. The time interval be-

tween the emergency notice and the power delivery usually spans from 30 minutes to two hours. This DR scheme is typically enabled when capacity response programs have failed to live up to the current power reserve requirements.

Critical Peak Rebate - CPR

The CPR operates counter to the critical peak pricing scheme. Participants are rewarded for reducing their electricity footprint instead of paying a high rate in case of peak loads. The reward is analogous to the amount of power reduction compared to the predicted consumption during a CPP event. Customers can be only benefit by their participation. This is the primary reason that CPR programs enjoy high acceptance rates by both customers and regulators.

2.2.3 Demand Response Benefits

All Smart Grid stakeholders could enjoy DR programs advantages [34]. The benefits could be grouped into direct and indirect ones, with customers participating in the DR programs enjoying immediate, mainly monetary advantages, whereas the whole society and other consumers could be benefited by the indirect DR effects.

DR schemes participants should notice evident savings in their bills which is highly correlated with their compliance and responsiveness to the particular DR scenario they are enrolled with. Due to the regulated nature of power generation in the context of DR systems, overall electricity price reduction could be expected. Furthermore, reliability of smooth power distribution could be perceived as an immediate aftermath of the DR utilization. Indeed, the DR deployments give the opportunity to the customers to assist in decreasing the outages risk and also avoid their own risk to face a power disruption. Last but not least, the whole market benefits from price reduction and reliability. This is due to the fact that small power demand reduction brings significant decrease in generation cost and consequently in electricity rates.

2.2.4 Demand Response Challenges

Baseline demand as well as response estimation are two quantities that play a significant role in the success of the Demand Response systems. As far as the former is concerned, it describes a demand profile prior its association

with a DR program. Such a demand profile is usually calculated with statistical methods over a set of customers. The accuracy of the baseline demand is critical for the performance measurement of the incentive-based DR programs, since this metric is used for calculating incentive payments for the customers as well as for estimating the load reduction. Two methodologies exist for calculating baseline demand [32]. The contractual approach engages the customer in the process of foreseeing the demand load. However, that is not always feasible and some of the customers are reluctant to participate in such a scenario. The alternative, is led by the utility and is known as the administrative approach. In that case, the energy provider estimates the baseline considering the average energy consumption rate of the previous periods.

Response estimation is another challenge in the DR environment. The participants must respond to DR signals, sent by the utility, during a certain time-frame. Otherwise, they will not serve their purpose. It becomes apparent that the utility needs to assess the number of the DR-enrolled customers who will be available on every control signal dispatch. Furthermore, it also needs to estimate the percentage of those who will finally respond. In that say, the utility can further refine the DR signals to achieve more effective load shedding.

In addition, other challenges are present. The service provider must be able to verify the customers' response so as to be able to evaluate the service quality as well as to remunerate the participants for their contribution. Other than this, the DR programs efficiency heavily depends on the users' engagement to the DR service, since the more reactive they are, the better for the operational state of the system. Finally, one could note that the inclusion of response automation in DR signals would increase the users' participation.

Last but not least, due to the amount of information exchanged between the consumers and the utility, justified privacy concerns arise. Customers would not like to leak personal information and interests through their interaction with the DR service provider. Such conditions are regarded unwelcome by the consumers and consequently hinder DR programs adoption from the electricity market.

2.3 Privacy in the Smart Grid

Smart Grid intends to substitute the traditional electrical grid while offering, at the same time, new opportunities for the utility companies. For instance, load monitoring could assist the operators to foresee future power consumption and to adjust their production and delivery approaches. AMI supports advanced SG operations such as load monitoring, power generation planning as well as demand response. These can be achieved by constantly sending consumption data towards the utility facilities and more specifically to Energy Management Control System (EMCS). Such a control system requires a continuous flow of information in order to perform efficiently. Nevertheless, frequently sent, fine-grained data transmission introduces new challenges that have to be addressed. Privacy related issues are of significant value since there have been proofs that individuals' privacy could be violated [25].

2.3.1 Non-Intrusive Load Monitoring

Non-intrusive load monitoring, usually abbreviated as NILM or NALM , refers to a set techniques that enable the identification of appliances usage in the customers' premises. The information that a NILM process can infer refers to the type of appliances and their state (on or off) associated with the respective time-frame. Much like an AMI records the power consumption, a non-intrusive sensor is needed in order to provide the necessary aggregated data for the identification procedure. The aggregated data corresponds to house-wide or room-wide power consumption information. Even though AMI and NILM sensors are technically similar, the NILM sensors record the power trace in a higher frequency, usually at second or sub-seconds intervals. NILM techniques are characterized as non-intrusive because they eliminate the need for outlet or appliance-level meters or other laborious and intrusive sensors in the household.

Non-intrusive load monitoring is used for a variety of reasons. First and foremost, it assist in analyzing the power consumption patterns and designing techniques to achieve energy demand reduction. Moreover, load forecasting can be supported while NILM algorithms also contribute information for energy saving audits. Collecting load data, designers of appliances can develop more environmental friendly apparatus whereas utilities can detect appliance failure. Last but not least, NILM technology enables demand side load management and contributes in the implementation of incentive programs for particular appliance usage patterns.

Various appliances or classes of appliances have distinct power consumption features which constitute their so-called *power signatures*. NILM mechanisms try to uncover these signatures from the aggregated power information in order to identify the appliances which have contributed to the power consumption. However, several actors in the SG can repurpose power consumption data and the extracted information might be used in ways other than originally intended. G.Q Hart [15] was the first to express such concerns. He claimed that NILM could be used as a surveillance technology. Furthermore, information extracted from load monitoring systems could help organized crime to better plan burglaries and marketers to conduct direct marketing campaigns for the consumers. In conclusion, NILM can undoubtedly assist towards a more efficient SG, but the privacy concerns are justified.

2.3.2 User mode detection

While NILMs identify the apparatus in use along with its schedule, use mode detection attempts to deduce the activity being performed with a particular device. Experiments have shown that TV channels and web browsing recognition is possible with high accuracy.

For instance, Greveler et al. [14] employ 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 analogous to the brightness of the displayed content. They have demonstrated the effectiveness of their method by showing high correlation between the viewed movie and the power consumption. The correlation is proved by a Pearson coefficient with values 0.93/0.94/0.98 for the three movies they have experimented with.

In a different setting, Clark et al. [7] attempt to detect the website that a computer is rendering on the browser from a collection of 8 webpages. To achieve this, they apply direct load monitoring on the computer with power recording frequency of 1 khz. Utilizing a set of classification techniques and coupling them together they managed good accuracy, of almost 60%, with the absence of false positives.

2.3.3 Behavior deduction

The behavior deduction of all classes of customers, namely residential, commercial or industrial, is beneficial for the design of Demand Response programs as well as for the prediction of the electricity demand under certain behavioral conditions. Nevertheless, NILM and activity detection methodologies reveal the appliances schedule and the associated activities that the customers undertake. On a higher level, this information can be used to extract customers' behavioral trends. Utilities can record behavioral patterns and this fact could work as an impediment for the DR adoption because of the privacy concerns it arises for the residential users, in particular.

Lisovich et al. [24] conducted an experiment in a students' residence in order to prove that repurposed energy consumption traces can reveal behavioral patterns and habits. The experiment lasted two weeks. They were constantly collecting electrical data while, for verification purposes, they also installed a video surveillance. After they undergone the behavior extraction module a training phase, they were able to detect load events and predict behaviors. The behaviors were divided in several categories, such as *presence, sleep schedule, meal times*. A degree of disclosure metric was then associated to each of those behaviors. As they claim, their behavior extraction system performed quite well indicating that privacy concerns are justifiable.

2.4 PETs Taxonomy

2.4.1 Anonymization

Anonymization attempts to decouple the consumption data from its *producer*. In that scenario, the data consumer receives some energy usage information which is not attributed to any particular identity. Thus, the data consumer can statistically process the data and perform the calculations it needs, such as inferring electric consumption trends. Nevertheless, these trends can not be associated with any particular customer.

Jawurek et al. [17] show that anonymization techniques are not sufficient to protect customers' privacy, on the condition that an adversary has access to some external indicators. In their paper, they demonstrate two attacks on the unlinkability of smart metering consumption traces. The first attack attempts to associate the identity of a household with pseudonymous

consumption traces via anomaly correlation. They show that the attack is feasible and it also permits deduction of the household behavior. In the second attack they demonstrate that the tracking of a consumption trace origin, hidden behind different pseudonyms, is also possible. The authors achieve this by exploiting patterns in the electricity consumption.

In conclusion, anonymization techniques robustness seem to suffer when an attacker has access to secondary data sources. Several aspects of the data items nature, such as the sampling rate, might give an adversary the opportunity to reveal data producer's identity or distinguish among them.

2.4.2 Trusted Computation

In trusted computation scenarios the data consumer does not have access to individual power consumption information. Instead, it only receives an aggregation of the private data items. The aggregate is computed either by the households themselves or an additional entity, the trusted third party (TTP), is introduced in the protocol to perform the aggregation. Under this setting, the threat model considers the individual power readings as the asset of the system. On the other hand, the disclosure of this individual data to the data consumer constitutes the main threat. The disclosure can be performed by the entities which are in charge of the aggregation. Due to this fact, the different protocols usually need to make strong assumptions in regards to the aggregator trustworthiness. If the assumptions fails, then the privacy guarantees are instantly canceled.

In the trusted computation approaches, we try to deprive the data consumer from accessing individual power readings. By publishing aggregate information, the data consumer can not recover end users' details. However, at the same time the data consumer gets sufficiently accurate aggregate data.

Power data aggregates are mainly either temporal or spatial. *Temporal aggregates* regard the power traces of a singe user over time, whereas *spatial* one regards power traces of multiple users at a certain time interval.

2.4.3 Cryptographic Computation

Under cryptographic computation protocols two variants exist: encryption schemes that rely on the homomorphic property, and protocols that use secret sharing schemes. In both cases, the data consumer can not decrypt

individual data items but only their aggregate.

In homomorphic encryption, the data producers encrypt their individual power consumption items with the public key of the data consumer. The encrypted data then undergoes a homomorphic operation before it is sent to the data consumer. In the case of aggregation, the individual data items usually undergo a homomorphic addition operation. The homomorphic addition is performed by the multiplication of the individual, homomorphically encrypted data items which corresponds to the addition of these values as if they were decrypted. In the end, the data consumer receives the aggregate and gets the decrypted result using its private key.

As far as the secret sharing schemes are concerned, a secret is divided in multiple parts and each part is given to a participant. In order to reconstruct the secret, all or a subset of the participants have to contribute their share. In the context of the SG, a share could be thought of as the private electricity consumption reading of a smart meter. Rottondi et al. experiment with Shamir's secret sharing algorithm, a secret sharing variant, in [28].

2.4.4 Perturbation

Privacy enhancing protocols that use perturbation techniques add appropriate noise to individual data items or to the final aggregate. In this way, the data consumer might still be able to compute the statistics it needs but on the other hand, the privacy of the data producer is preserved.

A special variant of the perturbation schemes is the differential privacy technique. In differential privacy literature the presence of a trusted data aggregator is assumed. The aggregator usually responds to users' queries or it publishes statistics in reference with a population. This technique assures that a user of a particular privacy-preserving statistical database can learn the properties of a population as a whole but it can not infer valuable conclusions on the individuals' properties.

Definition 1. ε-Differential privacy [8]*: A randomized function K gives ε-differential privacy if for all data sets D_1 and D_2 differing on at most one element, and all $S \subseteq Range(K)$, the following holds:*

$$Pr[K(D_1) \in S] \leq e^{\varepsilon} * Pr[K(D_2) \in S] \quad , \tag{2.1}$$

where the probability Pr is taken over the coin tosses of K, and the S is a singleton set, if the output of K is discrete random variable, or a small

range of reals, if the output of K is a continues random variable. Finally, $Range(K)$ is the output range of K.

Every randomized function K that complies with the above definition does not make any output significantly more or less likely in case an individual adds or removes her personal information. In other words, the presence or absence of an individual in the database does not significantly change the distribution of the output. Consequently, this property assures the privacy of the individual information. We note here that we go through a more detailed analysis of the differential privacy against homomorphic aggregation in section 3.3.

2.4.5 Verifiable Computation

In the verifiable computation paradigm, the aggregator, along with the aggregation result, it also provides a proof of the aggregate calculation correctness to the data consumer. Taking advantage of this property, the aggregation protocols of this kind can assume untrusted aggregators, whereas at the same time they can guarantee the integrity of the aggregation result. This particular PET scheme is well suited for protocols that aspire to provide billing capabilities. In fact, this observation holds due to the integrity and accuracy guarantees of the aggregate result that the verifiable computation protocols can provide.

Such protocols are usually developed based on the zero-knowledge proof (ZKP) principles [13]. Briefly, in zero-knowledge proof protocols, two parties, namely the prover and the verifier, interact. The verifier gets knowledgeable of a prover's statement who attempts to prove the validity of that statement without revealing additional information to the verifier, except for the actual statement. In the context of a electricity aggregation protocol, the data producer plays the role of the prover while the data consumer plays the role of the verifier. The data producer, e.g. a household, calculates the total energy consumption on a monthly or bimonthly basis. Then, it informs the the data consumer, e.g. the utility company, about the result. At the same time, the data producer can persuade the data consumer about the validity of the result without disclosing individual smart meter readings.

Chapter 3

Fitting PETs to DR model

A set of PET classes were presented in *Chapter* 2 that provide privacy guarantees in the context of the Smart Grid. Nevertheless, some of them require strong trustworthiness assumptions on third parties, whereas others reveal to little information rendering DR deployments rather infeasible. In this chapter, we initially identify the requirements that a privacy preserving protocol should live up in order enable DR schemes in the Smart Grid. Then, we analyze the advantages and disadvantages of some of the featured solutions.

3.1 Requirements in a DR environment

A privacy-enhancing solution design should meet a set of requirements in order to integrate effectively in a DR context. The effect of those requirements is twofold. First, they should protect consumers' privacy. Second, they should not hinder advanced SG operations, such as load monitoring and demand forecasting.

3.1.1 User-side requirements

The user's main concern is the potential leakage of privacy sensitive data as a result of participating in a DR program with a utility company. Thus, ideally a user would require the following.

Disclosure of consumption data collection

A DR participant should know when the utility company collects consumption data and give his approval for this activity. An even more important aspect of the data collection is the nature of data collected in terms of granularity, location and frequency. Indeed, such characteristics either individually

or collectively might pose severe threats in user's privacy [15]. Moreover, a DR subscriber should be informed about the duration of the data retention and should ideally control the rightful application of the privacy policy under which energy data was collected.

Control of reuse of consumption data

Control of data reuse assures that the data subject, in this case the electricity user, gives or denies his permission to the utility to process the collected information for new purposes. Such purposes include individual data disclosure to third parties or re-purposed data processing that was initially disallowed by the privacy policy between the utility company and the user.

Minimum possible personal data collection

There is a quite extended literature [24], [21], [27] that indicates the privacy implications of smart meter readings collection. Under certain circumstances, such collected data can be coupled with other auxiliary information that comes from other sources. As a consequence, the utility or any other entity that might have access to personal information will be able to derive meaningful conclusions regarding an individuals' habits and behavioral patterns. Thus, the amount of information collected by the utility has to be the minimum possible which in this case is dictated by the functional requirements of DR. Such functional requirements set the least amount of information needed by the utility company to achieve the goals of DR, such as cost savings and load shifting.

Authorized access to consumption data

The energy consumption data collector, in this case the utility company, has to make sure that unauthorized access to individual data is not permitted. That particular requirement concerns the storage of the information after collection and before utilization. The utility company has to keep the consumption readings encrypted and take all the appropriate measures to avoid data breaches.

As of now, the users' requirements presented were tightly connected to privacy. However, the electricity consumers, except for privacy guarantees, also expect power services reliability. In particular, consumers expect low likelihood of blackout incidents owing to inaccurate power demand estimation. As it has been noted in [32], users' acceptance is of high importance for the success of a DR program.

Reliable SG services

Smart Grid operators acknowledge the value of users' high participation rate in DR schemes and they try to keep their interest high by providing benefits. The users mainly expect orderly power flow as well as the financial gain as a result of their enrollment. Potential inability of the utility to meet customers' expectations will increase their discomfort and it will probably determine the success of a DR program.

3.1.2 Operator-side requirements

On the other side of a Demand Response program, the operator also sets a number of requirements that will allow it to run a DR scheme successfully.

Impactful baseline demand prediction

The baseline is the power demand profile that a user has in the absence of any DR scheme. Baseline demand is very important for the utility since in constitutes the metric that is used to calculate the users' power demand reduction. Intuitively, the demand reduction can be thought of as the baseline consumption minus the realized consumption. There are different approaches to calculate the baseline and all of them give varying results in terms of accuracy. This leads to inexact estimation of the DR performance indicators. As an immediate consequence, such inaccuracies affect the effectiveness of the DR program. It becomes apparent that the more accurate power data the utility has the better services it can provide.

Impactful DR response estimation

The DR participants' response estimation determines the DR signal policy that the utility company will follow to achieve the desired power reduction or shift result. More specifically, the utility needs to assess the portion of the DR participants that will respond to certain signal as well as the quantitative aspects of their response. The accuracy of the estimation will significantly affect the DR program performance. The estimation is mainly based on historical data by means of power consumption traces and the context of the power consumers. Once more, accurate historical data will significantly foster better DR response assessment at the cost of privacy.

User DR response verification

The Smart Grid operator needs to know with certainty if a particular DR program participant has actually responded to the DR signal sent. This becomes a necessity because of the contractual obligations between the two entities of the program. The user will request to be remunerated for his DR signal respond, while the utility needs to verify that the user has actually responded. Other than the monetary reasons, DR response verification also assists in measuring the quality of the given service as well as to the refinement of it to better meet the committed service objectives.

Fine-grained power consumption data

As explained in section 2.3, high resolution data poses a significant threat to the users' privacy. On the other hand, fine-grained energy readings gives to the utility a clearer insight on the power profile of each user. In that sense, the utility will be more successful in providing advanced Smart Grid services with certain quality guarantees. This is because of the accurate statistics that it can derive from the volume of consumption data that it has collected for a particular user or a group of users. Last but not least, A.Cárdenas et al. in [4] claim that there are other legitimate reasons for utilities to require fine-grained smart meter data, such as fraud detection and dispute resolution.

Billing capability

The Smart Grid operator needs to be able to charge the customers for their actual power consumption. The billing data needs to be fine-grained for the utility to be able to accurately calculate the energy cost. This is owing to some DR schemes being employed, such as the TOU program. As explained earlier, in TOU paradigm the pricing fluctuates over time. Thus, exact bills can only be derived if the consumption readings are available in high-resolution. However, the billing data needs to be sent to the utility only once per month or bimonthly. This is a basic difference compared to the smart meter readings sent to the operator for power grid monitoring and stabilization reasons. Last but not least, billing data needs to be attributable to a unique household.

3.2 Evaluation of featured PET mechanisms

Under this section we aspire to investigate the implications of featured privacy preserving techniques in the DR environment. We choose characteristic implementations based on the PET taxonomy presented in section 2.4. Then, analyze the privacy guarantees as well as the impact these mechanisms have on a DR program operation.

3.2.1 Escrow-Based Anonymization

Efthymiou et al. in [11] propose an escrow mechanism which plays the role of the mediator between households that participate in a DR program and the utility company that collects the power consumption data. The role of the escrow can be assigned either to the smart meter of each household or to a TTP. Every smart meter is equipped with two distinct IDs, namely the high frequency identifier, *HFID*, and the low frequency identifier, *LFID*. The former is used for communicating high sampling rate energy readings to the power grid operator for demand side management activities, while the latter is used for billing purposes. We note that the HFID is anonymous whereas the LFID is attributed to the respective household. Moreover, the escrow agent is the only entity that knows about the connection of a HFID/LFID pair.

The escrow agent service privacy guarantees are heavily based on two factors. First, the escrow should comply with a strong data policy that will only allow for keeping track of the HFID/LFID pairs. Hence, the escrow agent must be honest and never reveal to third parties the LFID which correspond to some particular HFID. Second, as the authors also point out, the privacy level of this scheme is dependent on the size of the anonymity set, the set of all subjects i.e. households. These somewhat strong assumptions lessen the reliability of the protocol. The escrow service, a single point of privacy failure, has to strictly follow the privacy policy, otherwise no anonymity is guaranteed. Furthermore, Jawurek et al. in [18] introduce de-pseudonymization techniques that further weaken the privacy properties of the escrow mechanism. Their attack vectors assume that the adversary has access to anonymous power consumption traces but, by taking advantage of secondary data sources, they still manage to create linkability. The linkability regards the connection between the power consumption traces and real household identities.

As far as the DR effectiveness is concerned, the escrow mechanism allows for transmission of *exact*, non-aggregated power data in frequent time intervals. It fosters advanced quality power grid services because of the original consumption data that the utility has at its disposal. Due to the utilization of the anonymous HFID, the operator could potentially offer personalized DR signals to individual customers, industrial or residential, further enhancing the DR program accuracy and users' acceptance. Last but not least, all DR schemes, such as TUO or CPP are feasible since no perturbation is imposed on the data sent, thus rendering estimation and verification of the customers' DR response possible. Nevertheless, all these advantages come at the cost of high privacy risk as indicated in the previous paragraph.

3.2.2 Homomorphic Secure Aggregation

Li et al. have proposed in [22] an aggregation protocol over a group of smart meters based on the Paillier cryptosystem [26]. The data consumer, the entity that is interested in making statistical analysis on the aggregate power consumption, publishes aggregation plans towards all participating smart meters which are organized in a spanning tree topology. The aggregation plans indicate the nature of the data to be collected, such as temperature or power and the time interval of the selected data. The plan ultimately controls the aggregation operation. At the root of a the spanning tree there is a *collector*. The collector is in charge of producing the public Paillier key that will be used by all the child nodes of the tree for encryption purposes. This privacy preserving solution utilizes homomorphic encryption to perform in-network aggregation in a bottom-up manner (Fig. 3.1), while it guarantees the privacy of the intermediate results. Each smart meter Paillier-encrypts its consumption traces with the collector's public key and homomorphically adds the result with the encrypted readings of its children in the spanning tree. In the end, the collector homomorphically decrypts the final tree aggregate and communicates the result to the data consumer. We note here that the Paillier cryptosystem is explained in more detail in section 4.3.

The authors have assumed an *honest-but-curious* adversary model. Thus, the smart meters will not tamper with the protocol execution process but at the same time they might try to read other parties' intermediate power aggregation results. As of the latter, 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 spanning tree has the collectors' private key. In the end, only the collector will be able to decrypt the Paillier-encrypted aggregate and transmit it to the data consumer, such

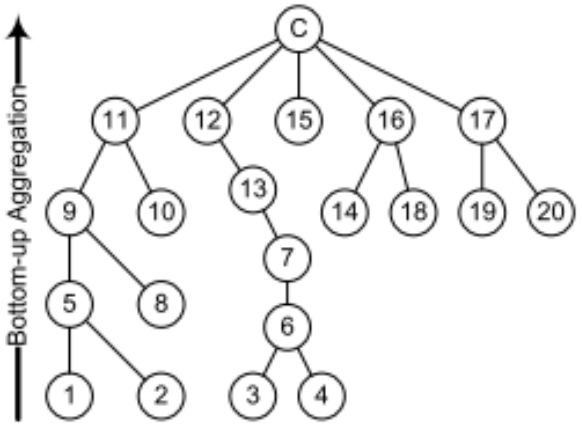

Figure 3.1: The aggregation spanning tree

as the utility. Furthermore, from a privacy perspective the spanning tree aggregate is well secured. The data consumer will not be able to distinguish individual consumption for any of the participants, unless it uses auxiliary information from secondary data sources, as noted in [3]. Conclusively, the users' privacy remains protected behind the electricity usage aggregate over a whole group of smart meters in a building or in a neighborhood. Nevertheless, privacy concerns may arise if the adversary obtains access to other auxiliary information.

Aggregation of power consumption information at multiple levels (building, neighborhood, town) is important for the orderly operation of the power grid [29]. Spatial aggregation is considered less privacy invasive since the information collected does not refer to a certain customer but to a group of customers instead. Although the last observation mitigates the privacy concerns that may arise in the smart grid, it also has an impact on the effectiveness of the DR programs. The aggregates that a data consumer receives characterize the consumption patterns of a number of participants and not individual power consumers anymore. A first observation is that the DR programs can now be offered on a neighborhood, district or town level. The homomorphic secure aggregation suggested by Li et. al calculates exact aggregates. Thus, the data consumer has an accurate picture of the power consumption trends for a particular group of customers. Exact aggregates increase the DR efficiency drastically since reliable statistics can be drawn at a group level. On the other hand, the shortage of more fine-grained knowledge on the energy consumption patterns per customer could downgrade the value of a DR scheme. As a matter of fact, the statistical models that can

be deduced from exact power aggregates describe an average energy user profile. However, different customers in a group have different habits and their response to DR signal might not follow the same patterns. In conclusion, exact power aggregates allow for effective DR schemes on a group level but challenges such as individual customers responsiveness and discomfort impact should be addressed.

3.2.3 Differentially Private Perturbation

Acs et al. [3] propose a privacy enhancing technique that ensures differential privacy [8] by utilizing a novel way for adding appropriate noise to the calculated aggregate. The distributed noise generation mechanism that they use allows for not relying on a TTP to act as an aggregator. Its role can be played by the grid operator or supplier. At the first step, the smart meters, represented as nodes, are grouped in clusters and they create pairwise keys $K_{i,j}$ with Diffie-Hellman key exchange. At a second phase, when the nodes need to report their power measurements to the utility operator or supplier, every node u_i calculates the following value :

$$\hat{X}_t^i = X_t^i + G_1(N, \lambda) - G_2(N, \lambda) \quad , \tag{3.1}$$

where X_t^i is the measurement of node u_i at time t, N is the size of the cluster, whereas $G_1(N, \lambda)$ and $G_2(N, \lambda)$ are random values independently drawn from the same gamma distribution. On the aggregate level, these random values of each node will accumulate to some Laplacian noise which constitutes the guarantee for the differential privacy of the aggregate. This feature is supported by a lemma which states that the Laplace distribution is divisible as the sum of independent and identically distributed random variables following a gamma distribution [3]. Even though this step provides differential privacy for the sum of the measurements of all nodes, the individual measurements are still exposed to privacy threats. The authors propose a third phase in the protocol to address this vulnerability. At the third phase, encryption takes place. Each pair of nodes in a cluster issues a dummy encryption key k using the pairwise key $K_{i,j}$ of the first phase. In the encryption, which is based on modulo addition, the first node of the pair adds k to its X_t^i while the other subtracts k from its own X_t^j. As a result, the aggregator can not decrypt the individual measurements. Nevertheless, at the last step, the aggregator adds all the encrypted measurements. The dummy keys cancel each other out and it retrieves the sum of the noisy measurements. The process that we have described is depicted in Fig.3.2.

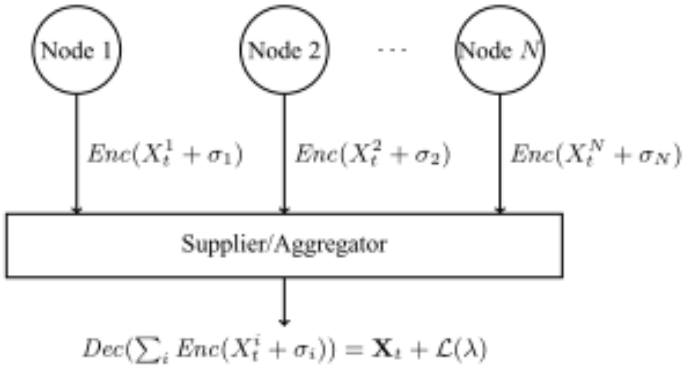

Figure 3.2: Distributed noise generation. $\sum_{i=1}^{N} \sigma_i = L(\lambda)$, where $\sigma_i = G_1(N, \lambda) - G_2(N, \lambda)$ and $L(\lambda)$ the Laplacian noise.

The first important observation in the operation of this protocol is the absence of a TTP. The participants do not need to rely on a separate entity to provide them with privacy assurance. The smart meters can one-by-one communicate their measurements to the power supplier without threatening the privacy of individual households. This is owing to the modulo addition-based encryption of the raw measurements with a pairwise key k between every two nodes of a cluster. Except for individual power readings, the aggregate power consumption of a cluster is also secure in terms of privacy because of the random noise that each node adds according to gamma distributions. In conclusion, the authors assert that their protocol guarantees the differential privacy of both individual smart meter readings and the final cluster aggregate.

As far as the impact of this protocol on Demand Response programs, the authors observe that the bigger the size of the cluster the smaller the error. Due to the differential privacy guarantees, the final cluster aggregate is noisy, hence is not accurate. The differential privacy schemes are flexible enough in terms of privacy-utility trade off. The *global sensitivity* of some function f can be easily calibrated to assure the appropriate balance between privacy and utility. Nevertheless, in differential private scenarios the noisy aggregate is expected to have some impact on the efficiency of DR programs. The perturbed data itself constitutes a challenge for the DR provider since, based on that data, it will calculate the statistics needed for the effectiveness evaluation of the program. For instance, in a demand and supply setting, prediction is of significant importance to match power production with demand. However, the prediction considers past consumption trends.

In case that such trends are not communicated to the DR program designer in an accurate way, the challenge that it faces to produce reliable power statistics becomes apparent. In conclusion, differential privacy techniques is a promising approach towards privacy preservation in the smart grid but careful calibration is needed, since high privacy guarantees could hinder the efficiency the employed DR schemes.

3.3 Homomorphic Aggregation vs Differential Privacy

Both privacy-enhancing techniques attempt to protect the individual information by returning results that concern the whole population that someone examines. We assume here that an adversary can initiate queries to a database where the individual information is stored. In homomorphic aggregation protocols, the aggregation results are usually exact, no matter if the computation is conducted on encrypted data, as we have pointed out in sections 2.4.2 & 2.4.3. As a consequence, the addition or the removal of a participant to or from the population is expected to affect the aggregation result noticeably.

For instance, lets suppose that an adversary wants to learn if a particular student is a member of some school S. In that case, the aggregate information that a query of the adversary would return is the number of the students in the school S. On the other hand, the membership or not, of a particular student, is considered private information. An adversary could learn if a student is a member of a certain school if he makes the two following queries:

1. Find the number of the students in school S
2. Find the number of the students in school S not named "Georgios Liassas"

Hence, it is apparent that exact aggregates can leak individual information. If the result of the first query differ from the second by 1, then the student named "Georgios Liassas" is a member of school S. Moreover, as E. Shi et al. point out [31], the individuals' privacy is not protected if an adversary has access to arbitrary auxiliary information about an individual of the population.

Briefly, the homomorphic exact aggregation can not guarantee the following:

- *It can not promise that one's data will not affect the aggregate result.* Every individual of a population should be expected to affect the aggregate result, otherwise the result would not have any utility. Nevertheless, it would be preferable some individual's data to affect the result in a controlled way, so that it would not disclose, with high probability, her contribution to the result.

- *It can not promise that the attacker will not be able to learn new information about an individual from looking at the result.* For instance, the queries responses might reveal a strong trend over the population. Thus, no matter if an individual is in the database or not, her personal information could be compromised. Moreover, an attacker might be aware of some function of an individual in reference with the aggregate result (e.g. an individual's age is exactly twice the average age of the population). This auxiliary information poses a threat on the individuals' privacy.

An individual would feel safe to submit her personal data to a database if she knew that her contribution would not change much the distribution of the query output. In other words, the query output should not change significantly if an individual decides to share her information. In fact, this is what differential privacy guarantees. It provides a rigorous guarantee about the privacy of personal information. The released response from a database, as a result of some query, gives minimal evidence about the possible contribution of an individual in a data set. In addition, differential privacy remains privacy-robust even in the case where an adversary has access to secondary sources of information, whereas it is independent of the computational power available to an adversary [10], [9].

In a differential privacy setting we have a trusted server which is simulated by a randomized algorithm f. We claim that f is ε-differentially private if for all datasets D_1 and D_2 ,that differ by only one record, the adversary can not guess, with high probability, the contribution of an individual to the result. To achieve this, the theory on differential privacy introduces the notion of *global sensitivity*. Global sensitivity of some function f is defined as the worse case difference or gap, in query responses, that is caused by adding or removing an individual's data from a dataset. We note that the

sensitivity is a property of the function f alone and it is not related to any database. Function $f(D) = X$ is a deterministic, non-privatized function over a dataset D which returns a vector X of k results.

Definition 2. Global Sensitivity [8]: *For $f : D \rightarrow R^k$, the sensitivity of f is :*

$$\Delta f = \max_{D_1, D_2} \parallel f(D_1) - f(D_2) \parallel_1 \quad , \tag{3.2}$$

for all D_1 and D_2 differing in at most one element.

For a release mechanism K, in order to provide the privacy guarantees and to be considered differential private, it needs to bridge the global sensitivity gap. The sensitivity gap bridging is achieved by adding noise. Random noise values are taken from a Laplacian distribution with standard deviation large enough to cover the sensitivity gap. Laplacian noise is not the only method to achieve differential privacy, though it considered the easiest one.

Hence, a differentially private release mechanism K will responds to queries with:

$$f(D) + (Lap(\Delta f/\varepsilon))^k \quad , \tag{3.3}$$

where $f(D)$ is the original answer and $Lap(\Delta f/\varepsilon)$ is the distribution of the noise added to the response. More specifically, the term $Lap(\Delta f/\varepsilon)$ describes the scaled symmetric exponential distribution with standard deviation $\sqrt{2}\Delta f/\varepsilon$. The probability density function is:

$$p(x) = \frac{e^{(-|x|/b)}}{2b} \tag{3.4}$$

and the cumulative distribution function:

$$D(x) = (1/2)(1 + sgn(x)(1 - e^{(|x|/b)})) \quad , \tag{3.5}$$

where $b = \Delta f/\varepsilon$.

We note that the parameter ε is a publicly known parameter which controls the trade-off between the accuracy of the computation (utility) and the robustness of the privacy guarantees (privacy). Lower ε indicates higher privacy but lower utility, and vice versa.

With the following proof [9] we aspire to show why a release mechanism K satisfies differential privacy by covering the sensitivity gap of a function f with Laplacian noise.

Proof. Lets assume a non-privatized function $f(D) = X$, with $k = 1$. At any $R \in Range(K)$, considering the formula 2.1, we have:

$$e^{\varepsilon} \geq \frac{Pr[K(D_1) \in S]}{Pr[K(D_2) \in S]}$$

$$\geq \frac{\frac{\varepsilon}{2\Delta f} * e^{-\frac{(|R-f(D_1)|\varepsilon)}{\Delta f}}}{\frac{\varepsilon}{2\Delta f} * e^{-\frac{(|R-f(D_2)|\varepsilon)}{\Delta f}}} \qquad \text{(probability density function)}$$

$$\geq \frac{e^{-\frac{(|R-f(D_1)|\varepsilon)}{\Delta f}}}{e^{-\frac{(|R-f(D_2)|\varepsilon)}{\Delta f}}} \qquad \text{(simplifying the fraction)}$$

$$\geq e^{-\frac{(|R-f(D_1)|\varepsilon)}{\Delta f} + \frac{(|R-f(D_2)|\varepsilon)}{\Delta f}}$$

$$\geq e^{\frac{\varepsilon}{\Delta f} * |f(D_1) - f(D_2)|}$$

$$\geq e^{\varepsilon * A} \qquad \text{(with A} \leq 1\text{)} \qquad \square$$

Chapter 4

Simulation Methodology

As mentioned in the previous sections, Demand Response programs have a great potential in the modern electrical grid. They can help handling power disturbances, mitigate blackouts and adapt load curves, among others. Nevertheless, at the same time, they also pose privacy risks on the users. Recent research efforts [19] have indicated that power consumption data can be repurposed by the utility companies and used for disclosing users' daily activities. Consequently, users' privacy protection is of significant importance for the wide adoption of DR programs. Decreasing users' discomfort, originating from privacy threats, is a key factor for the successful deployment of the Smart Grid. Multiple PETs have been proposed to equip DR solutions with privacy guarantees. Interested readers can refer to section 2.4, where a taxonomy of such privacy enhancing mechanisms is proposed.

4.1 Simulation Overview

This section intends to give an overview of the simulation, explain the purposes it serves and to define the objectives that we try to achieve via this process.

The Motivation

Although PETs mitigate privacy leakages caused by energy consumption data, they also cause other side-effects that could potentially hinder the utility of the Smart Grid advanced services. For instance, power aggregates on a neighborhood or building level might give an exact picture of the energy consumption but on the same time they diminish the consumption patterns accuracy per individual user. Furthermore, the utility has to address the

challenging problem of accountability. In a DR context where power aggregates are used, some users might underperform but still receive a notable reward.

We have considered the challenges that emerge owing to the use of power aggregates to protect users' privacy. We have simulated a DR system where users are grouped with variable group sizes. The main objective of the simulation is the evaluation of an incentive allocation mechanism used in the context of a DR scenario. We have assumed that the utility has an accurate baseline consumption for each household participating in the DR program. The baseline consumption indicates the electricity usage of each household before the initiation of any DR scenario. CER smart metering data has been used as the baseline [5].

We have designed a simple DR mechanism where the utility sets a monetary reward as an incentive for the participating households to take part in the DR program. As far as the users are concerned, they undertake the responsibility to reduce their energy consumption right after the reception of a DR signal from the utility. However, not all of the users in the groups respond the same dutiful way. Such a scenario introduces implications in the accuracy of an incentive allocation mechanism and affects the fairness of the process. We investigate those implications, try to identify the causes and mitigate the problem.

The Evaluation

The evaluation is performed based on three performance indicators, namely the homogeneity, the incentive allocation error and the privacy. As far as the homogeneity is concerned, it indicates how effectively the utility can distinguish between users who constantly respond to DR signals and those who cheat trying to hide their deed behind group's overall power curtailment. In regards to the incentive allocation error, it basically calculates the error in allocating the deserved incentive to the DR users who are grouped. Last but not least, privacy quantifies the uncertainty of an adversary, the utility company in particular, in reference with the correlation of a user with an observed action. In the context of the simulation the observed action, from the utility point of view, is the power aggregate of a group of users after the dispatch of a DR signal.

The Simulation Data

The simulation is trace-driven, meaning that the power usage information, that has been used in the experiments, is originated from a real system [5]. In 2012, Ireland's Commission for Energy Regulation (CER) published some smart meter electricity trial data. The data regards smart meter power consumption readings for 782 households on a basis of 30 minutes interval. The power readings span from January 1^{st} 2009 till December 31^{st} 2010. Every data point consists of three values. The first one denotes the smart meter ID. The second is a five digit code where the three first represent the day while the two last digits determine the time during a certain day. As for the last value, it determines the electricity consumed during 30 minutes interval, in kW.

4.2 System description

4.2.1 System Components

The Utility (U)

The utility company is the entity of the system which is interested in the statistical processing of the households' consumption data after the completion of a DR event. In the context of our simulation, a DR event is initiated with the dispatch of a DR signal towards the participating households and it is completed after 3 hours. More specifically, the utility sends to the households DR signals every day for a certain number of days. The DR event lasts 3 hours from 18:00 in the evening to 21:00 in the night. During that period the households taking part in the DR program have to shed their energy load. As a consequence, the utility will reward the household with a monetary incentive analogous to their power reduction effort.

The Household (H)

The households are those entities of the system that give feedback to the utility in reference with their power curtailment after the completion of a DR event. In order to alleviate their privacy concerns, the utility allows those households to be organized in groups and send to it their power aggregates instead of individual smart meter readings. However, the role of the group organizer is assigned to the utility. The size of the groups is common for all of them but variable. For instance, the utility can choose a number of 2, 4 or 36 households per group. Then all groups should be of the same size. In that setting, every household power consumption is encapsulated in

the respective group power consumption. As an immediate aftermath, some households might not comply with their obligation to reduce their power load as expected, believing that they can not be found accountable for their improper behavior. We have recognized this situation and we distinguish the households in two categories, the free-riders and the legitimate users. We suppose that in the former category the households do not reduce their energy footprint at all, while in the latter they reduce by some reduction rate $\gamma \neq 0$.

The Collector (C)

Between the utility and the groups of households we introduce the last entity of the system, the Collector. Every group is assigned a collector. The collector is in charge of setting a secure environment for the households to report their smart meter readings. Moreover, it also undertakes the task of forwarding the group's power aggregate to the utility. From the privacy point of view, the collector is the weakest link in the system we have described. In section 4.3 we elaborate on the privacy risks and we introduce a threat model of the system.

4.2.2 System Functions

Each of the system components, namely the utility, the households and the collector is equipped with a set of functions in order to serve its role in the context of the DR program. The following paragraphs describe each component functions in detail.

The Utility Functions

The utility has a twofold role. First, it is responsible to distinguish the free-riders from the legitimate users and isolate them in order to find themselves accountable for their behavior. Second and more importantly, the utility has to ensure high accuracy in the allocation of the incentive. Nevertheless, both tasks are challenging because of the privacy mechanism we employ. By reporting the group's aggregate power consumption to the utility company instead of individual smart meter readings, the utility lose valuable information granularity. To overcome these difficulties we have developed three different functions for the utility and we evaluate them in Chapter 5.

Random. The random function constitutes the simplest method we have developed and it is used as our baseline to compare the effectiveness of the other two functions we present later. The utility starts by splitting 512 out

of 782 households to groups of some group size. A subset of the households is used in order to help us form the groups of different sizes. As we have noted before, the DR program lasts for 535 days based on the available CER data. After the initial creation of the groups, the utility randomly shuffles the households in order to form the groups of the next DR round. At each round it allocates the incentive to the participating households as it is thoroughly explained in Section 4.2.3.

Smart. The main problem of the random function is that we can not control at any stage of the process the grouping of the households. More specifically, some legitimate users might be grouped with some free-riders and thus the group members will not receive a fair share of the incentive. Ideally, we want to give the chance to the legitimate users to be grouped together and not be mixed with the free-riders.

The smart function attempts to improve the incentive allocation of the next round by considering the amount of incentive of the previous DR round, for each of the households. Prior to the incentive allocation of the current DR round, the utility sorts the households in decreasing order based on the amount of incentive they received by the utility in the previous round. Here the intuition is that those households that have been ranked high in the list are probably legitimate, while those lower in the list are probably the free-riders. However, as we have explained previously, legitimate users might also be grouped with the free-riders. As a consequence, those households' ranking will be affected by free-riders infringing behavior. The utility needs a way to give the chance to those underprivileged legitimate users to be grouped with other legitimate users and finally receive the incentive they deserve. Furthermore, the utility wishes to discover the free-riders for accountability reasons.

Considering the aforementioned utility objectives we have designed a smarter algorithm to form groups. The smart function does not randomly pick group members among all households but from a certain range of households with some high probability and out of this range with some lower probability. In more detail, after the utility has sorted the list of households based on the incentive allocated from the previous DR round, it starts forming the new groups for the next round. For every new group it randomly picks a household from the sorted list. We refer to this first household as the indexed household. Then, it completes the formation of the group using a radius. The radius defines the neighboring households of the indexed one. We set the radius length to be equal to the half of the group size. While the group formation has not been completed, the utility picks one new household

a time from a certain range with high probability. The range is defined as the set of households belonging in the $[-radius, +radius]$ interval. Since we have initially assumed that the neighbors of the indexed household belong to the same category (legitimate or free-riders) we pick from inside the range new group members with relatively high probability. In that manner we attempt to keep legitimate and free-riders apart. To give the chance to the legitimate households, which have been ranked low in the sorted list, to be grouped with other legitimate, we pick outside from the range new group members with relatively low probability. After the completion of a certain group, a new indexed household is picked and the group formation continues as described above, until all groups are created. Finally, the incentive allocation mechanism run again and the incentives it distributes serve as the basis for the new DR round.

Fig. 4.1 illustrates a snapshot of the Smart algorithm state. There are 6 houses which must be grouped in groups of 2. There is one free-rider, the household numbered 2 and five legitimate. The total incentive is €100 which is split among the users. The households have been sorted based on the amount of the incentive received from the previous DR round. In the current round, the index that points to the household 1 is randomly picked and the radius has half the length of the size of the groups, thus the length of 1. The households 5, 1 and 3 belong inside the range, while the households 6,2 and 4 belong outside that range. The algorithm will choose the other member(s) of the group inside or outside the range with some probability, as explained earlier. We should note that the household which are already picked to participate in a group can not be used again.

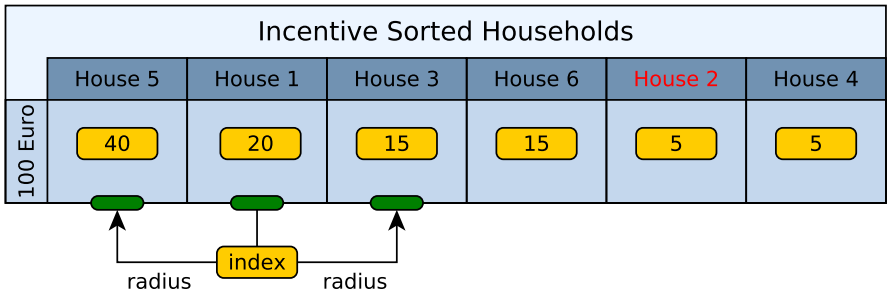

Figure 4.1: Group formation Smart algorithm for group size : 2

Mixed. Even though the Smart algorithm is characterized by a more structured and controlled way of shuffling the households and forming the

groups, it also shows some deficiencies. First and foremost, the utility might manage after some DR rounds to sufficiently separate the legitimate users from the free-riders. However, it has no indication that such a state has been reached. Hence, it might continue shuffling the households after some relatively optimal state has been reached. Such a tactic could once more pair the legitimate users with the free-riders and cancel all the effort that has been made, up to that point. As far as the incentive allocation mechanism is concerned, so far the utility does not exploit any historical data. To integrate the past knowledge in the allocation mechanism we have introduced a reputation metric. Every household participating in the DR program has its own reputation. The reputation is simply defined as the accumulated incentive for each household over the past DR rounds. Here, the intuition is simple. Due to the fact that we have assumed a consistent behavior for the users, once they are characterized legitimate or free-riders, they act analogously throughout the entire DR program lifecycle. Consequently, we expect that after every DR round the reputation metric will assist to better distinguish between the two users' category.

In the mixed algorithm we have two different phases which are interleaved. During the first phase the groups are formed based on the Smart algorithm presented above. The second phase sorts and groups the households based on their reputation. Those two phases run alternately throughout the DR program duration. The first phase allows for a convenient and purposeful shuffling of the DR participants. Thus legitimate users are grouped with legitimate, free-riders with free-riders and legitimate with free-riders as well. In the second phase we come to exploit the aftermath of shuffling of the first phase. The reputation tends to get higher for the legitimate whereas lower for the free-riders, in the long run. The more DR rounds the better for this disunion. The second phase sorts the households based on their reputation and forms the group picking households sequentially. Consequently, the legitimate are grouped with the legitimate and free-riders with free-riders. Last but not least, now the utility does not need to have an indicator when to stop shuffling the groups' households. It can simply take advantage of the fact that the separation between the two users' category improves as the DR program evolves.

The Household Functions

The household simulates the behavior of a user who is participating in a DR program. For each of the 512 households that the system randomly picks to participate in the DR program, we associate their consumption data with

the respective household and then we calculate the realized energy based on its γ, the reduction rate. We remind here that a free-riders has $\gamma == 0$ while the legitimate users have $\gamma > 0$. The formula 4.1 describes the computation of the amount of electricity that it is consumed by a certain household based on its γ.

$$Realized_h = (1 - \gamma) * \sum_t Power(t) \quad , \tag{4.1}$$

where $Power(t)$ indicates the power consumption at some time point t. On the other hand, formula 4.2 shows the amount of electricity reduction for each household. Obviously, the reduction of a free-rider equals zero.

$$Reduced_h = \gamma * \sum_t Power(t) \tag{4.2}$$

During the run of the incentive allocation, the household will forward the $Realized_h$ value to the utility encrypted via the collector, as it is thoroughly described in Section 4.2.3.

The Collector Functions

The collector undertakes all the operations of the group, as a group leader. It is in charge of calculating and dispatching the group power aggregate to the utility and allocate the incentive to the group members. As far as the aggregate calculation is concerned, it is explained in detail in section 4.3. In regards to the incentive allocation mechanism, the collector uses two methods which are described right below.

Naive Allocation. In the naive allocation, the collector does not use any historical information to estimate the effort of each group member. Thus, the group incentive (I_g) is split equally among the households of that group as in formula 4.3.

$$I_h = \frac{I_g}{S_g} \quad , \tag{4.3}$$

where I_h is the incentive of each household of the group g and S_g is the number of households participating in that group.

Advanced Allocation. Obviously, the naive allocation mechanism is not very efficient since it assumes that all the participants reduce the same

amount of energy. This is not true since we know that there are free-riders who are not reducing at all. To tackle this problem, we introduce a new method which takes advantage of the household's reputation. We assume that the higher the reputation of a certain participant, the higher the incentive it should receive. This idea is depicted in the formula 4.4

$$I_h = \frac{P_h}{P_g} * I_g \quad , \tag{4.4}$$

where P_h is the household's reputation and P_g is the group's reputation. The group's reputation is calculated as the sum of all households' reputation belonging in that group.

4.2.3 Incentive allocation

The incentive allocation mechanism is the most important function of our system. Its effectiveness has a significant impact on the orderly operation of the DR program. Users' discomfort from possible incentive misallocation could threaten its success. Moreover, it could deprive the utility, the households and the whole society from such a program benefits as described in Section 2.2.3. In the rest of this section we elaborate on the incentive allocation process.

At time point t_0 the utility sends to all the households a DR signal. The DR signal is to be applied by the households at the time point t_1 and will last till the time point t_3. Indeed, we have simulated the reception of the DR signal from the households at 18:00 in the evening while the DR event terminates at 21:00 in the night. At some later time point t_r the households are asked to report to the utility their power consumption recorded during the DR event.

As we have seen earlier, each household calculates the amount of power that it has consumed from t_1 to t_3 based on the formula 4.1. At t_r the utility asks all the collectors to report the group's total energy consumption. This is the initiation of the incentive allocation. In their turn, the collectors ask from all the households belonging to that group to report their individual electricity consumption measurements. Every collector computes the group power aggregate as explained in section 4.3, and reports back to the utility.

At this point the utility has collected the power aggregate information that it has been realized from all the groups. As we have mentioned before, we

have assumed that the utility company has a perfect electricity consumption baseline. In fact, a perfect baseline constitutes an accurate prediction of the amount of energy every household would consume in case none DR program was deployed. Having the perfect baseline, the utility can now calculate the amount of power that every group has reduced, as in the formula 4.5.

$$Reduced_g = Baseline_g - Realized_g \quad , \tag{4.5}$$

where $Baseline_g$ is the amount of energy that the group g would consume from t_1 to t_3. The utility then calculates the sum of these values as depicted in formula 4.6.

$$Reduced_T = \sum_g Reduced_g \quad , \tag{4.6}$$

where $Reduced_T$ is the total amount of energy that it has been curtailed due to the deployment of the DR program for a single run. The last function of the utility in reference with the incentive allocation is the distribution of the incentive based on the effort of each group. For simplicity, we have assumed that the utility incentive is constant and does not chance overtime. The formula 4.7 shows the way that the utility splits the incentive among the groups.

$$I_g = \frac{Reduced_g}{Reduced_T} * I_U \quad , \tag{4.7}$$

where I_g is the group incentive and I_U is the total incentive that the utility intends to distribute.

The incentive allocation terminates by distributing the incentive of each group (I_g) among the households of the group. The collector is the entity in the system that has undertaken this responsibility. In the ideal case, the group incentive should be analogous to each household's effort. Nevertheless, for privacy reasons the collector should never get access to individual power measurements of the households. To overcome this obstacle, we have introduced two methods, namely the naive and the advanced allocation which are already described in Section 4.2.2.

4.3 Threat Model

4.3.1 Assumptions

We start the threat analysis by making some fundamental assumptions. In a real world deployment all the system components need to communicate via a bi-directional communication channel. We assume that this channel is secure with respect to confidentiality, integrity and availability. We also assume that every component is secure enough to protect the data it stores or creates. In more detail, as far as the household's smart meter is concerned, we assume that it utilizes a trusted platform module (TPM). The TPM assures the integrity of power measurements as well as secure cryptographic operations that any household needs to conduct while the DR protocol is running. The same holds for the collector. We assume that the collector is also equipped with a TPM which is used for reliable cryptographic operations. Last but not least, we have assumed that the utility is capable to protect group power measurements it receives from the collectors.

4.3.2 Risks Analysis

The main threat we address in this study is the disclosure of individual power measurements to an adversary. The privacy of the users participating in the DR program is of significant importance to us. That boils down to the protection of the households' smart meter readings. In our DR scenario, the collector and the utility are the adversaries. In case we reveal fine-grained, individual power measurements to the utility, the users' behavioral patterns are exposed and this constitutes an apparent privacy risk [20].

To alleviate these concerns we have introduced the collector who is in charge of aggregating the power readings of the households that belong to the collector's group. In that sense, the utility only receives an aggregate and consequently can not infer with certainty the individual power footprints of the users. In addition, we have assumed a honest-but-curious adversary model for the collector. Essentially, the collector follows the protocol but it might try to infer electricity usage information related to some household taking advantage of the messages that are routed through it.

In order to protect against such a risk we utilize the Paillier cryptosystem. Its homomorphic properties allows for the protection of individual smart meter readings without affecting the accuracy of the final aggregate. The Figure 4.2 illustrates the aggregation tree for 10 households which are organized

in groups of two under each collector. The Paillier cryptosystem uses a public key (K_{pu}) and a private key (K_{pr}).

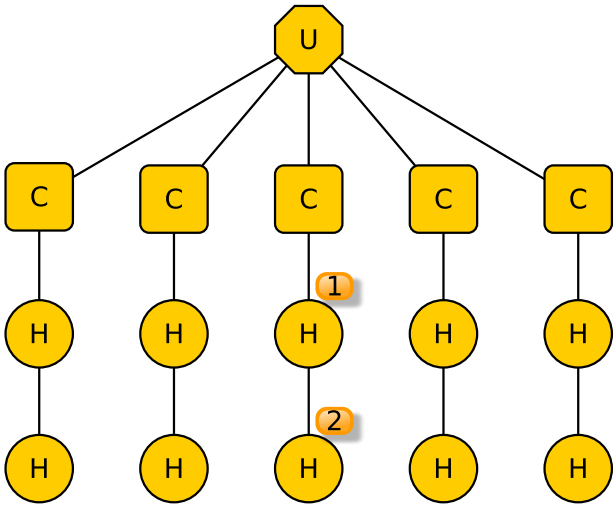

Figure 4.2: 10 households grouped by two under each collector

The K_{pu} is used by the households to encrypt their individual smart meter power readings. On the contrary, the K_{pr} is used from the collector to decrypt the aggregate of the individual energy measurements. In our system the collector is in charge of creating those keys and distributing the K_{pu} to the households it controls. We demonstrate how the collector calculates the power aggregate by the following concrete example.

The Homomorphic Aggregation Protocol

Let us suppose that participants 1 and 2 in the aggregation tree 4.2 want to report their power readings to their *collector* C. We name their measurements m_1 and m_2 respectively. Both participants encrypt their data with a public key K_{pu} that their collector has provided. Hence, $C_1 = E_{K_{pu}}(m_1)$ and $C_2 = E_{K_{pu}}(m_2)$. Now, in case every household attempts to report directly to the collector, its power measurement privacy will be compromised. This is owing to the fact that the collector owns the K_{pr} and consequently it can decrypt any smart meter reading which is encrypted with its K_{pu}.

To address the aforementioned issue, the household which is in charge of reporting to the collector (household 1) will not report anything unless it has received and Paillier-add the power reading of its adjacent household

(household 2). Thus, the aggregation protocol works in a bottom up manner. The household at the bottom of the aggregation tree transmits its encrypted power measurement to the household above. In its turn, the household above will retrieve its smart meter power measurement, it will encrypt it with the K_{pu} and it will Paillier add it with the encrypted measurement it received from the household at the bottom. This procedure will repeatedly happen until the protocol reaches the household which is in charge of reporting to the collector.

In the Paillier cryptosystem the additions in plaintext are translated to multiplication in the ciphertext. Thus, in our example, the household 1 before it reports to the collector will perform the following computation : $C_g = C_1 \times C_2$, where C_g is the group power consumption encrypted aggregate. Once the collector receives the aggregate C_g, it decrypts it with its private key K_{pr} and it retrieves the exact group aggregate, as shown in formula 4.8.

$$D_{K_{pr}}(C_g) = m_1 + m_2 \tag{4.8}$$

Paillier Homomorphic Encryption

Under this section we aspire to provide a concrete example of the way that Paillier cryptosystem works based on the aggregation tree illustrated in Figure 4.2.

- The collector of each group generates a public key K_{pu} and a private key K_{pr}

    1. the collector selects a number $k = 32$, the length in bits of the RSA modulus n
    2. the collector selects two random and distinct prime numbers $p = 59351$, $q = 55219$ of $k/2$ length each
    3. it computes the the RSA modulus $n = p * q = 3277302869$
    4. it computes $\lambda = lcm(p-1, q-1) = (p-1)*(q-1)/gcd(p-1, q-1) = 1638594150$
    5. the collector proceeds by selecting a random integer $g = 8943306254069481040$ in $Z*_{n^2}$, a set of integers coprime to n
    6. it computes $\mu = (L(g^{\lambda} mod n^2))^{-1} mod n = 1754157928$, where $L(u) = (u-1)/n$

7. finally the K_{pu}:(n, g) and the K_{pr}:(λ, μ)

- Participants' 1 and 2 current power consumption is $m_1 = 2$ and $m_2 = 3$ respectively
  1. participants 1 and 2 select a random integer r_1=1904648907 and r_2=1035629130 in $Z*_n$, one for every measurement they need to encrypt
  2. they compute their ciphertext as $C = E_{K_{pu}}(m) = g^m * r^n mod n^2$, thus $C_1 = E_{K_{pu}}(m_1) = 4878868962385258562$ and $C_2 = E_{K_{pu}}(m_2) = 4624922822985571729$
- Household 2 sends to household 1 its encrypted smart meter reading C_2. Household 1 will then homomorphically add this value to its own encrypted measurement.
  1. Ciphertext sum : $C_1 * C_2 mod n^2 = 2778590782834299795$
- Household 2 then sends to the collector the sum of ciphertexts. The collector decrypts the ciphertext sum and sends the aggregate to the utility
  1. $D_{K_{pr}}(C_1 * C_2 mod n^2) = 5$

Following the example of the power aggregate computation we can verify that at every step the privacy of the individual measurements is guaranteed. While the encrypted power readings are forwarded from one household to the other they remain confidential due to the fact that none of the houses holds the private key K_{pr}. Last but not least, individual smart meter readings are also protected from the honest but curious collectors. Even though every collector possesses the private key K_{pr} of its group, it only receives the aggregate power consumption of the group, encrypted. By decrypting this chiphertext it only gets access to the exact group aggregate. Hence the collector is not able, in the first place, to distinguish among atomic electricity consumption patterns.

Chapter 5

Experimental Results

The experiments aim to evaluate the effectiveness of the incentive allocation system. More specifically, the experiments differ on the size of the groups formed whereas they share other common experiment characteristics such as the total incentive offered by the utility company as well as the reduction rate γ. More details on the parameters of the experiments are presented later in this chapter.

The effectiveness of the proposed solution is measured in terms of privacy and incentive allocation error. Due to the presence of the free-riders in the population of the DR participants, we introduce a new metric, the homogeneity. This metric indicates how good our household shuffling methods perform in separating the legitimate users from the free-riders. Based on the group size, we calculate the homogeneity of the population for each DR day and we try to correlate the incentive allocation accuracy with it.

5.1 Performance Indicators

In order to calculate the incentive allocation error we first define the deserved incentive of a household. The deserved incentive is computed based on the formula 5.1.

$$D_h = \frac{Reduced_h}{Reduced_g} * I_g \quad , \tag{5.1}$$

where $Reduced_h$ is the amount of energy that the household h reduced because of the DR signal. Once we know the incentive that the household deserved because of its effort to reduce its power consumption by some amount,

we calculate the incentive error as shown in formula 5.2.

$$E_h = I_h - D_h \quad , \tag{5.2}$$

where I_h is the incentive the household h received from the incentive allocation mechanism. At this point, we note that the approach we use to calculate the error does not change, no matter which of the two incentive allocation mechanisms (naive or advanced) is used. It can be seen that the E_h can be either a positive or negative. It is positive when the utility overestimates the effort of the households to curtail their power footprint, whereas it is negative when it underestimates. We note here that the total allocation error that we present later in this section regards the total absolute error.

As far as the homogeneity metric is concerned, it shows the percentage of groups in the population which are free of free-riders. In order to calculate the homogeneity per DR day, we first compute the maximum number of groups needed to keep all the free-riders that lurk in the system.

$$N_{max} = \frac{N_f}{S_g} \quad , \tag{5.3}$$

where N_f is the total number of the free-riders in the system and S_g is the size of the group. Formula 5.3 presents the way we calculate the number of groups we need in order to isolate the free-riders. Once we know the N_{max}, we then compute the number of groups in the system which actually consist solely from free-riders after we have shuffled the households in the groups. We name this value N_{act}. The homogeneity is given by the formula 5.4.

$$M = \frac{N_{act}}{N_{max}} * 100 \tag{5.4}$$

We have used the Shannon entropy [30] as the privacy metric of the system. The entropy formula is described in formula 5.5. Shortly, this metric measures the average unpredictability in a random variable in respect with its information content.

$$H = \sum_{j=1}^{N} -p_j * log(p_j) \quad , \tag{5.5}$$

where N is the number of possible states of the system and p_j is the probability of the state j to occur. We express the state of the system N as a function of two parameters, namely the group consumption (Q) and the group size (S_g). A state $N(Q, S_g)$ of the system represents the number of the

weak integer compositions of Q into S_g parts. In its turn, the weak integer composition of an integer n is the number of ways that the n can be expressed as the sum of non-negative integers. In our setting, the weak integer compositions indicate the number of ways that a certain group power consumption Q could have been consumed among the S_g members of the group.

Based on [16], the N can be computed with the formula 5.6. Nevertheless, due to the use of factorials, the computation becomes hard for big Q and S_g. We overcome this problem using the binomial coefficients of the Colt library [6].

$$N = \binom{Q + S_g - 1}{S_g - 1} = \frac{(Q + S_g - 1)!}{(S_g - 1)! * Q!} \tag{5.6}$$

With the proper parameters in the binomial coefficient function implementation we are able to compute the N. Moreover, we have assumed that the weak integer compositions of Q into S_g have the same probability. Hence, the formula 5.5 is transformed in the formula 5.7.

$$H = N * \frac{-1}{N} * log(\frac{1}{N}) \tag{5.7}$$

5.2 Experimental Results and Discussion

In the context of the experiments 512 households participated in the DR program for a duration of 535 days. 410 or 80% of the households acted as legitimate while 102 or 20% of the households acted as free-riders. The reduction rate was set to $\gamma == 0.2$ for the legitimate users and $\gamma == 0$ for the free-riders.

As explained in section 4.1, a DR signal arrives at each home at 18:00 in the evening every day and lasts until 21:00. During this period the households respond to the DR event by reducing their power demand. After the DR event is completed, the households report their realized consumption to the utility.

In the context of this experiment we test all the utility functions, as described in section 4.2.2, for different group sizes and we evaluate their performance. We also consider the effectiveness of the utility functions in conjunction with the two incentive allocation methods that the utility is equipped with.

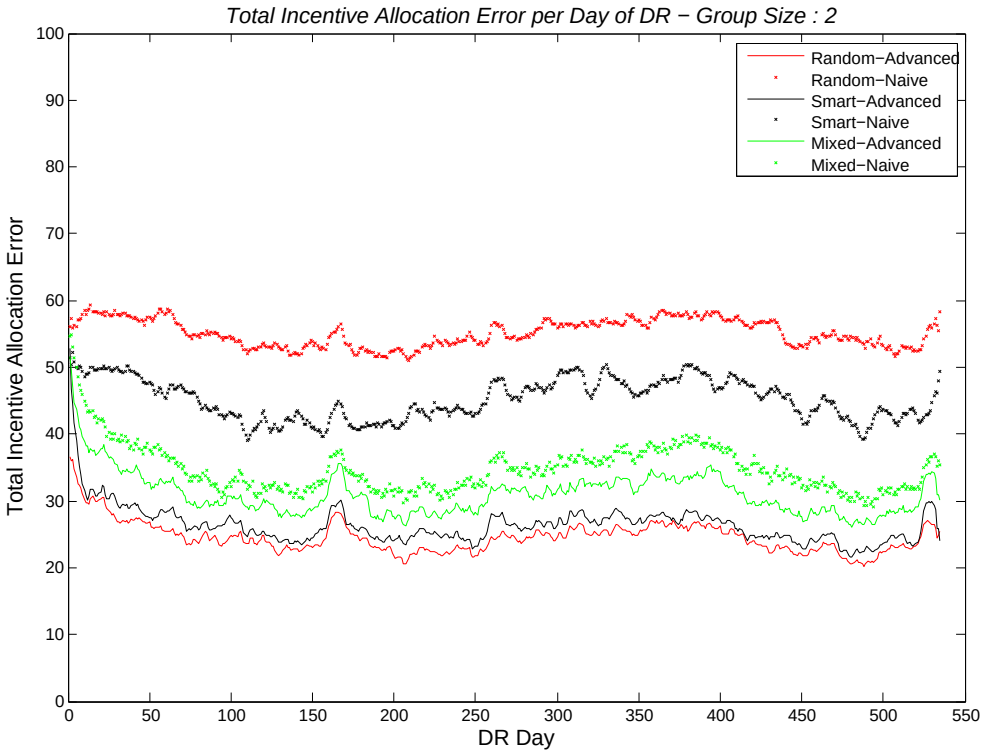

Figure 5.1: Total incentive allocation error for group size 2

The figure 5.1 illustrates the performance of the several shuffling methods of the utility in terms of the total incentive allocation error. This figure corresponds to the performance results for group size equal to 2. As it can be seen, for each shuffling method we run both naive and advanced incentive allocation mechanisms.

Random shuffling of the groups coupled with the naive incentive allocation have the worst performance in terms of allocation error. This combination is the simplest possible and we use it as our baseline in order to be able to compare other approaches against it. The bad performance is reasonable since we do not manage to separate the legitimate users from the free-riders by randomly shuffling the groups of households. Furthermore, we use naive incentive allocation which equally shares the group incentive among the group members. In fact, with the naive allocation mechanism we are not able to

fairly distribute the group incentive and the high allocation error is therefore expected. By shuffling with the utility smart function we manage a more impactful separation of the participants. This is depicted on the graph 5.1, where the smart shuffling combined with the naive allocation shows better performance than the baseline. The mixed function gives the best results among those which use the naive allocation mechanism. This is due to the way that the mixed method works. As it has been described in section 4.2.2, the mixed method uses the reputation metric in order to create new groups of households at each even DR day. In fact, it seems that this approach further decreases the allocation error.

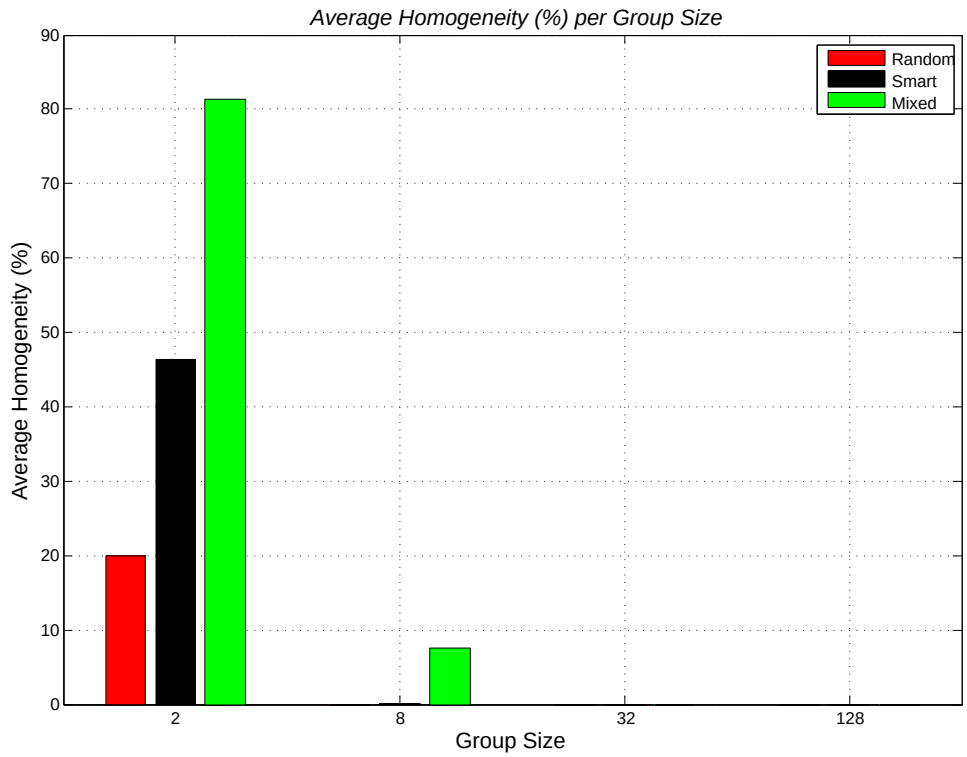

Figure 5.2: Average homogeneity per group size for all utility functions

Despite the improvement, the total incentive error still remains high. Surprisingly enough, the same utility functions follow the inverse trend when they are combined with the advanced incentive allocation mechanism. Hence,

the random function causes the lowest allocation error with the smart function coming second and the mixed one showing the poorest performance, as it can be seen in figure 5.1.

In figure 5.2, the average homogeneity over the 535 days of the experiment is plotted for every group size. As we have described earlier, the homogeneity is an indicator of the impactful separation of the legitimate users from the free-riders. The trend of the allocation error is inverse to the trend of the homogeneity when we consider the naive incentive allocation approach. In fact, for group size equal to 2, the higher the homogeneity, the lower the incentive allocation error, as it is illustrated in figure 5.1.

However, the same observation does not hold when we consider the advanced incentive allocation mechanism. This particular mechanism works based on the reputation of every household. The reputation is the accumulated incentive of all previous DR days that a household received. Intuitively, to get a reputation that accurately reflects the power consumption behavior of a DR participant we need to meet two requirements. First, we need to achieve high homogeneity. Second, we need as many group shuffles as possible. With the first requirement we do not assign reputation credits to the free-riders, whereas with the second requirement we further increase the reputation of the legitimate users. Based on these observations, one would argue that the combination of the mixed utility function with the advanced allocation mechanism should perform better than what is shown in figure 5.1. Nevertheless, the mixed function shuffles the groups of households half the times as other utility functions (random, smart). Even though mixed utility function achieves high homogeneity score, as it is illustrated in figure 5.2, its performance, in terms of total incentive allocation error, is worse compared to the random and smart utility functions.

Figure 5.3 depicts the total incentive allocation error for group size 8 and 32. We can see that the allocation error significantly increases while the performance of all the utility functions drops. Nevertheless, we can also note that when the utility functions utilize the advanced allocation mechanism, the error still remains less compared to the naive incentive allocation approach.

However, we can also observe that the incentive allocation error is tightly correlated with the group size, since the bigger the group the higher the error. This is expected due to the inability of the utility to distinguish among different DR participants' behavior when a group consists of a big number of

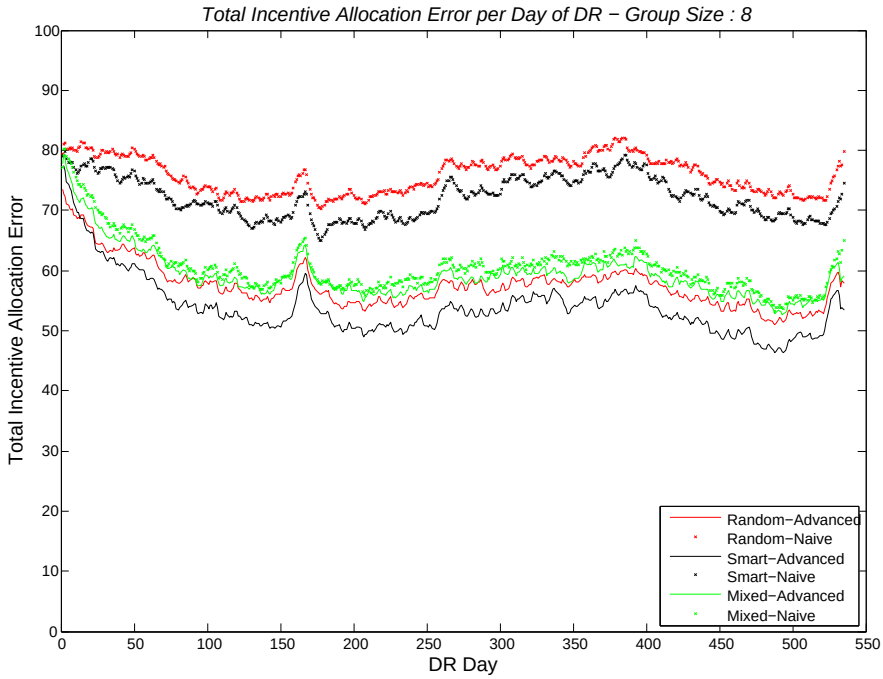

(a)

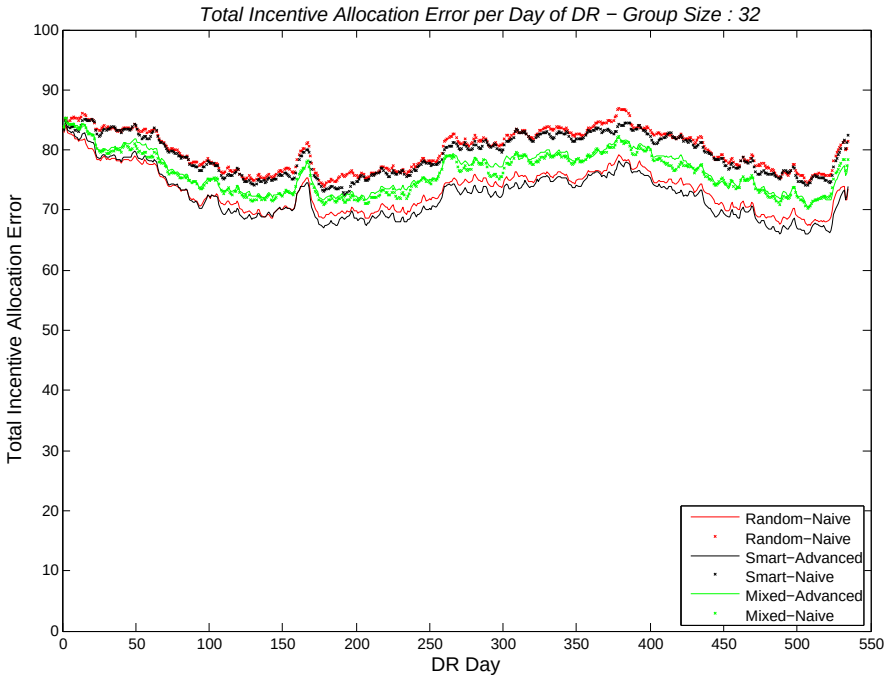

(b)

Figure 5.3: Total incentive allocation error for group size 8 & 32

households. Moreover, as it is shown in figure 5.2, our shuffling methods do not manage to totally separate the legitimate users from the free-riders for group size 32 and 128. Indeed, the zero homogeneity for big groups indicates that in every DR round every group has at least one free-rider. This fact affects the performance of our system for big group sizes as it is depicted in figures 5.3 & 5.4.

Last but not least, for all group sizes, the advanced allocation mechanism outperforms the naive one, no matter which utility function they are combined with, boosting in this way the performance of the system in terms of incentive allocation error. Nevertheless, as the group size increases, all the combinations of the shuffling methods with the utility functions gradually converge to high error rates.

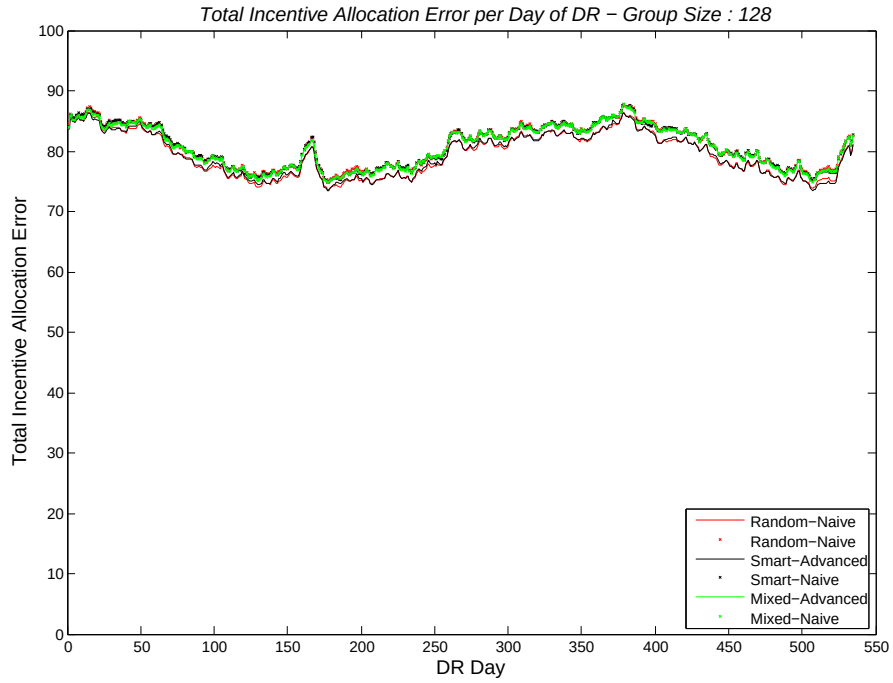

Figure 5.4: Total incentive allocation error for group size 128

The average allocation error per group size for all the 535 days of the DR program gives another view of our system performance. This is illustrated in figure 5.5. More specifically, in the subfigure 5.5(b), we present the average error as a consequence of the application of the naive allocation mechanism.

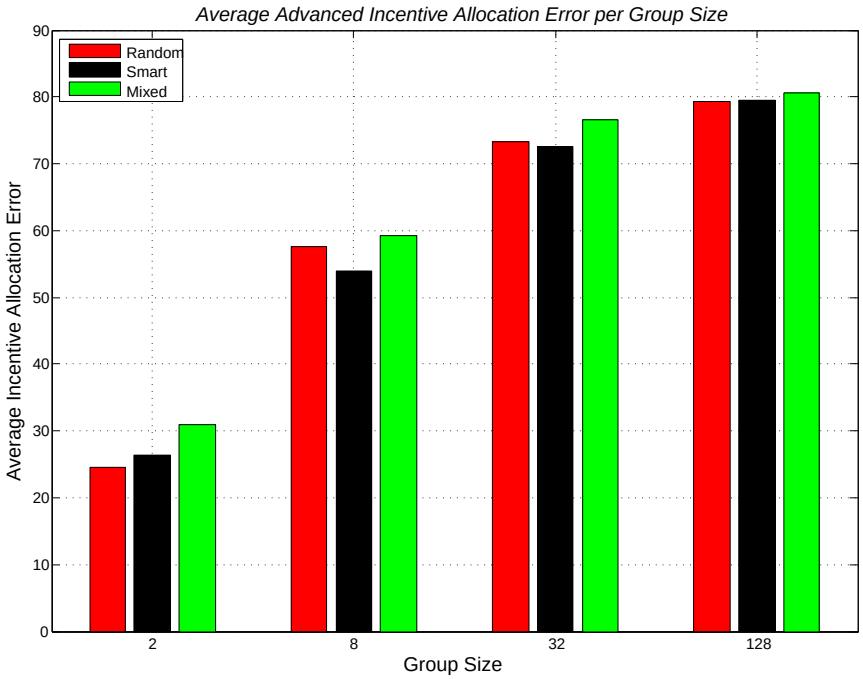

(a) Error of the advanced incentive allocation method

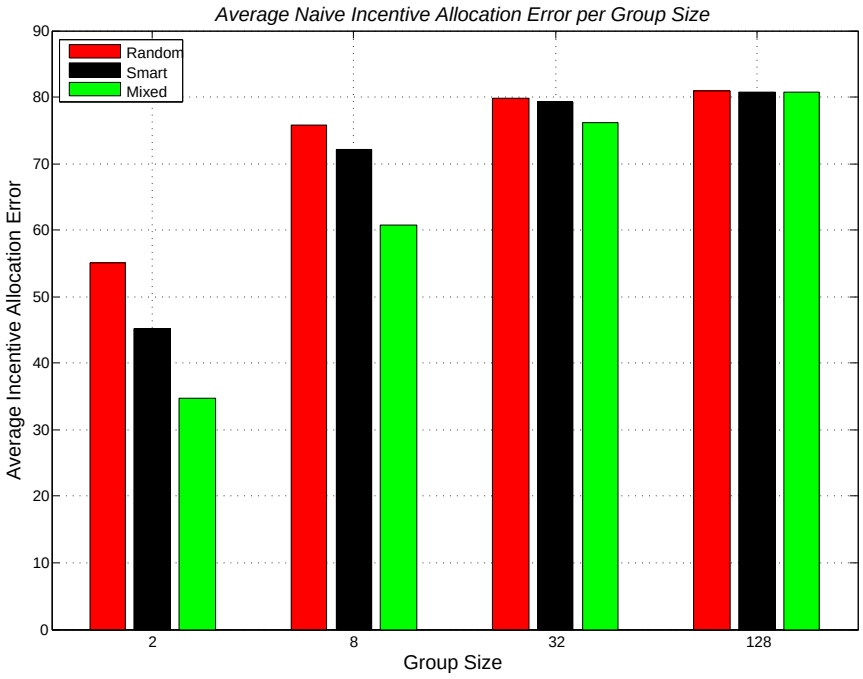

(b) Error of the naive incentive allocation method

Figure 5.5: Average allocation error per group size

As it can be seen, our smart and mixed utility functions outperform the random function (baseline), even though the gain in performance drops while the group size is getting bigger. As for the advanced allocation mechanism (subfigure 5.5(a)), the total average allocation error is constantly lower when compared to its naive counterpart (subfigure 5.5(b)).

As far as the privacy of our system is concerned, intuitively, the bigger the size of a group the higher the privacy. In fact, this is verified by the figure 5.6. We observe that as the group size gets bigger the entropy increases as well. Consequently, the privacy increases as the group size gets bigger.

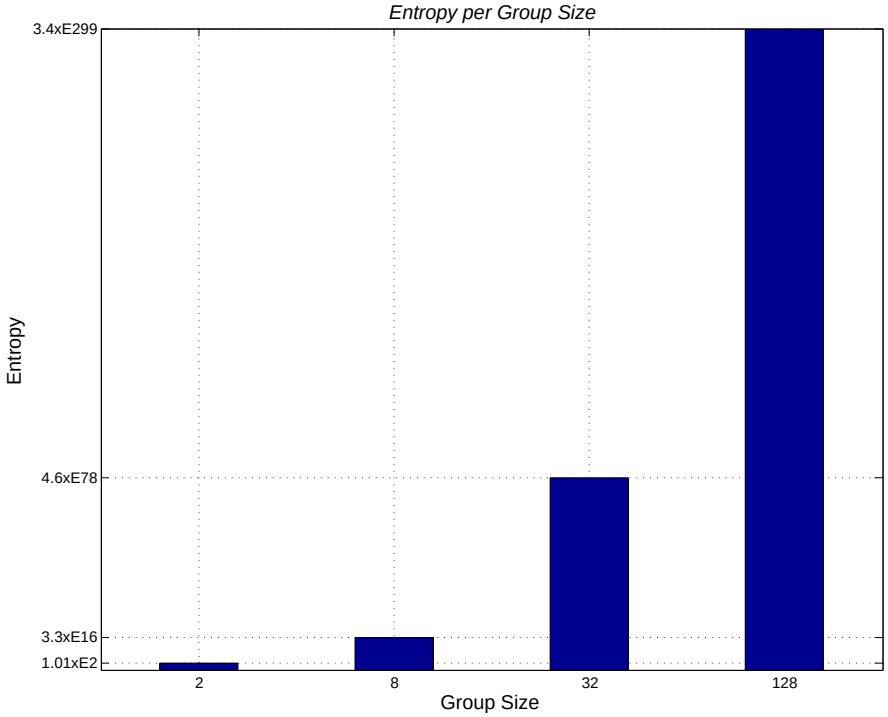

Figure 5.6: Total incentive allocation error for group size 128

The entropy values illustrated in figure 5.6 regard the average entropy of a group of a certain size over the 535 days that the DR simulation lasts. More specifically, for a certain group size we first calculate the entropy of every group for each DR day. Then, for every particular DR day we compute the average entropy over the groups, as shown in formula 5.8.

$$H_{day} = \frac{\sum_{g \in G} H_g}{S_G} \quad , \tag{5.8}$$

where g is a group of the set of groups G and H_g is the entropy of that group in reference with its power consumption. In addition, S_G is the number of groups in the system. Finally, we compute the average entropy over the 535 days of the DR for a fixed group size, as shown in formula 5.9

$$H_{exp} = \frac{\sum_{day=1}^{T} H_{day}}{T} \quad , \tag{5.9}$$

where T is the number of the DR rounds in days.

Chapter 6

Conclusions and further work

Smart Grid and Demand Response paradigm are technologies that can bring significant benefits, as we have described in Section 2.2.3. Nevertheless, important steps must be taken towards the improvement of DR programs before its full potential can be realized. The improvement should regard both the accuracy of the DR service as well as the valuable privacy of the individual electricity users.

In this thesis, several aspects of the DR paradigm have been studied, whereas a simulated DR system was developed with the goal of analyzing its performance, mainly, in terms of utility and privacy. The main contributions of the thesis are :

1. A background part setting the scene of the DR paradigm in the Smart Grid.
2. A taxonomy of the proposed PETs in the literature, mainly based on their underlying cryptographic techniques.
3. Enumeration and analysis of the requirements in a DR environment from the point of view of both the users and the utility company.
4. Evaluation of a set of proposed PET mechanisms in the DR context.
5. Analysis and comparison of homomorphic aggregation and differential privacy.
6. Implementation and evaluation of a DR system.

The background part elaborates on the Smart Grid stakeholders as well as the technological components and advancements that enable its deployment. Moreover, special attention has been given on the Demand Response

paradigm. The benefits that it brings are discussed along with the most significant challenges. In addition, various DR schemes are presented. In the background section, the reader will also find an extended discussion in reference with the privacy in the SG. Last but not least, several privacy-enhancing technologies are categorized and an analysis takes place in terms of their properties.

In chapter 3 we commence by presenting and discussing the requirements that both the DR participants and the energy provider need to live up, in order for a DR program to be successfully deployed. Then, we conduct an in depth analysis of some featured PETs that could be employed over a DR environment to protect individuals' privacy. We investigate their implications in terms of privacy as well as the impact they pose on the efficient operation of a DR scheme. This section concludes with a comparison between two promising PET classes, namely the homomorphic aggregation and the differential privacy. More specifically, we elaborate on the privacy guarantees of each class identifying their strengths and weaknesses and analyzing their operation.

In the last two chapters 4 and 5 we present the simulated DR system, its functions and components as well as the experimental results. In more detail, we describe the three functions that the DR system is equipped with in order to distinguish the free-riders from the legitimate users. Moreover, we elaborate on the two methods that the simulated utility company uses in order to split the incentive among the DR program participants. Finally, we evaluate the result in terms of the incentive allocation error, the DR participants homogeneity and the system entropy that is tightly correlated with the privacy.

6.1 Future work

Privacy enhancing mechanisms should be considered as an integral part of every DR program since privacy concerns could hinder the Demand Response wide adoption by the public. However, at the same time they pose challenges in unraveling its full potential. For instance, in our study we have shown that the accuracy of the incentive allocation mechanism is affected by the privacy protection technology that we have deployed on top of the DR mechanism.

Based on our research, it would be interesting for someone to investigate others methods that could possibly better separate the free-riders from

the legitimate users. In that respect, the homogeneity would be improved and a more fair incentive allocation could take place. Furthermore, another research path could be the improvement and optimization of the incentive allocation mechanisms itself. Needless to say that any mechanism proposed should carefully consider the users' privacy concerns.

Finally, along with the aforementioned ideas, it would also be interesting for someone to investigate what the impact of another PET would be one the system we have simulated in the context of this thesis. Then a comparison between the two systems, in terms of the accuracy, homogeneity and entropy indicators, could give useful insights and new ideas that could help to build better incentive allocation mechanisms.

Bibliography

[1] Benefits of demand response in electricity markets and recommendations for achieving them. Tech. Rep. February, U.S Department of Energy - Office of Electricity Delivery & Energy Reliability, 2006.

[2] The Smart Grid : An introduction. Tech. rep., U.S Department of Energy - Office of Electricity Delivery & Energy Reliability, 2010.

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

[4] Cárdenas, A., Amin, S., and Schwartz, G. Privacy-Aware Sampling for Residential Demand Response Programs.

[5] CER. Smart metering trial data publication, – http://www.ucd.ie/issda/data/commissionforenergyregulationcer/2012.

[6] CERN - European Organization for Nuclear Research. Official site "Colt Project Page", http://acs.lbl.gov/software/colt/index.html, 2013.

[7] Clark, S. S., Ransford, B., Sorber, J., Xu, W., Learned-miller, E., and Fu, K. Current Events : Identifying Webpages by Tapping the Electrical Outlet.

[8] Dwork, C. Differential Privacy. *33rd International Colloquium on Automata, Languages and Programming, part II (ICALP 2006) 4052* (2006), 1–12.

[9] Dwork, C. Differential privacy: A survey of results. *Theory and Applications of Models of Computation* (2008), 1–19.

[10] Dwork, C. The Promise of Differential Privacy: A Tutorial on Algorithmic Techniques. *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, D (Oct. 2011), 1–2.

[11] Efthymiou, C., and Kalogridis, G. Smart Grid Privacy via Anonymization of Smart Metering Data. *2010 First IEEE International Conference on Smart Grid Communications* (Oct. 2010), 238–243.

[12] Electric Power Research Institute. IntelliGrid Program, 2012.

[13] Goldwasser, S., Micali, S., and Rackoff, C. The knowledge complexity of interactive proof-systems. *Proceedings of the seventeenth annual ACM symposium on Theory of computing - STOC '85* (1985), 291–304.

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

[15] Hart, G. Residential energy monitoring and computerized surveillance via utility power flows. *IEEE Technology and Society Magazine 8*, 2 (June 1989), 12–16.

[16] Heubach, S., and Mansour, T. *Combinatorics of compositions and words.* Discrete Mathematics and its Applications. Boca Raton, FL: CRC Press. xxiii, 480 p., 2009.

[17] Jawurek, M., Johns, M., and Rieck, K. Smart metering depseudonymization. *Proceedings of the 27th Annual Computer Security Applications Conference on - ACSAC '11* (2011), 227.

[18] Jawurek, M., Johns, M., and Rieck, K. Smart Metering De-Pseudonymization. *ACSAC '11 Proceedings of the 27th Annual Computer Security Applications Conference* (2011), 227–236.

[19] Kalogridis, G., and Denic, S. Z. Data Mining and Privacy of Personal Behaviour Types in Smart Grid. *2011 11th IEEE International Conference on Data Mining Workshops* (2011), 636–642.

[20] Kursawe, K., Danezis, G., and Kohlweiss, M. Privacy-friendly Aggregation for the Smart-grid. *Privacy Enhancing Technologies 6794* (2011), 175–191.

[21] Li, F., and Luo, B. Secure Information Aggregation for Smart Grids Using Homomorphic Encryption. *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference* (2010), 327–332.

[22] Li, F., Luo, B., Liu, P., and Hammad, P. E. Secure Information Aggregation for Smart Grids using Homomorphic Encryption - Presentation.

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

[24] Lisovich, M. A., Mulligan, D. K., and Wicker, S. B. Inferring Personal Information from Demand-Response Systems. 11–20.

[25] Molina-Markham, A., Shenoy, P., Fu, K., Cecchet, E., and Irwin, D. Private memoirs of a smart meter. *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building - BuildSys '10* (2010), 61.

[26] Paillier, P. Public key cryptosystems based on composite degree residuosity classes. *Advances in cryptology, EUROCRYPT '99 1592* (1999).

[27] Rajagopalan, S., and Sankar, L. Smart meter privacy: A utility-privacy framework. *Smart Grid ...* (2011).

[28] Rottondi, C., Verticale, G., Capone, A., Milano, P., and Leonardo, P. A Security Framework for Smart Metering with Multiple Data Consumers.

[29] Sanders, W. H. Progress towards a resilient power grid infrastructure. *IEEE PES General Meeting* (July 2010), 1–3.

[30] Shannon, C. A mathematical Theory of Communication. *The Bell System Technical Journal XXVII* (1948).

[31] Shi, E., Chan, T., Rieffel, E., Chow, R., and Song, D. Privacy-preserving aggregation of time-series data. *Proceedings of NDSS* (2011).

[32] Wattalyst. Knowledge Gaps and Key Performance Indexes. Tech. rep., 2012.

[33] Wattalyst Consortium. Wattalyst Program, http://www.wattalyst.org/WattalystWebsite/index.html/, 2013.

[34] Zhu, N., Bai, X., and Meng, J. Benefits Analysis of All Parties Participating in Demand Response. *2011 Asia-Pacific Power and Energy Engineering Conference* (Mar. 2011), 1–4.