

Understanding Game-Based Privacy Proofs for Energy Consumption Aggregation Protocols

Andreas Unterweger, Sanaz Taheri-Boshrooyeh, Günther Eibl, *Member, IEEE*, Fabian Knirsch, Alptekin Küpçü, *Member, IEEE*, and Dominik Engel, *Member, IEEE*

Abstract—Despite the large number of privacy-preserving aggregation protocols in the Smart Grid, there is no common methodology for evaluating and comparing their privacy guarantees. Protocol discussion often lacks a formal evaluation of the proposed privacy guarantees. In order to transfer the well-established formal methodology of game-based proofs to the Smart Grid domain, in this paper, we present (i) a game-based privacy definition which addresses the privacy requirement to be captured in an aggregation protocol (the definition may be used or extended for other protocols); (ii) we exemplify our game-based proof technique for two aggregation protocols, and (iii) we provide a novel and compact way to visualize and easily compare the privacy guarantees of different protocols. We employ two sample protocols that reflect the basis of the most common approaches currently found in the energy aggregation literature. In summary, we contribute a guideline on how to conduct formal evaluations for protocol developers as well as an easy-to-understand way to assess the privacy guarantees of different aggregation protocols for non-experts.

Index Terms—Smart Grid, Aggregation, Privacy, Game-based Proof, Visualization

I. INTRODUCTION

FOR some use cases in the Smart Grid, e.g., grid stability and load forecasting, the total energy consumption of a neighborhood, city, or region is needed [1]. While aggregating, i.e., adding up, the individual consumption values of each household in an area seems to be a trivial task, straight-forward summation would expose each household's contribution although only the sum of all consumption values is needed [2]. This raises privacy concerns [3], especially if smart meters measure at high resolutions [4].

The concerns lead to a large number of proposals for privacy-preserving aggregation protocols, e.g., [5]–[8]. Protocols that protect customer privacy aim at reducing the data to the required minimum for the purpose of providing a particular service [9]. For aggregation protocols, this is reflected by providing data at minimum required spatial or temporal resolution needed for the use case, e.g., network monitoring or billing.

While the terms *privacy* and *security* are often used inconsistently, for the purpose of this paper, *privacy* is defined as protecting legally acquired data from illegal or unauthorized use (e.g., smart meters learning other smart meters' individual

energy consumption), whereas *security* refers to an external attacker affecting correctness (e.g., changing the aggregate). In this paper, our focus is privacy only. Protocols that protect customer privacy therefore aim at reducing the data to the required minimum for the purpose of providing a particular service [9]. For aggregation protocols, this is reflected by providing data at minimum required spatial or temporal resolution needed for the use case, e.g., network monitoring or billing.

Different methods to ensure privacy are used by various privacy-preserving aggregation protocols, e.g., homomorphic encryption [10], masking [6], and secret sharing [11]. Despite the indicated advantages and disadvantages of each method, choosing a protocol from the wide variety is difficult due to the vast differences in privacy guarantees that the corresponding publications make. There are two main reasons for these differences.

First, different publications assume different adversaries and adversary capabilities, i.e., who may attack the data to be aggregated and how. While some authors consider very powerful adversaries, e.g., one capable of manipulating data and colluding with other parties participating in the protocol [12], others only consider a subset of honest parties which must not collude, e.g., [10], [13]. This makes it very hard to compare different protocols.

Second, and more importantly, different publications use different levels of rigor to prove the privacy-preserving properties of the protocols they propose. “Proofs” range from short arguments in prose (e.g., [2], [12]) to actual game-based proofs (e.g., [14]–[16]). On the one hand, this limits the number of available protocols with rigorous proofs of their privacy-preserving properties, while, on the other hand, for non-experts, protocols with in-depth proofs are sometimes hard to follow (e.g., [6]) and thus difficult to classify in terms of their exact privacy guarantees.

To make the comparison of privacy guarantees of aggregation protocols easier, in this paper we present the following:

- First, we provide a rigorous game-based definition of the required privacy guarantees for aggregation protocols. Note that, while other proof techniques like simulation-based proofs also exist [17] in the context of aggregation protocols, for the sake of presentation and space, we leave them out of scope and as future work.
- Second, we provide exemplary formal game-based proofs of sample aggregation protocols, based on our formal definition and cryptographic methodology.
- Third, we describe *privacy levels* that reflect different amounts of effort required to break the privacy of a

A. Unterweger, F. Knirsch, G. Eibl and D. Engel are with the Center for Secure Energy Informatics, Salzburg University of Applied Sciences, Puch bei Hallein, Austria.

S. Taheri and A. Küpçü are with the Cryptography, Security and Privacy Research Group, Koç University, İstanbul, Turkey.

Manuscript received December 5, 2017; revised October 19, 2018.

customer participating in an aggregation protocol.

- Fourth, *maximal collusion sets* are elaborated that define which maximal subset of actors participating in a protocol may be adversaries so that a certain privacy level can still be guaranteed. Adding any other colluding party will break privacy.
- Fifth, a compact way to *visualize* the privacy levels achieved by a protocol, for each maximum collusion set, is presented. This illustrates how well the privacy offered by the protocol changes for different adversaries.

The last part of our five-part contribution, supported by the other four, simplifies the comparison of aggregation protocols with respect to privacy guarantees. Although there is literature on game-based proofs for aggregation protocols, e.g., [14], [16], such proofs often require advanced cryptographic knowledge. By contrast, this paper addresses domain experts with limited cryptographic knowledge, i.e., even those who are not willing or able to follow the details of rigorous proofs themselves may (i) assess protocols with proven privacy guarantees through an easy-to-understand visualization; or (ii), as protocol designers with more cryptographic knowledge, rely on our results to prove the properties of their own protocols.

When referring to privacy in the context of energy consumption aggregation, this paper relies on the concept of *unlinkability* [18], [19]. Simply put, the privacy of a smart meter is preserved if an adversary is not capable of retrieving individual measurements of said smart meter after the aggregation protocol has been executed. As such, we propose an experiment called *unlinkability* relying on the very well-known game-based proof techniques. Our proposed game captures all the privacy aspects of an aggregation protocol and shall be treated as a framework to enable a fair and rigorous privacy comparison among the existing aggregation protocols. A mathematical and concrete definition of unlinkability will be provided in the following sections.

Note that we describe ways to analyze and compare the privacy guarantees of arbitrary aggregation protocols. We do not propose an aggregation protocol, but employ two example protocols, which rely on the building blocks commonly used by aggregation protocols – homomorphic encryption [10] and masking [6].

This paper is structured as follows: First, we introduce some preliminaries in Section II and describe two sample aggregation protocols based thereon in Section III. Second, we define games for the privacy of these protocols in Section IV and subsequently prove and analyze them in Section V. Finally, in Section VI we present a way to visualize the privacy guarantees for the analyzed aggregation protocols, before we conclude our paper in Section VII.

II. PRELIMINARIES

The protocols presented in this paper build on homomorphic encryption and masking for privacy. In the following, these basic cryptographic schemes are presented, together with an overview of the involved parties and their goals.

A. Involved Parties

In aggregation protocols, multiple smart meters (customers) send measurements (e.g., energy consumption values) to a data concentrator, who computes the sum of all measurements. In some protocols, additional parties, e.g., aggregators, are present. A potential adversary who is capable of taking any of the participants' roles tries to find an individual smart meter's measurement (energy consumption) *after* the aggregation protocol is executed (i.e., we do *not* consider an adversary affecting the correctness of the protocol by providing incorrect measurement). The goals of the parties are the following:

- The **smart meter** does not want to reveal its measurements to any involved party. The privacy of the smart meter depends on the *unlinkability* of its measurement (for the formal mathematical definition, see Section IV). Each smart meter contributes a protected (e.g., encrypted) measurement which, together with the contributed measurements of all other smart meters in the network, can be used by the data concentrator, e.g., for load forecasting of neighborhoods or cities.
- The **data concentrator** is capable of computing the sum of protected measurements (of all smart meters) and gets the plain result of *only* the sum, e.g., for load forecasting. At this point, the goal of the aggregation protocol as well as our analysis ends. Yet, the data concentrator is not capable to extract individual measurements from the sum, if the number of smart meters is large enough [20].
- The **aggregator**, if present, performs mathematical operations on the protected measurements, but has no means to see the bare individual measurements. The aggregator passes the result of its computations (i.e., the sum of protected measurements) to the data concentrator.
- The **adversary**, if present, is capable of taking any role in the network, i.e., it may be another participating smart meter (or several), the data concentrator, the aggregator, or a combination of the aforementioned. The adversary follows the protocol specifications, i.e., he is honest but curious [10], and tries to use all collected information to break the smart meter's privacy by extracting or computing its measurement, i.e., by breaking unlinkability.

It is clear that the success of the adversary depends on which parties he controls and how many. One of the goals of this paper is to show how to prove the limit (maximal sets) of parties that the adversary can control without breaking the privacy of an individual smart meter. More powerful adversaries, e.g., those who can tap wires and/or disobey the protocol deliberately [21], may be capable of breaking privacy in cases an honest-but-curious adversary cannot. However, in our paper, we do not consider attacks on security and correctness (e.g., cybertampering as in [22]), and rather focus on privacy and analyze two relatively simple example protocols in terms of their privacy guarantees, i.e., what the limits of an honest-but-curious adversary are.

B. Homomorphic Encryption

Let $c_i = E_{pk}(m_i)$ denote the encryption of plaintext m_i with public key pk and $D_{sk}(c_i) = m_i$ denote the correspond-

ing decryption of ciphertext c_i with private key sk . For the purpose of this paper, only an *additively homomorphic* encryption is needed. For example, for the additively homomorphic Paillier cryptosystem [23], the property $D_{sk}(E_{pk}(m_1) \cdot E_{pk}(m_2) \bmod n^2) = m_1 + m_2 \bmod n$ holds, i.e., an operation exists, that if performed on two ciphertexts in the encrypted domain, corresponds to an addition of two plaintexts. This allows aggregating values in the encrypted domain. The Paillier cryptosystem has already been used in smart grid aggregation protocols, e.g., [2], [10], [24], [25]. In this cryptosystem, the multiplication of ciphertexts corresponds to the addition of plaintexts. The parameter n determines the security of the scheme and is recommended to be at least 2048 bits [26].

C. Masking

Masking is a lightweight scheme that adds an additive secret s_i to a plaintext m_i , i.e., $\tilde{m}_i = m_i + s_i \bmod k$. s_i is a random number that is uniformly drawn from the full range of the plaintext values, e.g., $0 \dots k-1$ using a cryptographically secure random number generator [15]. The random numbers are constructed in such a way that they either cancel each other out upon summation, i.e., $\sum_i s_i \bmod k = 0$, or the sum of the random numbers is subtracted at the end of the protocol, i.e., $\sum_i \tilde{m}_i - \sum_i s_i \bmod k = \sum_i m_i \bmod k = \sum_i m_i$. For the last equality to hold, k must be chosen large enough, i.e. $k > \sum_i m_i$. Masking allows to aggregate the data without revealing individual measurements, but yields the exact result at the end of the protocol. Masking is proposed for smart grid applications by, e.g., [6], [14], [27].

III. AGGREGATION PROTOCOLS

This section introduces the sample aggregation protocols we will analyze. Let the set of N smart meters be denoted by SM_1, \dots, SM_N . A smart meter i measures a value m_i at time t , usually in the range $[0, 2^{16} - 1]$. Depending on the protocol, smart meters report their measurements either to the aggregator \mathcal{A} , or the data concentrator DC , or to other smart meters. In the following, two basic types of aggregation protocols using (i) homomorphic encryption and a star network; and (ii) masking and a star and ring network, are described in detail. For ease of description here, the protocols are split into three phases, namely *initialization*, *measurement submission*, and *aggregation*. Measurement submission refers to sending the measurements for one point in time t , e.g., the aggregation of values in a 15 minutes interval.

A. Protocol I: Homomorphic Encryption with Star Network

For this type of aggregation protocols, all smart meters SM_1, SM_2, \dots, SM_N are arranged in a star network and report their additively homomorphic encrypted measurements to an aggregator \mathcal{A} . By exploiting the additively homomorphic property, the aggregator calculates an encrypted sum of the individual encrypted measurements and forwards this encrypted sum to the data concentrator DC . The setup is shown in Fig. 1.

a) *Initialization*: DC generates a pair of public/private keys for an additive homomorphic cryptosystem (e.g., Paillier cryptosystem [23]). The public key is shared with all SM_i and \mathcal{A} , whereas the private key is kept by DC .




Fig. 1. Star network aggregation protocol with homomorphic encryption: Each smart meter SM_i sends its encrypted reading c_i to the aggregator \mathcal{A} , which computes the ciphertext product $\prod_i c_i$ and sends it to the data concentrator DC , which decrypts this product to obtain the plaintext sum of the readings.




Fig. 2. Star and ring network aggregation protocol with masking: Each smart meter SM_i adds a random number s_i to its reading m_i , yielding \tilde{m}_i that is sent to the data concentrator DC . Additionally, the random shares are aggregated by the smart meters and forwarded to data concentrator DC , which uses the sum of shares to retrieve the plaintext sum.

b) *Measurement submission (for each t)*: Each SM_i encrypts its measurement m_i with the public key, yielding a ciphertext $c_i = E_{pk}(m_i)$ which is forwarded to \mathcal{A} .

c) *Aggregation*: \mathcal{A} computes the encryption of the total consumption of all smart meters, denoted as \tilde{M} , by calculating the ciphertext product of the additively homomorphic encrypted values c_i provided by the smart meters by $\tilde{M} = \prod_{i=1}^N c_i$. This value is then forwarded to DC for decryption. DC calculates the total consumption M by decrypting with its private key by $M = D_{sk}(\tilde{M}) = D_{sk}(\prod_{i=1}^N c_i) = \sum_{i=1}^N m_i$.

B. Protocol II: Masking with Star and Ring Network

This type of aggregation protocol is based on masking instead of additive homomorphic encryption. The protocol is based on a star and a ring topology. All smart meters and DC are connected in a ring and star network, and there is no distinct aggregator. Note that the ring network is only a logical link and is not necessarily reflected by the physical wiring. In practice, the end-to-end encrypted ring links can be physically routed via the data concentrator. The smart meters send their masked measurements directly to DC in a star network and they send the shares used for masking hop-by-hop to the following smart meter in a ring network. The smart meters aggregate their own share with the received shares and the last smart meter in the ring forwards the sum of the shares to DC . The principal setup is shown in Fig. 2.

a) *Initialization*: For initialization, the smart meters and DC are arranged in an ordered sending list L that determines the network (with DC being the first and last one in the sending list), e.g., $L = (DC, SM_1, SM_2, \dots, SM_N, DC)$. This list

is public and known to all participants. The modulus k is distributed to all parties.

b) *Measurement submission (for each t):* Each SM_i draws a random number s_i uniformly from $\{0, \dots, k-1\}$ using its pseudorandom number generator (as in Section II) and masks its measurement m_i by computing $\tilde{m}_i = m_i + s_i \bmod k$. The smart meters submit \tilde{m}_i to \mathcal{DC} over the star topology. Note that at this point \mathcal{DC} has no means of retrieving the aggregate, as only masked values have been received.

In order to allow \mathcal{DC} to recover the aggregate, the sum of random values is sent to \mathcal{DC} over the ring network. Therefore, \mathcal{DC} draws a random number $s_0 = S_0$ and submits this to SM_1 , the first smart meter in the sending list, which forwards $S_1 = S_0 + s_1 \bmod k$ (i.e., the sum of the received share and its own share) to the second smart meter in the sending list. All following SM_i , except the last one act as follows: Upon receiving S_{i-1} they calculate $S_i = S_{i-1} + s_i \bmod k$ and forward this to SM_{i+1} . The last smart meter (SM_N) calculates $S_N = S_{N-1} + s_n \bmod k$ and forwards this to \mathcal{DC} .

c) *Aggregation:* Using the obtained masked measurements \tilde{m}_i , $i = 1, \dots, N$, the sum of the shares $S_N = \sum_{i=0}^N s_i \bmod k$ and its own share s_0 , \mathcal{DC} gets the total consumption by calculating $M = \sum_{i=1}^N \tilde{m}_i - S_N + s_0 \bmod k = \sum_{i=1}^N \tilde{m}_i - \sum_{i=1}^N s_i \bmod k = \sum_{i=1}^N (\tilde{m}_i - s_i) \bmod k = \sum_{i=1}^N m_i \bmod k$.

IV. GAME-BASED PRIVACY DEFINITION AND PROOFS

In this section, we introduce the concept of game-based privacy proofs. These concepts are applied in the next sections for proving the privacy of the above aggregation protocols. Additionally, we introduce *privacy levels* to represent different privacy guarantees of aggregation protocols.

A. Game-Based Privacy Definition

In aggregation protocols, the privacy concern is to keep the confidentiality of individual measurements of smart meters while enabling the energy supplier to learn only the sum of the measurements M . Although information about an individual household's consumption could be mined from M , the question of which information can be mined exactly from the aggregate is not tackled in this paper, since the protocol's goal is to provide M in the first place. Privacy is only considered in terms of recovering exact measurements of individual smart meters, unlike, e.g., differential privacy [28]. Furthermore, correctness guarantees of aggregation protocols against malicious parties are out of the scope of this paper.

We model the unlinkability of the individual measurements of smart meters in the form of a game named *smart meters' data unlinkability*. The game is played between two parties named *challenger* and *adversary*. The challenger is an abstract entity, i.e., not an actual participant in the aggregation protocol. It represents all parties not controlled by the adversary, i.e., the challenger formally controls the entities who are concerned about their privacy and hence act honestly. The adversary is the party who acts on behalf of entities who aim at violating the privacy objective of the system (that is, they want to discover the individual consumption of smart meters).

Thus, in the game, we split the entities involved in the aggregation protocols (i.e., smart meters, the data concentrator, and the aggregator) into two *disjoint* subsets: the entities controlled by the challenger, which are called *honest*, and the ones controlled by the adversary, which are called *dishonest* or *colluding set*. In our game definition, we consider smart meters' dishonesty, unlike prior works [14], [16]. The adversary is assumed to have all the secret information of the entities under its control, as does the challenger.

In general, when the data concentrator is adversarial, independent of the aggregation protocol, the number of dishonest smart meters is limited to be at most $N-2$. It cannot be $N-1$ (meaning all but one smart meter are dishonest), since, in that case, the adversary could simply retrieve the consumption of the single honest smart meter by subtracting the consumption of the dishonest smart meters from the final aggregation value. Therefore, we assume that at least two smart meters, SM_i and SM_j , where $1 \leq i \neq j \leq N$, are honest.

The definition of unlinkability states that even when the adversary knows the individual measurements of the two honest smart meters independently, it cannot know which meter had which one of those known measurements after the aggregation protocol has been executed. The adversary's failure in the game indicates the privacy of that aggregation protocol against the presumed colluding set. Accordingly, in a game-based proof, we measure the privacy of the aggregation protocol by assessing the success probability of the adversary. This probability states how likely the adversary is to actually find the correct value of an individual smart meter's measurement, i.e., to break unlinkability and thus privacy. This is a very powerful honest-but-curious adversary: the adversary knows the individual measurements, but cannot link those measurements to the correct smart meter much better than random guessing. Since there are two honest smart meters (as the minimum requirement discussed above), random guessing has a probability of one half (0.5) of being correct.

In this paper, we consider full control of the adversary over his colluding set. One may assume different types of control, e.g., the adversary may only know the measurements of the smart meters, but not their internal states (e.g., their random choices) during the execution of protocols. It is straightforward to adapt our privacy definition to cover other levels of adversarial control.

Formally, the **smart meters' data unlinkability game** $\Gamma_{Adv}^{Unlink}(\lambda)$ for an aggregation protocol and an adversary Adv is defined as:

- 1) The *initialization* part of the aggregation protocol is run.
- 2) Adversary Adv outputs a pair of measurements m_0, m_1 within the measurement domain.
- 3) A random bit $b \in \{0, 1\}$ is chosen by the challenger Ch . Then Ch assigns m_b to SM_i and $m_{\bar{b}}$ to SM_j .
- 4) The remaining parts of the aggregation protocol are executed. Using the obtained data, the adversary Adv tries to determine b . His guess is denoted as b' .
- 5) The adversary outputs b' and *wins* if and only if his guess is correct, i.e., $b = b'$. In this case, the output of the game is defined as 1, and 0 otherwise.

Intuitively, an aggregation protocol provides smart meters' data unlinkability if an adversary has only a negligibly higher winning probability than random guessing, i.e.,

$$Pr[\Gamma_{Adv}^{Unlink}(\lambda) = 1] = \frac{1}{2} + \epsilon$$

Both, ϵ , which is non-negative, and the protocol Γ depend on a security parameter λ . λ is an integer value that is set when the scheme is initialized and is usually viewed as the length of the key (e.g., the recommended key length of the AES encryption scheme is at least 128 bits [29]). For the aggregation protocols described earlier, λ can be seen as the length of the encryption key in the homomorphic encryption protocol, and the length of the modulus k in the masking protocol, respectively. The running time of parties and the success probability of the adversary are all expressed as functions of λ . In particular, an aggregation protocol is privacy-preserving if an adversary that runs in time polynomial in λ wins the unlinkability game with an advantage (ϵ in the above formulation) that is no better than negligible in λ . ϵ can be made arbitrarily small by using large security parameters. The formal definition of this statement follows.

Definition 1: An aggregation protocol provides *computationally hard* smart meters' data unlinkability if for all probabilistic polynomial-time (PPT) adversaries Adv there exists a negligible function $negl(\lambda)$ such that:

$$Pr[\Gamma_{Adv}^{Unlink}(\lambda) = 1] = \frac{1}{2} + negl(\lambda)$$

Definition 2: A function $f(\cdot)$ is called negligible if for all positive polynomial functions $g(\cdot)$ there exists a constant K such that for all real numbers $k > K$, $f(k) < \frac{1}{g(k)}$ holds. The stronger guarantee where $\epsilon = 0$ holds for computationally unlimited adversaries and is called *information-theoretic* smart meters' data unlinkability.

B. Game-Based Proofs

As stated above, an aggregation protocol is defined as privacy-preserving if the adversary can win the unlinkability game only with negligible advantage. The idea of game-based proofs is not to directly *prove* the privacy definition, but instead to construct a reduction that *reduces* winning the unlinkability game to winning a low-level game. If the latter is only possible with negligible probability due to the properties of the underlying problem, the former is also only possible with negligible probability. Figure 3 demonstrates the high-level idea of a reduction proof [30]. The sample problem Γ^X indicates the low-level problem. Simply put, the aim of a reduction proof is to show that there is an efficient ADV who can solve any instance of Γ^X by a polynomial number of operations and calls to the black box adversary Adv (where Adv is able to break the data unlinkability game). Such reductions are in essence similar to NP-completeness type of reductions. Yet, additionally, ADV must construct Adv 's input in an indistinguishable manner and solve Γ^X with an advantage similar to the advantage of Adv in breaking the unlinkability game. As defined above, "breaking" the unlinkability game means that Adv guesses the bit value b correctly.

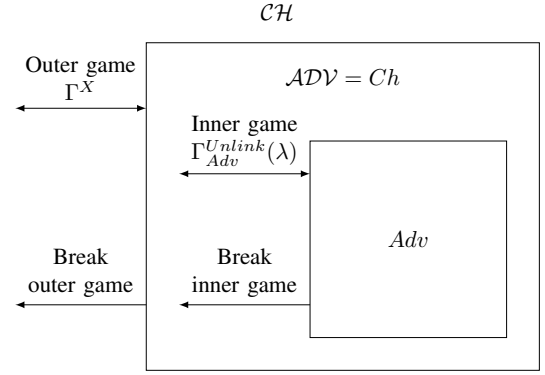


Fig. 3. Overview of a security proof by reduction: The outer adversary ADV can break the low-level game Γ^X by using the inner adversary Adv who breaks (wins) the data unlinkability game $\Gamma_{Adv}^{Unlink}(\lambda)$, which is based on the low-level game.

In the case of aggregation protocol I, the low-level problem is the chosen plaintext security of the homomorphic cryptosystem. In the case of protocol II, it is the indistinguishability of a pseudorandom generator from a truly random generator. These low-level problems are defined later in this section.

More precisely, the underlying security mechanism is formulated as an outer game in a reduction, where an outer adversary ADV calls the adversary Adv of the unlinkability game as his subroutine by simulating the role of the challenger Ch for him. By doing this, ADV benefits from Adv 's winning power to break the security of the underlying security mechanism, i.e., winning the outer low-level game which he plays with the outer challenger CH . Since $Ch = ADV$, only these three parties (CH , ADV , Adv) are involved. Note that Adv is free to perform any internal computation in order to win this game.

In the provided reduction proofs, three properties are shown: (i) the reduction, i.e., how ADV uses the success power of Adv to break the underlying security mechanism accompanied by a proof that ADV obtains non-negligible advantage in case Adv has non-negligible advantage; (ii) ADV simulates the role of a challenger for Adv indistinguishable from a real challenger. This involves ADV performing initialization, measurement submission and aggregation indistinguishable from Ch . In other words, it involves Adv having indistinguishable views when working on his own and when used by the inner game. If ADV behaves differently from Ch , it must be shown that Adv can distinguish this behavior with at most negligible probability. (iii) ADV runs in polynomial time also simulating the inner game for Adv . In the reduction proofs below this is always fulfilled and will not be shown, since Adv is assumed PPT and the reduction itself only involves simple computations and sending of messages.

In the following, we provide the security game of the underlying schemes, i.e., the CPA security game of the encryption scheme and the indistinguishability game of the pseudorandom generator. We will use these definitions in Section V for reductions.

Chosen plaintext attack (CPA) security game

Consider the following game $\Gamma_{ADV}^{CPA}(\lambda)$ which is defined for a public key encryption scheme π , consisting of key generation,

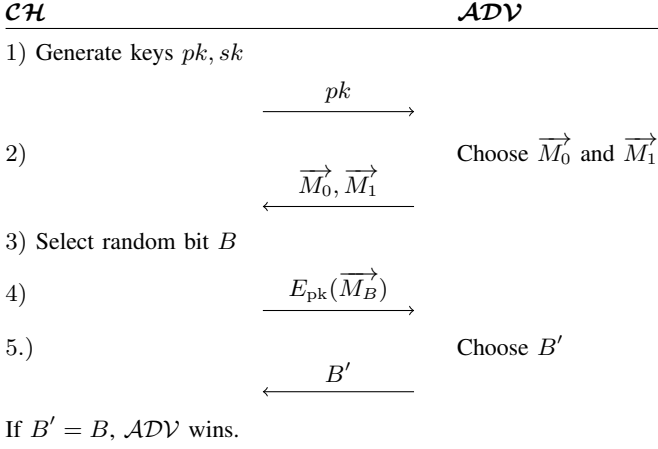


Fig. 4. Overview of the steps of the CPA game $\Gamma_{\mathcal{ADV}}^{CPA}(\lambda)$. Since the adversary is a black box, it is unspecified how he internally chooses \vec{M}_0, \vec{M}_1 and B' .

encryption, and decryption [30]. A visualization of the steps is additionally provided in Fig. 4.

- 1) Challenger \mathcal{CH} runs the key generation with the security parameter λ and obtains public and private keys pk and sk , respectively. \mathcal{ADV} is given the public key pk .
- 2) \mathcal{ADV} outputs two sets of messages $\vec{M}_0 = (m_{0,1}, \dots, m_{0,q})$, $\vec{M}_1 = (m_{1,1}, \dots, m_{1,q})$ of equal size¹, i.e., $|\vec{M}_0| = |\vec{M}_1| = q$ and $\forall i \in \{1, \dots, q\} : |m_{0,i}| = |m_{1,i}|$.
- 3) A random bit $B \in \{0, 1\}$ is chosen by \mathcal{CH} .
- 4) \mathcal{CH} sends the encryption of \vec{M}_B , i.e., $E_{pk}(\vec{M}_B) = (E_{pk}(m_{B,1}), \dots, E_{pk}(m_{B,q}))$ to \mathcal{ADV} who constructs B' .
- 5) \mathcal{ADV} outputs B' . If $B = B'$, \mathcal{ADV} succeeds and the output of the game is 1 and 0 otherwise.

Definition 3: A public key encryption scheme π has CPA security if for all PPT adversaries \mathcal{ADV} , there exists a negligible function $\text{negl}(\lambda)$ such that

$$\Pr[\Gamma_{\mathcal{ADV}}^{CPA}(\lambda) = 1] = \frac{1}{2} + \text{negl}(\lambda).$$

Pseudorandom generator [30]

Intuitively, a pseudorandom number is a deterministically created number $s \in \{0, 1\}^k$ that is indistinguishable from a truly random number $r \in \{0, 1\}^k$. Consider the following game where the adversary, modeled as a PPT distinguishing algorithm \mathcal{ADV} , is given a number S that depends on two equally likely values of B : (i) $S=s$, if $B = 1$; (ii) $S=r$, if $B = 0$. The pseudorandom number $s = P(u)$ is created by a deterministic, polynomial-time function $P : \{0, 1\}^\lambda \rightarrow \{0, 1\}^k$ with $\lambda < k$ given a uniformly random seed $u \in \{0, 1\}^\lambda$. P is a Pseudorandom Generator (PRG), if for all PPT distinguishers \mathcal{ADV} that output a guess $B' = \mathcal{ADV}(S)$ for B , we have:

$$|\Pr[B' = 1 | S = s] - \Pr[B' = 1 | S = r]| = \text{negl}(\lambda).$$

This also holds for sequences $\vec{S} = (s_i)_{i \in I}$.

C. Privacy Levels

We group privacy levels provided by the protocols by analyzing the hardness of the underlying problem. Some

¹Most encryption schemes used in practice reveal the size of the messages; hence the equal size requirement. For CPA security, single message security implies security for multiple messages [30].

problems are information-theoretically hard to solve (e.g., one-time pads). Other problems can be solved, but only with unrealistically high computational power (e.g., factorization of integers with only large prime factors).

For this latter type of privacy, we consider an adversary whose computational power is limited to being probabilistic polynomial time (PPT), and whose advantage for a privacy break is negligible (see [30] for more details and a motivation for the consideration of these properties). This means that a break can occur, but the probability of a break gets arbitrarily small for a sufficiently large security parameter λ , practically represented, e.g., by the bit length of a key. Summarizing, often *computational security* is achieved, meaning that for a sufficiently large security parameter, a computationally limited adversary can break privacy with only negligible probability.

In addition to low-level problems that are hard to solve in polynomial time, there are problems whose hardness is not well-analyzed and may rely on heuristics. We group the problems that the adversary has to solve in order to break privacy into different *privacy levels*, and present them from hardest (the strongest privacy guarantees) to easiest (the weakest privacy guarantees):

- **Information-theoretic:** Problems that cannot be solved even with infinite computing power, e.g., secret sharing [31] or the problem of obtaining data which is physically inaccessible (like an honest-but-curious smart meter cannot access data from the data concentrator in Protocol II). *Information-theoretic* smart meters' data unlinkability means that the advantage of the adversary in winning the unlinkability game is exactly zero.
- **Computationally hard:** Problems that can be solved only with an amount of computing power which is unrealistically high in the foreseeable future, e.g., homomorphic encryption [23] (as employed by Protocol I) with large modulus sizes (in general a large security parameter λ). Analogously, RSA with a 2048 bit modulus is considered secure until 2030 [32].
- **Heuristic:** Problems that are not trivial to solve in general, but may be easy for some hard-to-define cases or input data. For example, extracting (disaggregating) one distinct measurement m_d from an aggregation result $M = m_d + \sum_i m_i$ depends on the distribution of the measurements [20].

Each privacy level gives hardness guarantees about the problem the adversary has to solve in order to break privacy.

Moreover, an adversary may, in general, represent multiple parties in a protocol, i.e., by collusion. Thus, it is necessary to define privacy as a set of colluding parties which still gives a certain privacy guarantee that can be specified by a privacy level. If more parties collude, the privacy level potentially decreases. If different parties collude, the privacy level may change. In summary, privacy is specified by the maximal set(s) of colluding adversaries that *cannot* break the corresponding privacy game. For some low-level problems, the adversary needs to be computationally limited, which can be handled by choosing a sufficiently large security parameter.

V. GAME-BASED PRIVACY ANALYSIS

In this section, the privacy properties of the discussed aggregation protocols are proven formally. The task of the privacy analysis is to find the maximal set(s) of colluding adversaries that cannot break the corresponding privacy game. We assume secure and authenticated communication channels between parties throughout the paper.

Note that, while the following security proofs are presented for the maximal set of colluding adversaries, it is easy to verify that the protocol guarantees privacy under any subset of the given collusion set. This is the case since making the collusion set smaller always upholds the guarantees.

A. Privacy Analysis of Aggregation Protocol I

For Protocol I, three colluding sets are considered. In the first case, the aggregator and $N - 2$ smart meters collude: $Adv = \{\mathcal{A}, SM_i; l \notin \{i, j\}\}$, $Ch = \{DC, SM_i, SM_j\}$. In the second case, the data concentrator and $N - 2$ smart meters collude: $Adv = \{DC, SM_i; l \notin \{i, j\}\}$, $Ch = \{\mathcal{A}, SM_i, SM_j\}$. In order to illustrate that the adversarial set is maximal, it is finally shown that the collusion of the aggregator, the data concentrator, and $N - 2$ smart meters leads to a privacy break.

1) **Privacy against the aggregator and $N - 2$ smart meters:** $Adv = \{\mathcal{A}, SM_i; l \notin \{i, j\}\}$ and $Ch = \{DC, SM_i, SM_j\}$.

Theorem 1: If the encryption scheme π has CPA security, then Protocol I provides smart meters' data unlinkability against the data aggregator and $N - 2$ smart meters.

Proof: (i) First, we show the reduction (how ADV calls Adv as his subroutine to succeed in CPA security game) and then prove that ADV has non-negligible advantage in case Adv has non-negligible advantage.

- 1) In the outer CPA game, ADV is given the public key, i.e., pk (from \mathcal{CH}) which he passes to Adv in the initialization phase of the inner unlinkability game.
- 2) In the inner game, Adv outputs two measurements m_0 and m_1 of the same size to $Ch = ADV$, who sends the messages $\vec{M}_0 = (m_0, m_1)$ and $\vec{M}_1 = (m_1, m_0)$ to \mathcal{CH} . Note that \vec{M}_0 and \vec{M}_1 only differ in terms of their order.
- 3) \mathcal{CH} selects B , encrypts the corresponding message $\vec{C} := E_{pk}(\vec{M}_B) := (E_{pk}(m_B), E_{pk}(m_{\bar{B}}))$ and sends it to ADV . Through that, ADV receives the two ciphertexts since $\vec{C} = (c_B, c_{\bar{B}})$, so the order depends on B . ADV then associates c_B with SM_i and $c_{\bar{B}}$ with SM_j .
- 4) Adv and ADV run the measurement submission (with c_B used by SM_i and $c_{\bar{B}}$ used by SM_j). There, the adversary (since it controls \mathcal{A}) gets the ciphertexts c_B and $c_{\bar{B}}$. ADV also runs the aggregation with Adv .
- 5) Adv outputs a bit b' . ADV outputs the same bit as Adv , i.e., $B' = b'$. If $B = B'$, ADV succeeds and the output of the game is 1, and 0 otherwise.

Now we show that ADV wins exactly when Adv wins, i.e., ADV has non-negligible advantage if Adv has non-negligible advantage. Due to the definition of winning the CPA game and the choice of B' as b' in step 5,

$$Pr[\Gamma_{ADV}^{CPA}(\lambda) = 1] = Pr[B = B'] = Pr[B = b'].$$

Since \vec{C} was chosen as $(E_{pk}(m_B), E_{pk}(m_{\bar{B}}))$ in step 3, by comparison with step 3 of the unlinkability game, $B = b$, so

$$Pr[\Gamma_{ADV}^{CPA}(\lambda) = 1] = Pr[b = b'] = Pr[\Gamma_{Adv}^{Unlink}(\lambda) = 1].$$

(ii) After the reduction, Adv has the public key, the measurements of $\{SM_i; l \notin \{i, j\}\}$, the ciphertexts c_B and $c_{\bar{B}}$ as well as the product of all ciphertexts. Since Adv does not have the private key, the knowledge about the ciphertexts is of no use for him. In fact, due to the CPA security of the encryption scheme, it is computationally hard to distinguish between the encryption of different measurements i.e., c_B and $c_{\bar{B}}$. (iii) Since Adv is assumed PPT and the reduction has little overhead, ADV is also a PPT adversary. \square

2) **Privacy against the data concentrator and $N - 2$ smart meters:** $Adv = \{DC, SM_i; l \notin \{i, j\}\}$ and $Ch = \{\mathcal{A}, SM_i, SM_j\}$. A game-based reduction to the CPA game cannot be given since ADV needs the *private* key to fulfill his role as DC which decrypts the product of the ciphertexts. However, in the outer game he is only allowed to have the *public* key. Although the game-based reduction does not work, information-theoretic privacy guarantee is proven below.

Theorem 2: Protocol I provides information-theoretic smart meters' data unlinkability against the data concentrator and $N - 2$ smart meters.

Proof: The proof shows information-theoretic privacy, i.e., given the obtained information during the protocol, the success probability of Adv in the unlinkability game $\Gamma_{Adv}^{Unlink}(\lambda)$ must be exactly $\frac{1}{2}$. Stating step 4 in more detail, the challenger submits encrypted measurements for two honest smart meters SM_i and SM_j . The adversary also submits encrypted measurements for its own $N - 2$ smart meters $SM_l, l \neq i \neq j$. The challenger multiplies the encrypted measurements and sends the aggregated value to Adv who decrypts the aggregated value and only learns the sum of all N measurements, i.e., $M = \sum_{i=1}^N m_i$. Thus, at the end of this step, Adv can determine b' based on M only. So the success probability given the obtained information is

$$Pr(b' = b|M) = \frac{Pr(b' = 1, b = 1, M) + Pr(b' = 0, b = 0, M)}{Pr(M)}.$$

Next, we consider the chain rule of conditional probabilities

$$Pr(b', b, M) = Pr(b'|b, M) \cdot Pr(M|b) \cdot P(b).$$

Now the three terms are treated separately: (a) Adv needs to determine b' based on the available information (i.e., M , not b), so $Pr(b'|b, M) = Pr(b'|M)$; (b) due to the commutativity of the summation and the correctness of the additively homomorphic encryption scheme, M does *not* depend on the assignment of m_0 and m_1 to SM_i and SM_j , respectively (i.e., on b). Therefore $P(M|b) = P(M)$; (c) b is chosen uniformly by Ch , so $P(b) = \frac{1}{2}$. This leads to

$$Pr(b', b, M) = Pr(b'|M) \cdot P(M) \cdot \frac{1}{2}.$$

Plugging this term into the nominator terms above leads to

$$Pr(b' = b|M) = \frac{1}{2} (Pr(b' = 1|M) + Pr(b' = 0|M)) = \frac{1}{2}.$$

□ It should be noted that Adv can always calculate $m_0 + m_1$ as $M - \sum_{i,l \neq i \neq j} m_i$, but can not determine how each measurement is linked to SM_i and SM_j .

3) **Privacy violation by the data concentrator, the aggregator, and $N-2$ smart meters:** To show the privacy violation, one needs to provide the code for a PPT Adv who achieves non-negligible advantage. For $Adv = \{DC, \mathcal{A}, SM_i; l \notin \{i, j\}\}$ and $Ch = \{SM_i, SM_j\}$, the privacy violation of the honest smart meters' data can be easily proven formally with the unlinkability game as follows.

Theorem 3: Protocol I does not provide smart meters' data unlinkability against the data concentrator, aggregator and $N-2$ smart meters.

Proof: To show that a case where the protocol does not provide unlinkability, we provide a polynomial-time adversary that wins the unlinkability game with non-negligible advantage.

- 1) During initialization, Adv as DC generates the key pair.
- 2) Adv outputs a pair of measurements m_0, m_1 .
- 3) A random bit $b \in \{0, 1\}$ is chosen by Ch who then assigns m_b to SM_i and $m_{\bar{b}}$ to SM_j .
- 4) During measurement submission Adv as \mathcal{A} receives the individual encrypted measurements including $c_i = E_{pk}(m_B)$ and $c_j = E_{pk}(m_{\bar{B}})$ from smart meters i and j , respectively.
- 5) As DC , Adv can decrypt c_i and set

$$b' = \begin{cases} 0 & \text{if } D_{sk}(c_i) \text{ equals } m_0 \\ 1 & \text{if } D_{sk}(c_i) \text{ equals } m_1 \end{cases}$$

Success Probability: $b = 0 \Leftrightarrow m_0$ is assigned to SM_i in step 3 $\Leftrightarrow D_{sk}(c_i) = D_{sk}(E_{pk}(m_0)) = m_0 \Leftrightarrow b' = 0$, so due to correctness of the underlying encryption scheme, Adv always wins the game, i.e., $Pr[\Gamma_{Adv}^{Unlink}(\lambda) = 1] = 1$. That is, the advantage of the adversary is non-negligible i.e., $\epsilon = 1 - \frac{1}{2} = \frac{1}{2}$. □

B. Privacy Analysis of Aggregation Protocol II

For Protocol II, two colluding sets are considered. In the first case, the data concentrator acts as the sole adversary. This adversary is maximal, as in the second case, it is shown that a single colluding smart meter can help the data concentrator to break the privacy of one honest smart meter in a special case.

1) **Privacy against the data concentrator:** For $Adv = \{DC\}$ (i.e., Adv has no control over any smart meter) and $Ch = \{SM_1, \dots, SM_n\}$, privacy depends on the indistinguishability property of the PRG and is computationally hard.

Theorem 4: Creating the shares for masking using a PRG P , aggregation protocol II provides data unlinkability against the data concentrator.

Proof: Now we consider ADV playing an outside game. In the outer game (which is the indistinguishability of the PRG), ADV is given a vector $\vec{S} = (s_i, s_j)$ of two strings which are either pseudorandom (if $B = 1$) or truly random numbers (if $B = 0$). With ADV playing the role of Ch in the unlinkability game we construct a PPT distinguisher ADV who can distinguish a PRG from a truly random

generator if Adv wins the data unlinkability game with non-negligible advantage. The intuition behind the proof is that if \vec{S} contains pseudorandom numbers, then the inner protocol is the unlinkability game which can be won by Adv .

(i) We consider the following reduction:

- 1) The initialization is executed. The adversary Adv , i.e., the data concentrator, decides on the sending list L .
- 2) Adv outputs a pair of measurements m_0, m_1 .
- 3) A random bit $b \in \{0, 1\}$ is chosen by Ch (who is ADV) who then assigns m_b to SM_i and $m_{\bar{b}}$ to SM_j .
- 4) Adv delivers a random number S_0 . Having used s_i and s_j as the shares for masking the measurements of SM_i and SM_j , respectively, Ch submits all masked measurements. Adv then receives all masked measurements including $\tilde{m}_i = m_b + s_i$, $\tilde{m}_j = m_{\bar{b}} + s_j$ and the sum of masking values $S_N = S_0 + \sum_{t=1}^N s_t$.
- 5) Adv outputs a bit b' . ADV sets $B' = ADV(\vec{S}) = 1 \Leftrightarrow b' = b$, i.e. ADV guesses that he received a pseudorandom string $\vec{S} = (s_i, s_j)$ exactly when Adv wins.

(ii) After the reduction (step 4), Adv has received the same information as in the standalone unlinkability game apart from $S_0 = s_0$ which is only an *independently* created random share.

(iii) Since $B = 1 \Leftrightarrow \vec{S} = (s_i, s_j)$ this case corresponds to the protocol where pseudorandom shares are used. Due to the choice of B' in step 5 and assuming an *advantage* for Adv in winning the unlinkability game yields

$$Pr[B' = 1 | B = 1] = Pr[b' = b | \vec{S} = (s_i, s_j)] = \frac{1}{2} + \text{advantage}$$

If \vec{S} contains truly random values $\vec{S} = (r_i, r_j)$, the adversary Adv does not obtain [15] any information regarding b and hence wins the unlinkability game with probability $\frac{1}{2}$:

$$Pr[B' = 1 | B = 0] = Pr[b = b' | \vec{S} = (r_i, r_j)] = \frac{1}{2}.$$

Combining the preceding two equations we directly obtain the desired result

$$|Pr[B' = 1 | B = 1] - Pr[B' = 1 | B = 0]| = \text{advantage}.$$

Recalling that if *advantage* is non-negligible then ADV distinguishes between a pseudorandom generator and truly random one. Since this contradicts the security of the PRG, we conclude that *advantage* must be negligible. Thus, aggregation protocol II provides smart meters' data unlinkability. □

2) **Privacy violation by the data concentrator and a single smart meter:** For $Adv = \{DC, SM_{N-1}\}$, the privacy of smart meter N can be broken. This can be shown directly using the unlinkability game with $SM_i = SM_N$ as follows.

Theorem 5: Protocol II does not provide smart meters' data unlinkability against the data concentrator and a single smart meter.

Proof: The first 3 steps of the game run unchanged (see Section V-B1).

- 4) During measurement submission, SM_{N-1} being controlled by Adv sends S_{N-1} to SM_N who sends $\tilde{m}_N = m_N + s_N \bmod k$ to DC and also $S_N = S_{N-1} + s_N \bmod k$ to DC . Thus, Adv has \tilde{m}_N, S_{N-1} and S_N and

calculates $\tilde{m}_N - (S_N - S_{N-1}) \bmod k = m_N + s_N - s_N \bmod k = m_N \bmod k = m_N$, since k is chosen bigger than $M = \sum_i m_i$ and therefore also bigger than m_N .

5) *Adv* then sets

$$b' = \begin{cases} 0 & \text{if } m_N = m_0 \\ 1 & \text{if } m_N = m_1 \end{cases}$$

Thus, $b = 0 \Leftrightarrow m_0$ is assigned to $\mathcal{SM}_i = \mathcal{SM}_N$ in step 3 $\Leftrightarrow m_N = m_0 \Leftrightarrow b' = 0$, so $b = b'$, i.e., *Adv* always wins the game. \square

Similarly, it can be shown that $\mathcal{Adv} = \{\mathcal{DC}, \mathcal{SM}_2\}$ can break the privacy of $\mathcal{SM}_i = \mathcal{SM}_1$. Analogously, with two colluding smart meters \mathcal{SM}_{i-1} and \mathcal{SM}_{i+1} (in addition to \mathcal{DC}), the privacy of any smart meter \mathcal{SM}_i can be broken.

VI. VISUALIZATION

In this section, we propose a method to visualize and compare the privacy guarantees of aggregation protocols. We exemplify this visualization by depicting the privacy guarantees of the two aggregation protocols analyzed in this paper and provide an exemplary comparison of the two.

A. Visualization of Privacy Guarantees

In this section, we develop an intuitive graphical way to illustrate the privacy guarantees of aggregation protocols. We explain the visualization and exemplify it for the protocols presented in Section III. This visual representation is an alternate way of presenting the results of the proofs from Section V, but can be read and understood independently. Our visualization enables a quick and easy comparative analysis.

Fig. 5 (right) shows the colors used to illustrate the different privacy levels introduced in Section IV-C. The remainder of the figure (left) illustrates the privacy guarantees for the analyzed protocols from Section V with different maximal collusion sets. For each protocol, different sets (rows) of colluding parties (columns) are shown. Each row illustrates the maximum number of members from each party for which privacy at the respective level (color) is guaranteed for all honest parties.

B. Protocol Comparison based on the Privacy Visualization

For the aggregation protocol using homomorphic encryption (top panel), $N-2$ smart meters and the aggregator may collude (collusion set 2) so that attacks on privacy are polynomially hard (dark gray) as shown for Protocol I in Section V-A1. Alternatively, $N-2$ smart meters and the data concentrator may collude (collusion set 1) to provide information-theoretic privacy guarantees (Section V-A2). If more parties collude, these guarantees cannot be upheld (Section V-A3). Thus, the presentation covers maximal collusion sets. This is shown in Sections V-A1, V-A2 and V-A3.

Similarly, for the aggregation protocol using masking (bottom panel), an adversary controlling $N-1$ smart meters (collusion set 1) still allows for information-theoretic privacy guarantees for the remaining single honest smart meters' data. Intuitively, when \mathcal{DC} is honest, no malicious party gets any

quantity related to the measurement of the honest smart meter (Fig. 2, a formal proof is omitted due to lack of space). This guarantee cannot be upheld if any additional collusions occur. In an alternate adversarial constellation where only the data concentrator is dishonest (collusion set 2), privacy can only be guaranteed against a PPT adversary (Section V-B1), but not against additional dishonest smart meters (Section V-B2).

The visualization of the two aggregation protocols in Fig. 5 makes it easy to compare them in terms of their privacy guarantees. While the protocol relying on homomorphic encryption (top panel) is in any case secure against $N-2$ dishonest smart meters, this protocol also has an additional party, the aggregator, that could be potentially dishonest. In Protocol I, collusion of \mathcal{DC} or \mathcal{A} with $N-2$ smart meters does not break privacy. In contrast, in Protocol II, no aggregator is used, but privacy is broken if \mathcal{DC} and one \mathcal{SM} are dishonest. Thus, in a setting that smart meters are considered trustworthy (i.e., do not collude with \mathcal{DC} or \mathcal{A} , e.g., in common scenarios where smart meters are sealed and tamper resistant), one may prefer protocol II over protocol I. In fact, if we disregard the adversarial control over the smart meters, then Protocol II provides privacy without employing an extra aggregator (which is an advantage over Protocol I which requires an aggregator).

Conversely, if smart meters are more likely to be controlled by the adversary (e.g., when they can be manipulated in an easier way than \mathcal{DC} , which is located in a secured facility outside customer premises), then Protocol I is a better alternative to Protocol II since Protocol I is able to guarantee computational privacy against coalition of $N-2$ smart meters with \mathcal{DC} (or \mathcal{A}), whereas this privacy level is not achievable in Protocol II.

Thus, the visualization allows to draw the following conclusions: Protocol I (using homomorphic encryption) requires an additional aggregator, but is preferred in terms of privacy when smart meters are considered less trustworthy. Conversely, Protocol II (employing masking) is preferred when an extra entity (the aggregator) is undesired.

In summary, the proposed privacy visualization method allows illustrating and comparing the privacy guarantees of different protocols in a compact and graphical way. For each protocol, the maximum collusion sets and the associated privacy level are displayed from highest to lowest privacy level. For convenience, we provide a \LaTeX package to produce privacy visualizations like in Fig. 5. The package can be downloaded from <https://www.en-trust.at/downloads/> and freely used if attribution is given, i.e., this paper is cited.

VII. CONCLUSION AND OUTLOOK

This paper serves as a guideline on how to prove, visualize and compare the privacy guarantees of aggregation protocols. For two example protocols, our approach was applied, i.e., game-based proofs were elaborated and the results were visualized. It was illustrated how different collusions of parties participating in the respective protocols impact privacy and how privacy guarantees vary between protocols. In summary, this paper provides a basis for both protocol designers and implementers to evaluate the privacy impact of different aggregation protocols.

Protocol I	Max. collusion set	$\mathcal{SM}_{1..N}$	\mathcal{DC}	\mathcal{A}
	Collusion set 1	$N - 2$		
	Collusion set 2	$N - 2$		
Protocol II	Max. collusion set	$\mathcal{SM}_{1..N}$	\mathcal{DC}	
	Collusion set 1	$N - 1$		
	Collusion set 2			

Information-theoretic

Computationally hard

Heuristic

Fig. 5. Privacy visualization for the presented protocols: The row color indicates the privacy level achievable when all colored entities of the row collude, e.g., $N - 2$ smart meters and the data concentrator (first row in the first panel). Columns indicate entities and colors privacy levels as illustrated in the legend.

Although this paper focuses on privacy only, it is planned to construct an extended version for security and correctness proofs for energy consumption aggregation protocols so that more powerful adversaries can also be treated.

ACKNOWLEDGMENT

The financial support by the Austrian Federal Ministry of Science, Research and Economy, the Federal State of Salzburg, the Austrian National Foundation for Research, Technology and Development, EU COST Action IC1206, TÜBİTAK (the Scientific and Technological Research Council of Turkey) under project number 115E766, the Turkish Academy of Sciences, and the Royal Society of UK under Newton Advanced Fellowship NA140464 are gratefully acknowledged.

REFERENCES

- [1] E. McKenna, I. Richardson, and M. Thomson, "Smart meter data: Balancing consumer privacy concerns with legitimate applications," *Energy Policy*, vol. 41, pp. 807–814, 2012.
- [2] Z. Erkin and G. Tsudik, "Private Computation of Spatial and Temporal Power Consumption with Smart Meters," in *Proceedings of the 10th international conference on Applied Cryptography and Network Security*, ser. ACNS'12. Berlin Heidelberg: Springer, 2012, pp. 561–577.
- [3] S. B. Wicker and D. E. Schrader, "Privacy-Aware Design Principles for Information Networks," *Proceedings of the IEEE*, vol. 99, no. 2, pp. 330–350, 2011.
- [4] G. Eibl and D. Engel, "Influence of Data Granularity on Smart Meter Privacy," *IEEE Transactions on Smart Grid*, vol. 6, no. 2, pp. 930–939, 2015.
- [5] Z. Erkin, J. R. Troncoso-Pastoriza, R. L. Lagendijk, and F. Perez-Gonzalez, "Privacy-preserving data aggregation in smart metering systems: an overview," *IEEE Signal Processing Magazine*, vol. 30, no. 2, pp. 75–86, mar 2013.
- [6] K. Kursawe, G. Danezis, and M. Kohlweiss, "Privacy-friendly aggregation for the smart grid," in *Privacy Enhancing Technologies Symposium*, 2011, pp. 175–191.
- [7] G. Danezis, C. Fournet, M. Kohlweiss, and S. Zanella-Béguelin, "Smart Meter Aggregation via Secret-sharing," in *Proceedings of the First ACM Workshop on Smart Energy Grid Security (SEGS'13)*. New York, NY, USA: ACM, 2013, pp. 75–80.
- [8] G. Acs and C. Castelluccia, "I have a DREAM! (Differentially private smArt Metering)," in *Proc. Information Hiding Conference*. Prague, Czech Republic: Springer, 2011, pp. 118–132.
- [9] F. Knirsch, "Privacy Enhancing Technologies in the Smart Grid User Domain," *Information Technology, Thematic Issue: Recent Trends in Energy Informatics Research*, vol. 1, no. 59, pp. 13–22, 2017.
- [10] F. Li, B. Luo, and P. Liu, "Secure Information Aggregation for Smart Grids Using Homomorphic Encryption," in *Proceedings of First IEEE International Conference on Smart Grid Communications*. Gaithersburg, Maryland, USA: IEEE, 2010, pp. 327–332.
- [11] G. Danezis, M. Kohlweiss, and A. Rial, "Differentially Private Billing with Rebates," in *International Workshop on Information Hiding*. Berlin Heidelberg: Springer, 2011, pp. 148–162.
- [12] F. Borges and L. A. Martucci, "iKUP Keeps Users' Privacy in the Smart Grid," in *IEEE Conference on Communications and Network Security (CNS)*. IEEE, 2014, pp. 310–318.
- [13] D. Li, Z. Aung, J. J. R. Williams, and A. Sanchez, "Efficient Authentication Scheme for Data Aggregation in Smart Grid with Fault Tolerance and Fault Diagnosis," in *Proceedings of the Innovative Smart Grid Technologies (ISGT)*, 2012, pp. 1–8.
- [14] F. G. Marmol, C. Sorge, R. Petrlc, O. Ugus, D. Westhoff, and G. Martinez Perez, "Privacy-enhanced architecture for smart metering," *International Journal of Information Security*, vol. 12, no. 2, pp. 67–82, 2013.
- [15] C. Castelluccia, A. C. Chan, E. Mykletun, and G. Tsudik, "Efficient and Provably Secure Aggregation of Encrypted Data in Wireless Sensor Networks," *ACM Transactions on Sensor Networks (TOSN)*, vol. 5, no. 3, pp. 20:1–20:36, 2009.
- [16] J.-M. Bohli, C. Sorge, and O. Ugus, "A Privacy Model for Smart Metering," in *2010 IEEE International Conference on Communications Workshops (ICC)*. IEEE, 2010, pp. 1–5.
- [17] A. Rial and G. Danezis, "Privacy-preserving smart metering," in *Proceedings of the 10th annual ACM workshop on Privacy in the electronic society*. ACM, 2011, pp. 49–60.
- [18] S. T. Boshrooyeh, A. Küpcü, and Ö. Özkasap, "PPAD: Privacy Preserving Group-Based ADvertising in Online Social Networks," in *IFIPNetworking 2018*. Zurich, Switzerland: IEEE, 2018, pp. 541–549.
- [19] V. Daza, J. Domingo-Ferrer, F. Sebé, and A. Viejo, "Trustworthy privacy-preserving car-generated announcements in vehicular ad hoc networks," *IEEE Transactions on Vehicular Technology*, vol. 58, no. 4, pp. 1876–1886, 2009.
- [20] N. Buescher, S. Boukoro, S. Bauregger, and S. Katzenbeisser, "Two Is Not Enough: Privacy Assessment of Aggregation Schemes in Smart Metering," *Proceedings on Privacy Enhancing Technologies*, vol. 2017, no. 4, pp. 118–134, 2017.
- [21] A. J. Menezes, P. C. van Oorschot, and S. A. Vanstone, *Handbook of Applied Cryptography*, 5th ed. CRC Press, 2001.
- [22] Y. Guo, C. W. Ten, and P. Jirutitijaroen, "Online data validation for distribution operations against cyber tampering," *IEEE Transactions on Power Systems*, vol. 29, no. 2, pp. 550–560, 2014.
- [23] P. Paillier, "Public-Key Cryptosystems Based on Composite Degree Residuosity Classes," in *Advances in Cryptology — EUROCRYPT '99*. Berlin Heidelberg: Springer, 1999, vol. 1592, pp. 223–238.
- [24] D. Engel, "Wavelet-based Load Profile Representation for Smart Meter Privacy," in *Proc. IEEE PES Innovative Smart Grid Technologies (ISGT'13)*, Washington, D.C., USA, 2013, pp. 1–6.
- [25] Z. Erkin, "Private Data Aggregation with Groups for Smart Grids in a Dynamic Setting using CRT," in *2015 IEEE International Workshop on Information Forensics and Security (WIFS)*, Rome, Italy, 2015.
- [26] E. Barker, W. Barker, W. Burr, W. Polk, M. Smid, and C. S. Division, "NIST 800-57: Computer Security," pp. 1–147, 2012.
- [27] F. Knirsch, G. Eibl, and D. Engel, "Error-resilient Masking Approaches for Privacy Preserving Data Aggregation," *IEEE Transactions on Smart Grid*, vol. 9, no. 4, pp. 3351–3361, 2018.
- [28] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of cryptography conference*. Springer Berlin Heidelberg, 2006, pp. 265–284.
- [29] E. Barker and A. Roginsky, "Transitions: Recommendation for transitioning the use of cryptographic algorithms and key lengths," *NIST Special Publication*, vol. 800, p. 131A, 2011.
- [30] J. Katz and Y. Lindell, *Introduction to Modern Cryptography*, 1st ed. Chapman & Hall/CRC, 2007.
- [31] A. Shamir, "How to Share a Secret," *Communications of the ACM (CACM)*, vol. 22, no. 11, pp. 612–613, 1979.
- [32] A. Barker, "NIST Special Publication 800-57: Recommendation for Key Management - Part 1: General (Revised)," 2016. [Online]. Available: http://csrc.nist.gov/publications/nistpubs/800-57/sp800-57-Part1-revised2_Mar08-2007.pdf