

Evaluation of the Precision-Privacy Tradeoff of Data Perturbation for Smart Metering

Marco Savi, Cristina Rottondi and Giacomo Verticale

Abstract—Smart grid users and standardization committees require that utilities and third parties collecting metering data employ techniques for limiting the level of precision of the gathered household measurements to a granularity no finer than what is required for providing the expected service. Data aggregation and data perturbation are two such techniques. This paper provides quantitative means to identify a tradeoff between the aggregation set size, the precision on the aggregated measurements, and the privacy level. This is achieved by formally defining an attack to the privacy of an individual user and calculating how much its success probability is reduced by applying data perturbation. Under the assumption of time-correlation of the measurements, colored noise can be used to reduce even further the success probability. The tightness of the analytical results is evaluated by comparing them to experimental data.

I. INTRODUCTION

According to the novel Smart Grid paradigm, Smart Meters are connected to the communication network of the electricity grid through the Automatic Metering Infrastructure (AMI) and send their measurements about energy consumption to the power suppliers or other External Entities (EEs) such as grid managers or third party service providers. There is ample literature showing that electricity usage readings leak users' sensitive information: it has been shown [1], [2] that external subjects accessing these data might infer private informations about the users by exploiting the customers' behavior and even to determine which household appliances are being used during a given time span. This could potentially lead to threatening consequences: for example, burglars could detect whether houses are unoccupied before attempting burglaries, vendors could select potential targets for their marketing campaigns, insurances could infer the health status or the likelihood for an individual to cause accidents at home.

Therefore, both users and standardization bodies [3], [4] require that a secure and privacy-friendly collection framework is implemented to ensure privacy and confidentiality of data collected from the Smart Meters.

There is a significant body of work discussing how individual data can be aggregated in a cryptographically secure way so that entities providing ancillary services can only access to a function, generally the sum, of the measurements gathered from several meters (see, for example, [5], [6], [7], [8]). None of these works, however, answers a fundamental question: how large the aggregation set must be in order to ensure privacy? Clearly, if the aggregation set is too small, though an observer

will not have access to the exact measurement from each meter, he will be able to make reasonable guesses about the usage of specific appliances. These guesses can grow more and more precise as the observation interval becomes longer. Further, in order to give an estimate of the minimal aggregation set size for guaranteeing privacy, one must make assumptions on the behavior of the other users, which are hard to make and justify.

The second most popular approach for providing privacy to metering data is to modify the meter readings in order to conceal the fingerprints of appliances, thus making deanonymization or non-intrusive load monitoring less effective. This concealment is achieved either by using batteries or by directly adding random noise to the measurements (see, for example, [9], [10], [11]). This approach has the drawback of distorting the metering results, which will thus become less useful also for the legitimate recipient. It is worth noting, however, that the two above mentioned approaches can be combined. On one hand, the addition of random noise to an aggregate measurement makes the privacy of each individual independent of the behavior of other house dwellers. On the other hand, the aggregation of multiple meters makes it possible to reduce the impact of the added noise, resulting in more precise aggregate measurements.

Dimensioning the noise to be added to the aggregate is a crucial issue to make privacy-friendly approaches feasible for deployment in a large-scale grid. Some researchers have tackled this problem by using differential privacy [12], utility function [13], or information-theoretic metrics [11]. We introduce a definition of privacy that makes it easy to calculate the amount of noise necessary to ensure privacy and explicitly consider the beneficial impact of aggregation.

The contributions of this paper are the following:

- The formalization of the notion of ϵ -Privacy as a measure of the privacy of the users generating the measurements;
- The analytical derivation of the ϵ parameter for privacy guarantees, under the assumption that additive Gaussian noise is used for data perturbation;
- The evaluation of how privacy can be improved by employing suitably colored noise.

The derived results are optimal in the sense that no attack to privacy can further increase the value of ϵ without additional knowledge. These result are extremely useful in the design of Smart Metering deployments, which are expected to run for several years. Therefore possible privacy risks must be anticipated and tackled with, even without precise knowledge of the attackers' strategies. We prove the effectiveness of our proposal by evaluating its performance with real Smart Meters measurements from the Smart* dataset [14].

A preliminary version of this paper appeared in C. Rottondi, M. Savi, D. Polenghi, G. Verticale and C. Krauß, "A Decisional Attack to Privacy-friendly data aggregation in Smart Grids", IEEE Globecom 2013 - Symposium on Selected Areas in Communications (GC13 SAC), Dec. 2013

The remainder of the paper is structured as follows: Section II provides an overview of the related work. Our proposed data aggregation architecture is discussed in Section III. The formalization of a Privacy Challenge and of the resulting attack scenario are discussed in Section IV. Section V reports a theoretical analysis of the probability of success of such attack for two different characterizations of the injected noise, whereas Section VI shows numerical results obtained using real measurement traces. Section VII concludes the paper.

II. RELATED WORK

Our attack scenario builds upon the notion of differential privacy, which was first introduced by Dwork *et al.* in [15] and aims at guaranteeing that the removal or addition of a single item in a statistical database has negligible impact on the outcome of any query on that database. The authors give a formal definition of differential privacy as a measure of the tradeoff between the precision of the aggregate data and the probability of identifying the contributions of individual data inside the aggregate. Moreover, the authors describe how to achieve a target level of differential privacy by means of noise injection in the users' data. Our attack scenario for time series is based on the same principle. However, our approach is more focused on the specific characteristics of Smart Grid time series, resulting in simpler definitions.

The same authors of [15] present in [16] some applications of differential privacy to statistical data inference and learning theory, while in [17] they extend the study to the statistical characterization of the noise to be injected in general query functions, proving that privacy can be preserved by calibrating the standard deviation of the noise according to the function sensitivity.

Some other papers combine cryptographic schemes with differential privacy techniques in order to compute aggregate statistics: in [18], Dwork *et al.* propose a protocol for the distributed generation of random noise, aimed at the distributed implementation of privacy-preserving statistical databases by means of a verifiable secret sharing scheme.

Danezis *et al.* [19] apply differential privacy to a billing protocol to evaluate the monetary amount that customers should add to their bill in order to provably hide their activities, and proposes a cryptographic protocol for oblivious billing which imposes no additional expenditure.

Rastogi *et al.* [20] design a protocol for differentially private aggregation of temporally correlated time series, which is achieved by perturbation of the Discrete Fourier Transform of the data and by distributed noise addition. The protocol scales efficiently with the number of users, since it requires a computational load per user of $O(1)$. Our solution also relies on noise addition, which is performed directly on the individual metering data.

Some works [21], [12], [13], [22], [23], [24] apply the general notions expressed in [15] to the Smart Grid context. Jelasyt *et al.* [23] extend the application of differential privacy to streams of time-consecutive queries and analytically characterize the power to be injected in order to guarantee the privacy of individual measurements, when observed for

an unlimited amount of time. In our paper, we also evaluate the impact of the time observation intervals on the achieved privacy level. Acs *et al.* [12] define a scheme in which an electricity supplier is allowed to collect aggregate smart-metering measurements without learning anything about the energy consumption and the household activities of individual users, and discuss how differential privacy is affected by considering multiple time slots. However, this paper does not deal with temporal correlation of metering time series. Conversely, our proposal considers this feature, which can be exploited to reduce the level of privacy of the users' data.

Chan *et al.* [22] deal with a scenario in which an untrusted aggregator collects user data to periodically compute aggregate statistics. The proposed solution is resilient to users' failures and compromises and supports dynamic joins and leaves. We also assume untrusted aggregation nodes, but we focus our attention on a static scenario.

Shi *et al.* [21] define a model of data aggregator capable of obtaining statistics about aggregate data without compromising the privacy of single users. The authors introduce a formal noise injection model and a new distributed data randomization algorithm in order to ensure users' differential privacy, assuming the existence of malicious users that reveal their statistics to the data aggregator. Moreover, the authors define an error bound for aggregated data and evaluate the tradeoff between data utility and privacy. The same tradeoff evaluation is discussed by Rajagopalan *et al.* [13], who propose to filter low-power frequency components of smart-metering time series, in order to perform data obfuscation without compromising its statistical significance.

Zhao *et al.* [24] combine differential privacy with a battery-based load scheduling algorithm. However, such approach can be applied only in case the households are equipped with storage devices, which could lead to considerable installation and maintenance costs. Our proposed framework does not require any storage equipment.

In [25], Zhang *et al.* propose a new noise model to achieve data perturbation. The noise is correlated to time series in order to increase users' privacy without compromising data utility by means of a noise generation algorithm called PESP (Privacy Enhanced State-dependent Perturbation).

Although our work builds upon the notion of differential privacy, it differs from the previously mentioned approaches because it introduces a new definition of noise injection. In particular, Acs *et al.* [12] consider a Laplace distributed noise, whereas Shi *et al.* [21] consider a noise with geometric symmetric distribution. Such models are chosen to fit the differential privacy analytical results obtained in [15]. Conversely, we consider Gaussian noise, which allows for a comprehensive theoretical analysis in support to numerical results.

All the noise processes considered by the previous literature in the context of differential privacy are white noise processes, i.e., they refer to statistically uncorrelated sequences. In our work, referring to the approach proposed in [23], [25], we extend data perturbation to colored (i.e., correlated) noise sequences and study the assumptions under which colored noise leads to better performance in terms of privacy preservation with respect to white noise.

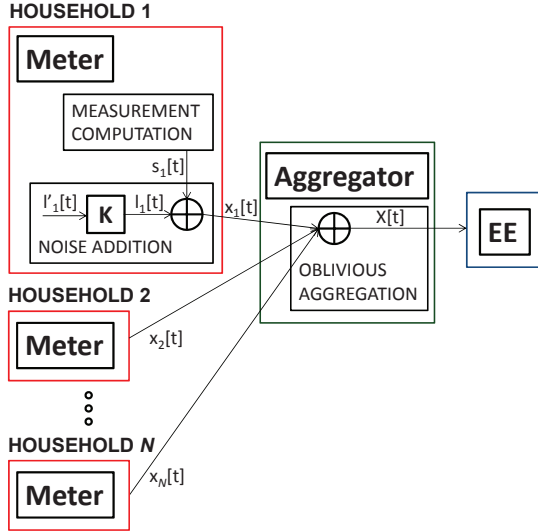


Fig. 1. The data perturbation and aggregation procedure

III. THE AGGREGATION ARCHITECTURE

We consider an AMI network with N Meters, each one conveying its measurements to an Aggregator, which is responsible for the data aggregation and for communicating the final aggregate to an EE. At each round t , the Meters generate the measurements $s_i[t]$, where i is the Meter number, add a perturbation noise $l_i[t]$ and deliver them to an Aggregator. The Aggregator sums the individual perturbed measurements $x_i[t] = s_i[t] + l_i[t]$ and delivers to the interested EE the quantity:

$$X[t] = \sum_{i=1}^N x_i[t] = \sum_{i=1}^N s_i[t] + \sum_{i=1}^N l_i[t] = S[t] + L[t]$$

where $S[t] = \sum_{i=1}^N s_i[t]$ and $L[t] = \sum_{i=1}^N l_i[t]$. Note that in the remainder of the paper, referring to temporal measurements, we will use upper case letters to indicate aggregate values and lower case letters to refer to individual values. A pictorial view of our considered scenario is proposed in Figure 1.

The noise $l_i[t]$ is a *Gaussian colored noise* with zero-mean and variance σ_i^2 obtained by filtering a white Gaussian noise $l'_i[t]$ with a Linear Time Invariant (LTI) filter K . The output of the aggregation procedure leads to an aggregate zero-mean Gaussian colored noise $L[t]$ with variance $\sigma_L^2 = N\sigma_i^2$. Note that this is equivalent to filtering the aggregate white noise $L'[t] = \sum_{i=1}^N l'_i[t]$ (having variance $\sigma_{L'}^2$), with the filter K . In case of *Gaussian white noise* addition, the filter K is absent and $L[t] = L'[t]$. The Meters, the Aggregator and the EE execute a cryptographic protocol that makes operations oblivious (e.g. as the one proposed in [26]), meaning that the Aggregator node can perform additions on the measurements, but cannot access the individual inputs nor the generated noise. Depending on the available technology and computational power at the Meters, the Aggregator can be implemented either centralized [21] or distributed [26].

A well designed system should provide a low σ_L^2 , while ensuring a required level of privacy. Some categories of EEs,

e.g. the billing service, impose strict tolerance thresholds on σ_L^2 . Other categories of EEs, such as operators on the wholesale market who need to collect aggregate metering data for statistical purposes, may tolerate much higher values of σ_L^2 .

IV. ADVERSARY AND PRIVACY MODEL

A. Adversarial Model

We assume that both the Aggregator and the EE behave according to the *honest-but-curious* attacker model, i.e., they correctly execute the protocol, but store all their inputs and process them in order to obtain additional information about the individual measurements $s_i[t]$. These nodes are expected to use an optimal algorithm to achieve their goal. Moreover, we assume that a secure communication channel is available between the nodes and, thus, the system is protected against eavesdroppers and other external attackers. In addition, the EE can collude with one or more meters in order to obtain the individual measurements of another meter. By collusion, we mean either that those meters are compromised, or simply that their behavior is very easy to predict, which is not uncommon in this scenario.

B. Privacy Model

The Privacy Challenge described below leverages notions from modern cryptography and differential privacy. We start by observing that we cannot easily give a mathematical definition of privacy. However, we can say that an aggregate measurement that does not include data from user a clearly respects the privacy of a to the maximum possible extent, since a is not even present. Now, suppose that the curious attacker is shown two noiseless aggregates. One includes data from a and the other one does not. Instead of a , the aggregate contains data from a user randomly chosen from the same population of a . All other users in the aggregate behave the same. If the attacker is able of guessing even a single bit of information about a , then it can easily tell in which aggregate a is.

Now, suppose that the curious attacker is shown the same two noiseless aggregates, but each aggregate is added to a different noise. Even if the attacker has some information about a , it can be difficult to guess in which aggregate a is, because of the injection of random noise. Even further, if the attacker has complete knowledge about a , it may not be able to guess because of noise. Therefore, we measure the loss of privacy of a as the ability of an observer to tell apart the behavior of a from the behavior of another user from the same population. More formally, we will say that the aggregation system is privacy-friendly if the probability that the Attacker (or Adversary) makes a correct guess is close to 1/2. A significant deviation from 1/2 means that the system is not very privacy-friendly.

Definition 1 (Privacy Challenge). *The challenger computes the noisy aggregate time series $X_a[t]$ and $X_b[t]$, within the*




Fig. 2. Definition of the Privacy Challenge

observation window $0 \leq t \leq N_s - 1$, calculated over N participants as follows:

$$X_a[t] = \sum_{\substack{1 \leq i \leq N \\ i \neq a, i \neq b}} s_i[t] + s_a[t] + L_a[t]$$

$$X_b[t] = \sum_{\substack{1 \leq i \leq N \\ i \neq a, i \neq b}} s_i[t] + s_b[t] + L_b[t]$$

The challenger sends the sequences $X_a[t]$ and $X_b[t]$ in a random order and the smart metering data $s_a[t]$ for $0 \leq t \leq N_s - 1$. The Adversary wins the challenge if it guesses whether the user a participates in the noisy aggregate time series $X_a[t]$ or $X_b[t]$.

Definition 2 (ε -Privacy). *The aggregation architecture provides ε -Privacy if no decision algorithm can give the correct answer to the Privacy Challenge with probability larger than $1/2 + \varepsilon$ for any pair of users (a, b) in the set.*

Note that the two aggregates $X_a[t]$ and $X_b[t]$ differ only by a single participant: the former contains the measurements $s_a[t]$, whereas in the latter the measurement of user a has been replaced by the measurement $s_b[t]$ of another user b from the same population. The optimal strategy for the Adversary is applying a *decision algorithm* that calculates the correlation between the time series $s_a[t]$ and $X_a[t]$ and between $s_a[t]$ and $X_b[t]$ as follows:

$$R_{s_a, X_a} = \sum_{t=0}^{N_s-1} s_a[t] X_a[t] \quad (1)$$

$$R_{s_a, X_b} = \sum_{t=0}^{N_s-1} s_a[t] X_b[t] \quad (2)$$

The Adversary chooses the noisy aggregate time series that results in the highest correlation.

Clearly, the higher is the noise variance σ_L^2 , the smaller is the difference between R_{s_a, X_a} and R_{s_a, X_b} , thus making the probability of correct guess approach $1/2$. In other words, defining:

$$R = R_{s_a, X_a} - R_{s_a, X_b}$$

we can say that:

- The Adversary makes the right choice when $R > 0$;
- The Adversary makes the wrong choice when $R \leq 0$.

Considering the linearity property of the operation of correlation, the decision algorithm assumes that the inequality $R_{s_a, s_a} > R_{s_a, s_b}$ always holds. This means that the Adversary

TABLE I
MAIN SYMBOLS

Sequence	Definition	Characterization
$s_a[t]$	Measurements of user a	Deterministic
$s_b[t]$	Measurements of user b	Deterministic
$S_a[t]$	Measurements of aggregate a	Deterministic
$S_b[t]$	Measurements of aggregate b	Deterministic
$L_a[t]$	Gaussian noise process for aggregate a	Stochastic
$L_b[t]$	Gaussian noise process for aggregate b	Stochastic
$X_a[t]$	$S_a[t] + L_a[t]$	Stochastic
$X_b[t]$	$S_b[t] + L_b[t]$	Stochastic

always assumes that the correlation of $s_a[t]$ with itself is greater than the correlation between $s_a[t]$ and $s_b[t]$.

It is worth noting that $s_a[-t]$ can be seen as the impulse response of a LTI filter named *correlation filter*. Then, $\mathcal{F}\{s_a[-t]\} = S_a^*(\varphi)$ is the frequency response of the correlation filter, where \mathcal{F} indicates the Discrete-Time Fourier Transform (DTFT) and $0 \leq \varphi \leq 1$ indicates the normalized frequency. This expression is analogous to a Matched Filter, which is optimal in case of white noise injection [27].

It is important to note that the attacker is defined to use the optimal strategy in presence of Additive White Gaussian Noise (AWGN), therefore if a system can be proved to be privacy-friendly under this strategy, then it is resistant to the whole wide family of deanonymization or load-monitoring attempts such as the attacks in [28] and [9], to name a few of the documented ones, and even to any other attack not yet described belonging to the same family. In fact, if any of these attacks could extract any information from the aggregate, then these pieces of information could be used to win the challenge. As a consequence, if the success of the challenge can be proved to be close to $1/2$, that means that no attack can extract useful information.

In the following Section we provide a theoretical analysis of the Privacy Challenge considering two different case studies, which differ in the characterization of noise injection and correlation.

V. ANALYTICAL RESULTS

Table I lists the main symbols used in the reminder of the Section. Note that the two processes $X_a[t]$ and $X_b[t]$ are Gaussian, since they are obtained by summing deterministic signal $S_a[t]$ (resp. $S_b[t]$) to a Gaussian process $L_a[t]$ (resp. $L_b[t]$). Therefore, the quantities R_{s_a, X_a} and R_{s_a, X_b} defined in Section IV, as well as their difference R , can be statistically characterized as Gaussian random variables. We derive the following theorems.

Theorem 1 (White noise). *If $L[t]$ is a Gaussian white noise with zero-mean and variance σ_L^2 , then the aggregation architecture provides ε -Privacy with*

$$\varepsilon = \left| \frac{1}{2} \operatorname{erfc} \left(- \frac{\sum_{t=0}^{N_s-1} (s_a[t]^2 - s_a[t]s_b[t])}{2\sigma_L \sqrt{\sum_{t=0}^{N_s-1} s_a[t]^2}} \right) - \frac{1}{2} \right| \quad (3)$$

with a and b chosen from the population so that ε is maximum.

The proof is provided in Appendix A.

Theorem 2 (Colored noise). *Let $L[t]$ be a Gaussian colored noise with zero-mean and Power Spectral Density $\text{PSD}_L(\varphi) = \sigma_L^2 |\mathbf{K}(\varphi)|^2$, with $\int_0^1 |\mathbf{K}(\varphi)|^2 d\varphi = 1$. Then, the aggregation architecture provides ε -Privacy with*

$$\varepsilon = \left| \frac{1}{2} \operatorname{erfc} \left(-\frac{\sum_{t=0}^{N_s-1} (s_a[t]^2 - s_a[t]s_b[t])}{2\sigma_L \sqrt{\int_0^1 |\mathbf{K}(\varphi)|^2 |\mathbf{S}_a(\varphi)|^2 d\varphi}} \right) - \frac{1}{2} \right| \quad (4)$$

with a and b chosen from the population so that ε is maximum.

The proof is provided in Appendix B.

The design of the filter K in case of colored noise addition is an important issue to address, since the energy spectrum $|\mathbf{K}(\varphi)|^2$ influences the value of the denominator in Equation (4). Specifically, the parameter ε decreases when the denominator increases, thus providing a better ε -Privacy for the users.

Note that the filter K is normalized, i.e., it has unitary energy ($\int_0^1 |\mathbf{K}(\varphi)|^2 d\varphi = 1$). This way, we guarantee that the variance of the additive noise is the same both for white and colored noise, since:

$$\sigma_L^2 = \int_0^1 \text{PSD}_L(\varphi) d\varphi = \sigma_L^2 \int_0^1 |\mathbf{K}(\varphi)|^2 d\varphi = \sigma_L^2,$$

VI. PERFORMANCE EVALUATION

We consider real home energy consumption traces taken from the SMART* dataset [14]. The main parameters of the dataset are summarized in Table II. The dataset consists in power consumption traces of 400 different homes sampled every minute for a period of 24 hours.

The privacy preserving framework is agnostic with respect to the purposes of the metering data collection. Different applications have different requirements in terms of data accuracy. The accuracy of the aggregated data is quantified in terms of a *Perturbation Coefficient* ψ defined as $\psi = \sigma_L / (NP_{\text{ave}})$, where σ_L is the standard deviation of the perturbation noise, N is the number of aggregated users, P_{ave} is the per-user average power consumption. The Perturbation Coefficient is a tunable parameter that can be set to different values depending on the specific application: the lower the Perturbation Coefficient, the higher the accuracy of the data and, on the other hand the worse the privacy. It is worth noting that several commonly used metering devices show an error in the order of $\pm 2\%$ (see, for example, [29]). Therefore, we expect that an aggregate error of about $\psi = 10^{-2}$ is compatible with the tolerance of most of the commonly deployed smart grid ancillary services (see [30] for a reference list). However, some applications may tolerate higher values of ψ . In this case, the size of the

TABLE II
SMART* DATASET [14] DESCRIPTION

Parameter	Description	Value
N_s	Number of samples per trace	1440
T_s	Sampling period	1 min
T_{obs}	Duration	24 h
P_{ave}	Average power	1.22 kW
P_{min}	Minimum power	0.089 kW
P_{max}	Peak power	14.69 kW
P_{std}	Std. Dev. of power per trace	0.91 kW




Fig. 3. Value of the privacy parameter, ε , versus the size of the aggregation set N both in case of analytical evaluation and simulation results ($T_s = 15$ min, $T_{\text{obs}} = 24$ h, $\psi = 10^{-2}$)

aggregation set required to achieve a target privacy level can be reduced (see Figures 3 and 4 and related comments).

For the considered traces, where $P_{\text{ave}} = 1.22$ kW (see Table II), $\psi = 10^{-2}$ corresponds to a noise with standard deviation $\sigma_l = 12.2$ W per meter.

A. Comparison with simulation results

We consider a scenario in which the EE knows the aggregate noisy time series $X_a[t]$, $X_b[t]$ and the corresponding samples of the time series $s_a[t]$. We have re-sampled the traces with a sampling period of $T_s = 15$ min, which is more in line with the current smart-metering technical specifications, thus each time series is composed of $N_s = T_{\text{obs}}/T_s = 96$ samples.

We choose to design the filter K for noise coloring as similarly as possible to the estimated spectral characterization of the Smart Meter measurements $s_i[t]$. This way, it is possible to hide the additive noise and to reduce the effectiveness of the Matched Filter used to exploit the decision algorithm described in Section IV. Hence, the chosen energy spectrum of K has the following squared absolute magnitude:

$$|\mathbf{K}(\varphi)|^2 = \frac{|\widetilde{\mathbf{S}(\varphi)}|^2}{\int_0^1 |\widetilde{\mathbf{S}(\varphi)}|^2 d\varphi} \quad (5)$$

where $|\widetilde{\mathbf{S}(\varphi)}|^2$ is the average spectrum of the consumption traces of the dataset.

Simulation results are obtained by randomly choosing pairs of meters, calculating the challenge sets according to Definition 1 (which includes the generation of random noise), then calculating the correlations using (1) and (2). Finally, the adversary chooses the aggregate with the largest correlation. Analytical results are obtained by applying (3) and (4).

Figure 3 shows the resulting Privacy, ε , versus the size of the aggregate set, N , for both white and colored noise. Results have been obtained both analytically and by means of simulations.

Simulation results closely approach the analytical estimations for both white and colored noise. Simulations also




Fig. 4. Value of the privacy parameter, ϵ , versus the size of the aggregation set N both for *white* and *colored* noise addition considering different values of the Perturbation Coefficient ψ ($T_s = 15$ min, $T_{obs} = 24$ h)

confirm that, for given values of ψ and N , the usage of colored noise allows for a consistent improvement of the user privacy with respect to white noise. Conversely, for a target level of privacy and a given value of ψ , colored noise allows for a reduction of the number of users N . For example, if we want to guarantee that $\epsilon < 10^{-2}$, in case of white noise addition the aggregation set must contain more than $N = 65000$ users, while in case of colored noise the aggregation set size can be reduced to about $N = 8000$. It is worth noting that our proposed aggregation mechanism works for any arbitrary choice of N , at the price of lowering the achieved privacy level. Therefore, the remainder of the Section thoroughly investigates the tradeoff between aggregation set size, data precision and privacy level.

B. Impact of tunable parameters

Figure 4 shows the privacy parameter, ϵ , versus the aggregation set size, for various values of the Perturbation Coefficient, ψ , for both white and colored perturbation noise.

The graph shows that decreasing the value of ψ (i.e., increasing data precision) leads to a worse ϵ -Privacy, both in case of white and colored noise.

It also shows that the value of ϵ guaranteed by white noise and $\psi = 10^{-1}$ is comparable to the one provided by colored noise and $\psi = 10^{-3}$. This means that colored noise makes it possible to use a lower noise variance to achieve a target privacy level, thus leading to a higher precision of the aggregated data, while maintaining the same size of the aggregation set.

Figure 5 shows how the ϵ -Privacy of the users is affected by the length of the observation windows T_{obs} . In case of white noise addition, increasing the observation window T_{obs} (i.e., expanding the time period during which the EE collects the aggregated data) leads to a smaller ϵ -Privacy for the users, since the adversary obtains more information about the correlation between the samples of the traces.

Conversely, in case of colored noise, increasing T_{obs} does not significantly affect the ϵ -Privacy. This happens because




Fig. 5. Value of the privacy parameter, ϵ , versus the size of the aggregation set N in case of *white* noise or *colored* noise addition for different values of the observation interval T_{obs} ($T_s = 15$ min, $\psi = 10^{-2}$)




Fig. 6. Value of the privacy parameter, ϵ , versus the size of the aggregation set N in case of *white* noise or *colored* noise addition for different values of the sampling period T_s ($T_{obs} = 24$ h, $\psi = 10^{-2}$)

the coloring filter K is specifically designed to make the perturbation noise match the signal energy spectrum, thus making the noise more difficult to be removed (see (5)).

Figure 6 shows how the ϵ -Privacy of the users is affected by the duration of the measurement sampling period T_s . In case of white noise addition, decreasing T_s (i.e., providing the EE with aggregated data of finer granularity) leads to a lower ϵ -Privacy for the users, since the adversary can exploit more information about the correlation of the time series measurements. Conversely, in case of colored noise, lowering T_s does not significantly affect the user ϵ -Privacy.

VII. CONCLUSION

This paper defines the concept of ϵ -Privacy, which makes it possible to evaluate the privacy guarantees for the users of a Smart Grid aggregation architecture as a function of the perturbation noise added to the Smart Meter aggregate measurements. We provide formulas for calculating the achieved ϵ -Privacy in case of Gaussian white and colored noise. We show with simulations that noise addition can achieve arbitrary levels of ϵ -Privacy. Further, by suitably coloring noise, it is possible to achieve the same level of privacy with a much lower noise variance and, thus, higher data precision.




Fig. 7. Theorem 1: The adversary's decision algorithm in case of *white noise* addition

We also investigate how various system parameters, such as the aggregation group size, the observation interval and the sampling time, impact on the tradeoff between data precision and user privacy. The results are optimal, meaning that no attack to privacy can be more effective than what is predicted by using this framework.

ACKNOWLEDGMENTS

The work in this paper has been partially funded by Regione Lombardia under grant no. 40545387 Smart Campus as Urban Open LABs (SCUOLA).

APPENDIX A PROOF OF THEOREM 1

Proof: We want to compute:

$$\varepsilon = \Pr \{R > 0\} - 1/2 \quad (6)$$

To do so, we need to calculate mean and variance of the random variable $R = R_{s_a, X_a} - R_{s_a, X_b}$ (see Figure 7). The correlation R_{s_a, X_a} can be expressed as:

$$R_{s_a, X_a} = R_{s_a, S_a} + R_{s_a, L_a}$$

The first term can be further rewritten as follows:

$$\begin{aligned} R_{s_a, S_a} &= R_{s_a, s_a} + R_{s_a, S_a - s_a} \\ &= \sum_{t=0}^{N_s-1} s_a[t]^2 + \sum_{\substack{1 \leq i \leq N \\ i \neq a, i \neq b}} \sum_{t=0}^{N_s-1} s_a[t] s_i[t] \end{aligned}$$

Moreover, R_{s_a, L_a} is a Gaussian random variable, obtained by filtering the Gaussian white random process $L_a[t]$ with the LTI correlation filter and sampling the output in $\tau = 0$, as depicted in Figure 7, according to the Matched Filter theory [27]. Considering that $s_a[t]$ is a real-valued time series and thus $|s_a[t]|^2 = s_a[t]^2$, by exploiting the Parseval's Theorem, the probability distribution of R_{s_a, L_a} is:

$$R_{s_a, L_a} \sim \mathcal{N} \left(0, \sigma_L^2 \sum_{t=0}^{N_s-1} s_a[t]^2 \right)$$

where the notation $X \sim \mathcal{N}(\mu, \sigma^2)$ means that the random variable X is normally distributed with mean μ and variance

σ^2 . It follows that R_{s_a, X_a} is also a Gaussian random variable defined as:

$$R_{s_a, X_a} \sim \mathcal{N} \left(\sum_{t=0}^{N_s-1} s_a[t]^2 + \sum_{\substack{1 \leq i \leq N \\ i \neq a, i \neq b}} \sum_{t=0}^{N_s-1} s_a[t] s_i[t], \sigma_L^2 \sum_{t=0}^{N_s-1} s_a[t]^2 \right)$$

The same considerations hold for R_{s_a, X_b} , therefore:

$$R_{s_a, X_b} \sim \mathcal{N} \left(\sum_{t=0}^{N_s-1} s_a[t] s_b[t] + \sum_{\substack{1 \leq i \leq N \\ i \neq a, i \neq b}} \sum_{t=0}^{N_s-1} s_a[t] s_i[t], \sigma_L^2 \sum_{t=0}^{N_s-1} s_a[t]^2 \right)$$

Thus, R is a Gaussian random variable with mean μ_R and variance σ_R^2 defined as follows:

$$\mu_R = \sum_{t=0}^{N_s-1} s_a[t]^2 - \sum_{t=0}^{N_s-1} s_a[t] s_b[t], \quad (7)$$

$$\sigma_R^2 = 2\sigma_L^2 \sum_{t=0}^{N_s-1} s_a[t]^2 \quad (8)$$

Therefore, ε can be computed as:

$$\varepsilon = \Pr \{R > 0\} - \frac{1}{2} = \frac{1}{2} \operatorname{erfc} \left(-\frac{\mu_R}{\sigma_R \sqrt{2}} \right) - \frac{1}{2} \quad (9)$$

By substituting (7) and (8) into (9), we obtain:

$$\varepsilon = \frac{1}{2} \operatorname{erfc} \left(-\frac{\sum_{t=0}^{N_s-1} (s_a[t]^2 - s_a[t] s_b[t])}{2\sigma_L \sqrt{\sum_{t=0}^{N_s-1} s_a[t]^2}} \right) - \frac{1}{2}$$

Note that ε is positive (negative) iff μ_R is positive (negative). If $\varepsilon > 0$, the *decision algorithm* defined in Section IV is the best algorithm the adversary could apply, since the assumption $R_{s_a, s_a} > R_{s_a, s_b}$ holds. If this assumption does not hold, i.e., $\varepsilon \leq 0$, it is possible to define a *dual decision algorithm*. Considering such algorithm, the decider makes the right (wrong) choice when $R \leq 0$ ($R > 0$). Therefore, by exploiting the dual decisional algorithm, the adversary increases its probability of correct guess with respect to a coin toss also when $\mu_R < 0$. Hence, we can eventually calculate ε as:

$$\varepsilon = \left| \frac{1}{2} \operatorname{erfc} \left(-\frac{\sum_{t=0}^{N_s-1} (s_a[t]^2 - s_a[t] s_b[t])}{2\sigma_L \sqrt{\sum_{t=0}^{N_s-1} s_a[t]^2}} \right) - \frac{1}{2} \right|$$




Fig. 8. Theorem 2: The adversary's decision algorithm in case of colored noise addition

APPENDIX B PROOF OF THEOREM 2

Proof: By applying the same procedure described in Appendix A to prove Theorem 1, we obtain:

$$R \sim \mathcal{N}\left(\mu_R = \sum_{t=0}^{N_s-1} s_a[t]^2 - \sum_{t=0}^{N_s-1} s_a[t]s_b[t],\right. \\ \left.\sigma_R^2 = 2\sigma_{L'}^2 \int_0^1 |K(\varphi)|^2 |S_a(\varphi)|^2 d\varphi\right)$$

In this case (see Figure 8), we can evaluate $\text{PSD}_{R_{s_a, L_a}}(\varphi)$ and $\text{PSD}_{R_{s_a, L_b}}(\varphi)$, i.e., the output PSDs of the colored noise processes $L_a[t]$ and $L_b[t]$ once filtered by the LTI correlation filter, in the following way:

$$\text{PSD}_{R_{s_a, L_a}}(\varphi) = \text{PSD}_{R_{s_a, L_b}}(\varphi) = \text{PSD}_L(\varphi) |S_a(\varphi)|^2 \\ = \sigma_{L'}^2 |K(\varphi)|^2 |S_a(\varphi)|^2$$

Hence, it follows that:

$$\sigma_R^2 = 2 \int_0^1 \text{PSD}_{R_{s_a, L_a}}(\varphi) d\varphi = 2\sigma_{L'}^2 \int_0^1 |K(\varphi)|^2 |S_a(\varphi)|^2 d\varphi$$

The rest of the proof is analogous to Theorem 1. ■

REFERENCES

- [1] G. Hart, "Nonintrusive appliance load monitoring," *Proceedings of the IEEE*, vol. 80, no. 12, pp. 1870–1891, Dec. 1992.
- [2] C. Laughman, K. Lee, and R. e. a. Cox, "Power signature analysis," *Power and Energy Magazine, IEEE*, vol. 1, no. 2, pp. 56–63, Mar. 2003.
- [3] National Institute of Standards and Technology (NIST), "Guidelines for smart grid cyber security," NIST Interagency Report 7628, Aug. 2010. [Online]. Available: <http://www.nist.gov>
- [4] H. Kreuzmann, S. Vollmer, N. Tekampe, and A. Abromeit, "Protection profile for the gateway of a smart metering system (Gateway PP)," Federal Office for Information Security Germany, Aug. 2011.
- [5] M. Jawurek, F. Kerschbaum, and G. Danezis, "Privacy technologies for smart grids - a survey of options," Microsoft, Tech. Rep. MSR-TR-2012-119, November 2012.
- [6] Z. Erkin, J. Troncoso-Pastoriza, R. Lagendijk, and F. Perez-Gonzalez, "Privacy-preserving data aggregation in smart metering systems: an overview," *Signal Processing Magazine, IEEE*, vol. 30, no. 2, pp. 75–86, March 2013.
- [7] G. Danezis, C. Fournet, M. Kohlweiss, and S. Zanella-Béguélin, "Smart meter aggregation via secret-sharing," in *Proceedings of the First ACM Workshop on Smart Energy Grid Security*, ser. SEGS '13. New York, NY, USA: ACM, 2013, pp. 75–80.
- [8] C. Rottondi, G. Verticale, and A. Capone, "Privacy-preserving smart metering with multiple data consumers," *Computer Networks*, vol. 57, no. 7, pp. 1699–1713, 2013.
- [9] S. McLaughlin, P. McDaniel, and W. Aiello, "Protecting consumer privacy from electric load monitoring," in *Proceedings of the 18th ACM Conference on Computer and Communications Security*, ser. CCS '11. New York, NY, USA: ACM, 2011, pp. 87–98.
- [10] S. Ye, F. Wu, R. Pandey, and H. Chen, "Noise injection for search privacy protection," in *Computational Science and Engineering, 2009. CSE '09. International Conference on*, vol. 3, Aug 2009, pp. 1–8.
- [11] L. Sankar, S. Rajagopalan, S. Mohajer, and H. Poor, "Smart meter privacy: A theoretical framework," *Smart Grid, IEEE Transactions on*, vol. 4, no. 2, pp. 837–846, June 2013.
- [12] G. Acs and C. Castelluccia, "I have a DREAM! (differentially private smart metering)," in *The 13th Information Hiding Conference (IH)*, 2011.
- [13] S. Rajagopalan, L. Sankar, S. Mohajer, and H. Poor, "Smart meter privacy: A utility-privacy framework," in *Smart Grid Communications, 2011 IEEE International Conference on*, oct. 2011, pp. 190–195.
- [14] S. Barker, A. Mishra, D. Irwin, E. Cecchet, P. Shenoy, and J. Albrecht, "Smart*: An open data set and tools for enabling research in sustainable homes," in *The 1st KDD Workshop on Data Mining Applications in Sustainability (SustKDD)*, 2011.
- [15] C. Dwork, "Differential privacy," in *Automata, Languages and Programming*, ser. Lecture Notes in Computer Science. Springer Berlin / Heidelberg, 2006, vol. 4052, pp. 1–12.
- [16] —, "Differential privacy: A survey of results," in *Theory and Applications of Models of Computation*. Springer, 2008, pp. 1–19.
- [17] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of Cryptography*, ser. Lecture Notes in Computer Science, S. Halevi and T. Rabin, Eds. Springer Berlin Heidelberg, 2006, vol. 3876, pp. 265–284.
- [18] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, "Our data, ourselves: privacy via distributed noise generation," in *Proceedings of the 24th annual international conference on The Theory and Applications of Cryptographic Techniques*, ser. EUROCRYPT'06. Berlin, Heidelberg: Springer-Verlag, 2006, pp. 486–503.
- [19] G. Danezis, M. Kohlweiss, and A. Rial, "Differentially private billing with rebates," in *Information Hiding*, ser. Lecture Notes in Computer Science, T. Filler, T. Pevn, S. Craver, and A. Ker, Eds. Springer Berlin Heidelberg, 2011, vol. 6958, pp. 148–162.
- [20] V. Rastogi and S. Nath, "Differentially private aggregation of distributed time-series with transformation and encryption," in *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*. New York, NY, USA: ACM, 2010, pp. 735–746.
- [21] E. Shi, T. Chan, and e. a. Rieffel, "Privacy-preserving aggregation of time-series data," in *NDSS Symposium*, Aug. 2011.
- [22] T. Chan, E. Shi, and D. Song, "Privacy-preserving stream aggregation with fault tolerance," in *Financial Cryptography and Data Security*, ser. Lecture Notes in Computer Science, A. Keromytis, Ed. Springer Berlin / Heidelberg, 2012, vol. 7397, pp. 200–214.
- [23] M. Jelasity and K. P. Birman, "Distributional differential privacy for large-scale smart metering," in *Proceedings of the 2Nd ACM Workshop on Information Hiding and Multimedia Security*. New York, NY, USA: ACM, 2014, pp. 141–146.
- [24] J. Zhao, T. Jung, Y. Wang, and X. Li, "Achieving differential privacy of data disclosure in the smart grid," in *INFOCOM, 2014 Proceedings IEEE*, April 2014, pp. 504–512.
- [25] F. Zhang, L. He, W. He, and X. Liu, "Data perturbation with state-dependent noise for participatory sensing," in *INFOCOM, 2012 Proceedings IEEE*, March 2012, pp. 2246–2254.
- [26] C. Rottondi, G. Verticale, and C. Krauss, "Distributed privacy-preserving aggregation of metering data in smart grids," *Selected Areas in Communications, IEEE Journal on*, vol. 31, no. 7, pp. 1342–1354, July 2013.
- [27] J. G. Proakis and M. Salehi, *Digital Communications, 5th Edition*. McGraw-Hill, 2007.
- [28] M. Jawurek, M. Johns, and K. Rieck, "Smart metering de-pseudonymization," in *Proceedings of the 27th Annual Computer Security Applications Conference*. ACM, 2011, pp. 227–236.
- [29] "Smart meters and smart meters systems: A metering industry perspective," EEI-AEIC-UTC White Paper, Mar. 2011.
- [30] G. Heffner, C. Goldman, B. Kirby, and M. Kintner-Meyer, "Loads providing ancillary services: Review of international experience," U.S. Department of Energy, Ernesto Orlando Lawrence Berkeley National Laboratory, Tech. Rep., 2008.



Marco Savi is a PhD student in the Department of Electronic, Information and Bioengineering (DEIB) at the Politecnico di Milano, Italy. He received his B.Sc (2010) and M.Sc (summa cum laude, 2012) in telecommunication engineering from Politecnico di Milano. His research interests mainly focus on the security and privacy issues in Smart Grids, on the performance evaluation and energy-efficient strategies definition for content caching in the access network, and on the Network Function Virtualization paradigm. He is a student member of IEEE.



Cristina Rottondi is postdoctoral researcher in the Department of Electronics, Information, and Bioengineering (DEIB) at the Politecnico di Milano, Italy. She received both Bachelor and Master of Science degrees "cum laude" in telecommunications engineering from Politecnico di Milano. Her research interests include communication security and demand side management in smart grids, internet traffic analysis and classification, and optical network planning.



Giacomo Verticale is assistant professor at Politecnico di Milano, Italy. He obtained his PhD in Telecommunication Engineering in year 2003 from Politecnico di Milano defending a thesis on the performance of packet transmission in UMTS. In the years 1999-2001 he was with the research center CEFRIEL, working on the Voice-over-IP and ADSL technologies. He was involved in several European research projects advancing the Internet technology. His current interests focus on the security issues of the Smart Grid and on Network Function Virtualization. He is a member of IEEE and ACM.