Spectral Differential Privacy: Application to Smart Meter Data

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY 4.0

SUBMISSION DATE / POSTED DATE

03-02-2021 / 06-02-2021

CITATION

Parker, Kendall; Barooah, Prabir; Hale, Matthew (2021): Spectral Differential Privacy: Application to Smart Meter Data. TechRxiv. Preprint. <https://doi.org/10.36227/techrxiv.13706665.v1>

DOI

[10.36227/techrxiv.13706665.v1](https://doi.org/10.36227/techrxiv.13706665.v1)

Spectral Differential Privacy: Application to Smart Meter Data

Kendall Parker, *Student Member, IEEE*, Matthew Hale, *Member, IEEE*, and Prabir Barooah, *Member, IEEE*

Abstract—We present Spectral Differential Privacy (SpDP), a novel form of differential privacy designed to protect the frequency content of time series data that come from wide sense stationary stochastic processes. This notion is motivated by privacy needs in applications with time series data over unbounded time, such as smart meters. First, a notion of differential privacy on the space of (discretized) spectral densities is introduced. A Gaussian-like mechanism for SpDP is then presented that provides differential privacy to the spectral density. Next, a novel streaming implementation is developed to enable real-time use of the proposed mechanism. The privacy guarantee provided by SpDP is independent of the time duration over which data is collected or shared. In contrast, time-domain trajectory-level differential privacy (TrDP) will require noise with large variance to provide privacy over an extended time duration. The technique is numerically evaluated using smart meter data from a single home to compare the utility of SpDP to that of time-domain trajectory-level differential privacy. The noise added by SpDP is substantially smaller than that added by time-domain TrDP, particularly when privacy over long time horizons is sought by TrDP.

Index Terms—differential privacy, spectrum, smart meter, trajectory

I. INTRODUCTION

The Internet of Things (IoT) is a central hub of interconnected, data-driven technologies supporting a variety of critical infrastructure. In the power grid, smart meters connect to the IoT for smarter energy management. They can benefit utility companies in billing, consumption monitoring and load forecasting. For consumers, they can be a tool to plan for conservation or monitor electricity use [1]. However, smart meter data is often collected at high temporal resolutions and can reveal sensitive information about its source. Such data has been well documented to reveal consumer presence, absence, or lifestyle patterns [2–4]. Hence, the nature of this data creates serious privacy concerns.

Differential privacy (DP) is a formal privacy framework that can strike a balance between the need for detailed data and addressing privacy concerns. Using a statistical notion of privacy, DP adds carefully calibrated noise to sensitive data (or functions thereof) to protect it [5]. For instance, authors in [6] develop a mechanism to improve DP for high dimensional data with linear queries. The value of DP to smart metering and related grid applications has been recognized by many

researchers. DP protects the location of smart homes from traffic analysis attacks in [7]. A data aggregation scheme that ensures DP in the presence of general measurement report failures was explored in [8].

While DP originated in the context of databases, it was extended to data in the form of trajectories or signals, termed Trajectory-level Differential Privacy (TrDP) [9]. A challenge in TrDP is adequately protecting sensitive events as data length grows. In many applications, such as analytics with smart meter data, there may be instances when an upper bound on the required time duration for an analytic is unknown, and thus a need may arise to protect many sensitive events over time. Privacy noise grows with the duration and magnitude of sensitive events to privatize, and thus the protection of many events over arbitrarily long time horizons can require arbitrarily large noise. Indeed, the noise scale in TrDP may grow to the point that the privatized trajectory is useless. This pitfall has been addressed for classical DP, where the author weakens privacy guarantees in the distant past [10].

In this work, we propose a new notion of differential privacy targeting this weakness of TrDP, by going to the frequency domain. In our approach, which we term *Spectral Differential Privacy* (SpDP), a signal's power spectral density (PSD) is treated as sensitive information, and a new definition and approach to privacy are defined to protect the spectral representation of sensitive events, rather than the time-domain representation. A privatized PSD in this setting is the result of a privacy mechanism applied directly to a PSD. By computing the PSD offline, this mechanism can be implemented offline. Our motivation is privacy of smart meter data, and a meter must transmit time-domain data in a streaming (as opposed to batch) fashion. We therefore also provide a streaming implementation to compute and share a signal in real time so that the PSD of the transmitted signal is the same as the privatized PSD that the SpDP mechanism generates.

Motivation comes from the observation that frequency content of smart grid signals are highly sensitive because of the capability to exploit usage patterns and routines in customers [11–13]. Through analysis of frequency content, the energy distribution of smart meter data per unit time can be classified into frequency bins, thus identifying prominent times of activity. Frequency-domain representation of DP has only been seen in [14] where authors study local DP with singular spectrum analysis.

The main advantage of the proposed privacy framework is that the noise in the resulting time domain data that is shared, i.e., the output of the streaming implementation, is bounded irrespective of data length. In contrast, noise in the privatized data shared by TrDP grows without bound as the

KP, MH, and PB are with the Department of Mechanical & Aerospace Engineering at University of Florida. Emails: {kendallparker, matthewhale, pbarooah}@ufl.edu.

Kendall Parker was partially supported by the Florida Education Fund (FEF) and GEM Consortium. Matthew Hale was supported by the AFOSR Center of Excellence on Assured Autonomy in Contested Environments and by NSF CAREER Grant #1943275.

time duration increased. Data privatized with SpDP therefore has more utility for downstream analytics when long time intervals are involved.

Two additional contributions of this paper are: (1) a novel non-i.i.d. additive Gaussian mechanism and (2) a data-based calibration method for choosing the adjacency parameters for SpDP and time domain TrDP. A related work that does not provide DP guarantees but privatizes smart meter data with correlated noise is [15]. The adjacency parameter is a design choice, and few guidelines exist in the literature on differential privacy how to choose its numerical value.

The rest of this paper is organized as follows: Section II summarizes TrDP and its weaknesses, defines the required mathematical preliminaries for SpDP and explains the problem solved by SpDP. Following this, Section III defines all elements of SpDP along with streaming implementation. Section IV discusses numerical results, followed by conclusions and future work in Section V.

II. PRELIMINARIES

The symbols $\mathbb{R}, \mathbb{R}_+, \mathbb{Z}, \mathbb{Z}_+$ denote the sets of real, non-negative real, integers and non-negative integers, respectively. A discrete sequence x is a function $x : \mathbb{Z} \rightarrow \mathbb{R}^n$, i.e., $x_k \in \mathbb{R}^n$ for every $k \in \mathbb{Z}_+$. The prototypical signal of interest is power demand of a customer. For that signal, the index k in x_k is a discrete time index, corresponding to the time ticks when data is sent.

The ℓ_2 norm of a sequence x is $\|x\|_{\ell_2} = \left(\sum_{k=0}^{\infty} \|x_k\|_2^2\right)^{1/2}$, and the symbol ℓ_2 denotes the set of all sequences x that have finite ℓ_2 norm. Further, the notation $\tilde{\ell}_2^n$ denotes the set of all sequences $x : \mathbb{N} \rightarrow \mathbb{R}^n$ such that every finite truncation of the sequence has a finite ℓ_2 norm. In other words, $x \in \tilde{\ell}_2^n$ if $\|x_k\|_2 < \infty$ for all k . Often we will consider a finite truncation of a sequence, such as $x_{1:n} := (x_1, \dots, x_n)$. For consistency, we will use $\|x_{1:n}\|_{\ell_2}$ to denote the 2-norm of the vector of truncated values, i.e. $\|x_{1:n}\|_{\ell_2} := (\sum_{i=1}^n \|x_i\|_2^2)^{1/2}$.

A. Summary of TrDP

The purpose of differential privacy (DP) is to mask the differences between adjacent pieces of data by ensuring that they produce approximately indistinguishable outputs when a mechanism is applied to them. That is, given some output sequence, it should be unlikely for its recipient to make meaningful distinctions between input sequences that could have produced it. We refer to the input x as the *sensitive* data and the output — a realization of the DP mechanism $M(x)$ — as the *privatized* data. In this work, a system will directly add noise to its outputs before sharing them. This is the input perturbation approach to DP, and has the advantage of masking sensitive data before it is shared.

The essential notions from trajectory-level differential privacy (TrDP) with the Gaussian mechanism are stated in the following proposition; see [9] for a thorough exposition. We first define adjacency using an adjacency parameter $B > 0$, which is a design parameter. Two trajectories are declared

(a) A consumer's power demand.

(b) Distance between demand data from the same consumer.

Fig. 1: (a) Sensitive power demand d and (b) the distance between $d_{1:K}$ and $d_{K+1:2K}$ as a function of duration K , which illustrates the trouble with TrDP. Data: Pecan Street Project [16].

adjacent if their ℓ_2 -distance is bounded above by B . An (ϵ, δ) -differential privacy mechanism makes adjacent signals approximately indistinguishable.

Proposition 1. Fix $B > 0$.

- 1) Two sequences $x, y \in \tilde{\ell}_2^n$ are said to be adjacent if $\|x - y\|_{\ell_2} \leq B$.
- 2) A mechanism M is (ϵ, δ) -differentially private with respect to this adjacency relationship if

$$\mathbb{P}[M(x) \in A] \leq e^\epsilon \mathbb{P}[M(y) \in A] + \delta \quad (1)$$

for all measurable $A \subseteq \tilde{\ell}_2^n$ and all adjacent x and y .

- 3) The Gaussian mechanism $M(x) = x + w$ is (ϵ, δ) -differentially private with $w(k) \sim \mathcal{N}(0, \sigma^2 I)$, if

$$\sigma \geq \frac{B}{2\epsilon} \left(Q^{-1}(\delta) + \sqrt{Q^{-1}(\delta)^2 + 2\epsilon} \right) \quad (2)$$

and $Q(a) = \frac{1}{\sqrt{2\pi}} \int_a^\infty \exp(-u^2/2) du$ is the Gaussian tail integral.

The interpretation is as follows: the parameter ϵ controls information leakage about a sensitive data, and smaller values of ϵ imply less leakage and hence stronger privacy. The parameter δ can be interpreted as the probability that ϵ -differential privacy fails. In the literature, typical values are $\epsilon \in (0, \log 3)$ and $\delta \in [0, 0.5]$ [5, 9].

To apply the notion of TrDP to time-domain demand data sequences, consider two power demand trajectories (in kW)

of length K from the same consumer¹: $d_{1:K} := \{d_k\}_{k=1}^N$ and $d_{K+1:2K} := \{d_k\}_{k=K+1}^{2K}$. As K increases, the distance $D(K) := \|d_{1:K} - d_{K+1:2K}\|_2$ increases without bound. Figure 1 is evidence in support of this claim for electrical power demand data collected at a 5-minute interval from a single home.

Since the time series used to find the distance $D(K)$ are from the same consumer over two consecutive time intervals, a reasonable notion of adjacency should qualify them as adjacent. Figure 1b shows that the time-domain distance between these trajectories is large and grows without bound over time. Thus, in a TrDP framework, a large adjacency parameter B is required to qualify the time series as adjacent. For given privacy parameters ϵ and δ , a large B necessitates the use of high-variance noise for privacy, cf. Proposition 1. As privacy is demanded for data of even larger duration K , the value of $B(K)$ must also be chosen larger. Since the variance of privacy noise is proportional to B^2 , the accuracy of any analytics with the privatized data degrades as K increases. Consequently, the utility of privatized data decreases commensurately as K increases. *To provide (ϵ, δ) -DP for a fixed ϵ and δ independent of the time duration K , the adjacency parameter — and thus the noise added — must be infinitely large. As a result, the utility of privatized data for analytics becomes zero. If B is held fixed while the length of trajectories increase without bound, then only a small number of events and/or events of short duration (relative to the length of the trajectory) can be protected through time domain TrDP.*

Authors in [9] address adjacency parameter selection for arbitrary finite-horizon (and finite-dimensional) settings. The subtlety there is that the B must be fixed based on the events that need protection; as the duration of interest increases, a new adjacency parameter must be fixed to protect events in the longer trajectory.

This discussion shows the difficulty of using standard TrDP to provide privacy to smart meter data, or to any smart grid applications requiring time-domain data. The difficulty comes from the fact that two data streams generated by the same behavior and same consumer will have small differences that do not vanish as time increases. Thus, the distance between these two time series will grow without bound as the time duration of the series grows. An exception is time series x with asymptotically vanishing x_k 's so that $x \in \ell_2$, but such time series are not relevant to smart grid applications.

B. Mathematical Preliminaries on PSDs

Here we summarize mathematical preliminaries needed to formally define the problem that is the subject of this work. Throughout this work we assume sensitive time-domain data is a wide sense stationary (WSS) stochastic process to make use of well-established spectral theory for stationary time series. A stochastic process x is called WSS if $\mathbb{E}[x_k] = \mu$ and $\mathbb{E}[x_k x_{k+m}] := R_{xx}(k, k+m) = R_{xx}(m)$ for all k , where $\mathbb{E}$ denotes expectation. A pure WSS model may not be

appropriate for many types of smart grid data due to seasonal and other quasi-periodic variations. However, such data can be modeled as a deterministic time-varying mean plus a WSS process.

The power spectral density of a zero mean WSS process x is the Fourier transform of its autocorrelation, or, alternatively, a limit of the average of the square of the Fourier transform of the truncated data:

$$\Phi_{xx}(\omega) := \sum_{m=-\infty}^{\infty} R_{xx}[m] e^{-j\omega m} \quad (3)$$

$$= \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\left| \sum_{k=1}^T x_k e^{-j\omega k} \right|^2 \right], \quad (4)$$

where ω is the (continuous) frequency variable. The equality between (3) and (4) is the Wiener-Khinchin Theorem [17]. The process x is assumed to be zero mean throughout the paper for notational convenience. Otherwise, every definition that involves an expectation must subtract the mean.

When the PSD exists, it is non-negative, 2π -periodic, and even, which makes it uniquely specified by its values over $[0, \pi]$ [17]. We therefore focus all operations on PSDs on the interval $[0, \pi]$ for the remainder of the paper.

We will work with a sampled version of the function $\Phi_{xx}(\omega)$. We choose an integer N and sample the PSD at $2N$ frequency points, where the n -th frequency ω_n is:

$$\omega_n = \frac{2\pi F_s}{2N} n \quad (\text{rad/sample}), \quad n = 0, \dots, 2N-1 \quad (5)$$

where F_s is the sampling frequency of the time-domain data in samples per unit time. Denote the part that goes up to the Nyquist frequency as ϕ^N :

$$\phi^N = [\Phi_{xx}(e^{j\omega_0}), \Phi_{xx}(e^{j\omega_1}), \dots, \Phi_{xx}(e^{j\omega_N})]^T \in \mathbb{R}_+^{(N+1)}. \quad (6)$$

Because of the periodic and even nature of the PSD $\Phi(\omega)$, which is retained by the sampled version, the $2N$ samples of the PSD are uniquely determined by the first $N+1$ samples, that is, by the array ϕ^N . Therefore, we will identify the sampled PSD with $2N$ samples by the array ϕ^N .

The time index on the sample path of x is unbounded, but the length of the array ϕ^N is fixed at $N+1$. The length of the sampled PSD is a design variable selected based on the frequencies one wants to resolve in the PSD. The superscript N will be often omitted to reduce clutter, and the sampled PSD will be referred to as ϕ .

There are several reasons why it is meaningful to consider the sampled PSD as the sensitive data to be privatized. First, as long as N is large enough to resolve the frequencies of interest, the sampled version will contain most of the information from the function $\Phi_{xx}(\omega)$. The second reason considers computation needs. All numerical algorithms to estimate the PSD do not estimate the continuous function $\Phi_{xx}(\omega)$. Rather, they estimate the vector ϕ for some N . The expectation of this unbiased estimate approaches the true spectral density of the process x for large K [18]. In addition, in order to have good estimation accuracy N must be far less than K . Moreover since the streaming implementation of

¹We use x to denote an arbitrary time series or stochastic process. Here, we consider demand signals specifically, which is indicated by a change in notation to d for the signal.

Fig. 2: Estimates of the power spectral density of a consumer's power demand computed from data of varying lengths. Data: Pecan Street Project [16].

SpDP developed in Section III and the analysis of privatized data require the sampled PSD, privacy guarantees provided directly on ϕ are useful.

C. Problem Definition

The privacy goal we consider in this work is masking trajectories from a single consumer to prevent adversaries from exploiting repetitive behaviors to uncover sensitive information. As previously mentioned, time horizons can be of arbitrary length and our DP framework is motivated by this characteristic. Since the frequency content of signals is sensitive information [11–13], we focus on privatizing the power spectral densities (PSD) of the power demand rather than the time-domain data itself. Informally, the PSD $\Phi_{xx}(\omega)$ of a stochastic process x is an average of a square of the Fourier transform of the sequence. In light of the weakness of time-domain DP described previously, the key advantage of privacy in the Fourier domain is that the PSD is defined over a frequency interval that is independent of the length of time involved. The highest frequency is the Nyquist frequency, which is half of the sampling frequency of the data. Thus, the noise needed to privatize intuitively adjacent PSDs is not dependent on (and does not grow with) the time interval.

To illustrate this advantage, Figure 2 shows three different PSD estimates of a consumer's power demand. These estimates are computed from varying lengths of time-domain data, but the frequency range only goes up to the Nyquist frequency, which is $0.5 \times \frac{1}{300} = 1.67 \times 10^{-3}$ Hz ($= 6 \text{ hour}^{-1}$) in this case since the data is sampled every 5 minutes. In addition, one can see from the figure that PSD estimates obtained from various data traces of the same consumer do not differ much. Later it will be shown that this feature allows a uniform B (irrespective of time interval involved) to quantify adjacency for DP when the PSD is considered the sensitive data to be privatized.

With this, the following problem is the subject of this work:

Problem 1. Given a time-domain signal $x = \{x_k\}_{k \in \mathbb{N}}$ and its sampled power spectral density (PSD) estimate ϕ , do the following:

- 1) Design an (ϵ, δ) -differential privacy mechanism $\mathcal{M}$ that privatizes the frequency-domain sensitive data ϕ in an appropriate sense to produce privatized PSD $\tilde{\phi}$.

- 2) Develop a streaming implementation of $\mathcal{M}$ that generates samples $\tilde{x}_k$ in real-time so that the PSD of $\tilde{x}$ is $\tilde{\phi}$. The time-domain data that the streaming implementation produces must still be useful in downstream analytics.

Problem 1.1 is solved with correlated Gaussian mechanism and the SpDP mechanism in Section III-B. Problem 1.2 is solved in Section III-C.

D. A Correlated Gaussian Mechanism for TrDP

The Gaussian mechanism mentioned in Proposition 1 uses i.i.d. noise which is standard in the literature on TrDP [5]. We now present an extension to the non-i.i.d. case, which will be useful in our sequel.

Proposition 2. Fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and let d be a sensitive signal and let privacy parameters $\epsilon > 0$, $\delta \in (0, 0.5)$ be given. Consider the correlated Gaussian mechanism $\mathcal{M} : \tilde{\ell}_2^n \times \Omega \rightarrow \tilde{\ell}_2^n$ defined by $\mathcal{M}(d) = d + w$ with $W \sim \mathcal{N}(0, \Sigma)$. This mechanism is (ϵ, δ) -differentially private if $\lambda_{\min}(\Sigma)$, the minimum eigenvalue of the covariance matrix Σ , satisfies $\lambda_{\min}(\Sigma) \geq \underline{\lambda}$, where

$$\underline{\lambda} := \frac{B^2}{(-A + \sqrt{A^2 + 2\epsilon})^2}, \quad A = \mathcal{Q}^{-1}(\delta) \quad (7)$$

Due to its technical nature, proof of the proposition is in the Appendix. With correlated noise, the probability of having a large difference between ϕ and its private counterpart is lessened compared to a more common i.i.d. mechanism. With this, the output time-domain signal from streaming implementation will maintain better correlation with the time-domain sensitive data when correlated noise is used. This will be further discussed in Section III-C.

Although there are many choices, we provide a specific procedure for generating a correlated noise vector w with a controllable degree of correlation among its entries to be in development of SpDP.

- 1) Consider a periodic and even function $m(\cdot)$ over integers $m : \mathbb{Z} \rightarrow \mathbb{R}_+$ defined as $m(r) = e^{-\beta \bmod(r, n)}$ for $r > 0$, and $m(-r) = m(r)$, $m(0) = 1$, where $\beta > 0$ is a constant. Define a $n \times n$ matrix as $M[i, j] = m(j - i)$, which is symmetric by construction.
- 2) Define $\Sigma = \frac{\underline{\lambda}}{\lambda_{\min}(M)} M$, where $\lambda_{\min}(M)$ is the minimum eigenvalue of M .
- 3) Compute $L := \Sigma^{-1/2}$, the unique symmetric positive definite square root of Σ , and generate $z \sim \mathcal{N}(0, I_n)$ and compute $w = Lz$.

The next result shows that all steps of the procedure can be implemented if M is positive definite, and that the resulting random vector w has the desired distribution.

Proposition 3. If $\beta > 1$, the procedure above generates a vector $w \sim \mathcal{N}(0, \Sigma)$ where Σ is a symmetric positive definite matrix, the correlation coefficient between w_i and w_j is $e^{-\beta|i-j|}$, and $\lambda_{\min}(\Sigma) \geq \underline{\lambda}$ where $\underline{\lambda}$ is given in (7).

Proof. Since every diagonal entry of $M_{n \times n}$ is 1, and symmetry of M tells us all eigenvalues of M are real, the Gershgorin circle theorem says the eigenvalues of M are in

the union of the intervals $[1 - R_\ell, 1 + R_\ell]$, where the radius $R_\ell = \sum_{j=1, j \neq \ell}^n e^{-\beta|i-j|}$ [19]. Consider two cases: n is odd and n is even. For the case n is odd, the sum in R_ℓ is maximum for the middle row, with index $\ell^* := (n+1)/2$. This is because this row has the most slowly decaying exponential terms. The corresponding radius is

$$R_{\ell^*} = \sum_{k=1}^{(n-1)/2} e^{-\beta k} < \int_0^{(n-1)/2} e^{-\beta t} dt$$

where the inequality follows by viewing the sum as the total area of unit-width rectangles under the curve $e^{-\beta t}$ starting from $t = 0$ to $(n-1)/2 - 1$. We now have

$$R_{\ell^*} < \int_0^{(n-1)/2} e^{-\beta t} dt < \int_0^\infty e^{-\beta t} dt = \frac{1}{\beta}$$

Thus, if $\beta \geq 1$, $R_\ell < 1$ for every $\ell = 1, \dots, n$, and thus every eigenvalue is positive. This proves positive definite nature of M . For the case n is even, the maximum radius R_{ℓ^*} occurs at two distinct row indices $(n-1)/2$ and $(n+1)/2$. The rest of the argument is the same as the odd case. Since M is symmetric positive definite, so is Σ . The scaling involved in the definition of Σ ensures that its minimum eigenvalue is $\underline{\lambda}$. Since Σ is symmetric positive definite, $L = \Sigma^{-1/2}$ exists and is symmetric positive definite. Now, the covariance of w is $\mathbb{E}[ww^T] = \mathbb{E}[L^2] = \Sigma$, and since linear combination of i.i.d. Gaussian r.v.s is Gaussian, we get $w \sim \mathcal{N}(0, \Sigma)$. The statement about the correlation coefficient follows from the observation that the correlation coefficient between w_i and w_j is precisely $M[i, j]$. $\square$

Larger values of β result in a random vector w whose entries w_i and w_j have less correlation as the lag $i - j$ increases, and smaller values of β have the opposite effect. Note that although the sufficient condition in Proposition 3 is quite conservative and numerical investigation shows positive definiteness of Σ is obtained at all positive values of β . We leave that as a conjecture.

III. SPECTRAL DIFFERENTIAL PRIVACY (SpDP)

This section develops Spectral Differential Privacy including an adjacency relation between the sampled power spectra, a formal mechanism, and streaming implementation to generate a private time series from the differentially private PSD. We use the input perturbation approach to privacy, which can provides privacy to a database with a single entry.

A. Definitions

Indistinguishability of signals in SpDP is made parallel to TrDP using an adjacency relation on the PSD estimates rather than the time series. Subscripts i, j indicate PSD estimates of two different trajectories from a single consumer.

Definition 1. Fix $B > 0$. The adjacency relation Adj_B is defined $\forall \phi_i, \phi_j \in \mathbb{R}_+^N$ for $p \geq 1$ as

$$Adj_B(\phi_i, \phi_j) = \begin{cases} 1 & \|\phi_i - \phi_j\|_p \leq B \\ 0 & \text{otherwise.} \end{cases}$$

Fig. 3: Distances between pairs of time series and sampled PSD estimates. Data: Pecan Street Project [16].

Following the privacy goal, each user applies this definition to mask sensitive events in their data individually. Therefore, the user specifies the boundedness defining Adj_B . Namely, a user selects an adjacency parameter $B > 0$, and their true PSD is made approximately indistinguishable from all other PSDs within distance B by the privacy mechanism.

The benefits of an adjacency relation on the PSDs is best illustrated in comparison with TrDP. Figure 3 shows the distance between PSDs of demand as a function of duration of data used to estimate the PSD. The same figure also shows the distance between demand data (in the time-domain), which increases monotonically as the data duration increases. The 2-norm was used in both cases to compute this distance. For PSD estimates, the distance does not grow as one considers longer durations of data. In fact, the distance between PSDs appears to settle down to a constant. A goal in this work is to classify two PSDs as adjacent if the same consumer under identical or nearly identical behaviors produced them. With a fixed length PSD, the adjacency parameter needed to do so will not increase with time series length. This results in constant distances between PSDs and bounded noise added to the PSD for differential privacy.

A mechanism $\mathcal{M}(\cdot)$ in SpDP is a randomized mapping with domain and co-domain $\mathbb{R}^N$. We define

$$\tilde{\phi} = \mathcal{M}(\phi) \quad (8)$$

so that $\mathcal{M}$ provides (ϵ, δ) -differential privacy to frequency-domain sensitive data ϕ for given ϵ and δ . We formally define SpDP below with probability space $(\Omega, \mathfrak{F}, \mathbb{P})$, and we use the Borel σ -algebra over $\mathbb{R}^N$, denoted $\mathcal{B}_N$.

Definition 2. A mechanism $\mathcal{M} : \mathbb{R}^N \times \Omega \rightarrow \mathbb{R}^N$ preserves (ϵ, δ) -differential privacy if, for all adjacent PSD signals $\phi_i, \phi_j \in \mathbb{R}^N$,

$$\mathbb{P}[\mathcal{M}(\phi_i) \in \mathcal{S}] \leq e^\epsilon \mathbb{P}[\mathcal{M}(\phi_j) \in \mathcal{S}] + \delta \quad \forall \mathcal{S} \in \mathcal{B}_N.$$

SpDP is made parallel to TrDP in that The privacy mechanism is applied to the PSD in in a parallel manner so that (1) each realization of the private PSD is itself a valid PSD and (2) the aforementioned definition is satisfied. This is consistent with the standard interpretation of differential privacy applied to time-domain data; with SpDP, an adversary will be unlikely to make meaningful distinctions between signals' frequency content. The claim is that SpDP provides

spectral differential privacy to all PSDs (and likewise provides its attendant immunity to post-processing and robustness to side information) with improved utility over TrDP because it requires less noise.

B. A SpDP Mechanism

Algorithm 1 describes a mechanism for SpDP. It uses concepts from positive dynamical systems. A positive dynamical system is one that, if the initial condition and input are non-negative, has non-negative states and outputs [20].

Algorithm 1: Mechanism $\mathcal{M}_{SpDP}$ to provide (ϵ, δ) -Spectral Differential Privacy SpDP.

Input: Sensitive PSD of demand ϕ , noise covariance matrix Σ , positive filter $P(z)$

Output: A differentially private PSD $\tilde{\phi}$

- /* Apply Gaussian mechanism */
 - 1 Set $\phi' \leftarrow \phi + \eta$, with $\eta \sim \mathcal{N}(0, \Sigma)$. /*
 - /* Make values non-negative */
 - 2 Set $\phi' \leftarrow (\phi')^+$ for all $\omega \in [0, \pi]$ /*
 - /* Apply $P(z)$ non-causally */
 - 3 Set $\tilde{\phi} \leftarrow P(z)[\phi']$
-

We compactly represent all the steps involved in the mechanism as

$$\mathcal{M}_{SpDP}(\phi) := P(z)[(\phi + \eta)^+], \quad (9)$$

where $P(z)[y]$ indicates the filter $P(z)$ is used on the signal y and the notation $(y)^+$ denotes the negative values thresholded to 0. This algorithm provides DP to the signal ϕ , and thus solves Problem 1a, as summarized in the following theorem.

Theorem 1. For a given adjacency parameter $B > 0$ and privacy parameters $\epsilon > 0$ and $\delta \in (0, 0.5)$, Algorithm 1 provides (ϵ, δ) -differential privacy to the PSD ϕ and produces a valid PSD $\tilde{\phi}$, if $\lambda_{\min}(\Sigma) \geq \underline{\lambda}$ where $\underline{\lambda}$ is defined in (7).

Proof. Consider the intermediate array $\phi' := \phi + \eta$, where η is a vector of N samples from a Gaussian distribution with mean 0 and covariance Σ . Because of the eigenvalue condition in the hypothesis, (ϵ, δ) -differential privacy of ϕ' follows immediately from Proposition 2. However, the sum $\phi + \eta$ will be negative at some indices (frequencies) with non-zero probability due to the zero mean nature of η , and hence the sum may not be a valid sampled PSD. Applying post-processing by thresholding negative values to zero makes the signal non-negative at all frequencies, which ensures the output is a valid PSD. Filtering with a positive dynamical system makes the result smoother while maintains non-negativity, which ensures that the output of the mechanism is a valid PSD. That $\tilde{\phi}$ is (ϵ, δ) -differentially private follows from immunity to post-processing of differential privacy: ϕ' is (ϵ, δ) -differentially private, and subsequent operations are merely post-processing, which means their outputs have the same level of privacy [5]. $\square$

Figure 4 demonstrates a numerical example of $\mathcal{M}_{SpDP}$ applied with Algorithm 1. The noise was generated using

Fig. 4: Numerical example of an SpDP mechanism: a sampled PSD estimate of a consumer's electric power demand is the sensitive data, while the privatized PSD is the output of the SpDP mechanism of Algorithm 1. The sensitive PSD is obtained from a consumer demand data in Pecan Street Project [16].

the procedure described in Section II-D. The figure shows the frequency-domain sensitive data ϕ , which is the sampled PSD of a consumer's demand and the privatized PSD data $\tilde{\phi}$ obtained by applying the mechanism $\mathcal{M}_{SpDP}$. These PSDs were estimated using one month of time-domain data. Values of the parameters used in this numerical example are provided in Table I. The positive dynamical system used in this example is a discrete time low-pass filter with cutoff frequency ω_N :

$$P(z) = K_p \frac{\omega_N z^{-1}}{1 + (\omega_N - 1)z^{-1}}. \quad (10)$$

C. Streaming Implementation

Estimates of a consumer's PSD from power demand data can be determined from past usage, and hence privatizing the PSD occurs offline. However, meters must transmit time-domain data, not PSDs, so a streaming implementation is necessary for a time-domain application of SpDP. Thus, a streaming implementation generates values in time whose PSD is the same as the private PSD that the SpDP mechanism generates. For streaming implementation, we assume the frequency-domain sensitive data ϕ , the privatized PSD $\tilde{\phi}$, and the time-domain sensitive data x are available to the smart meter that will perform the streaming implementation. While ϕ , $\tilde{\phi}$ are known *a priori*, x is available only in real-time. Also, development of a streaming implementation requires smart meters that are tamper resistant and trusted with the ability to perform filtering. Figure 5 illustrates the streaming implementation described in the following algorithm.

The next proposition shows the steps above are a valid streaming implementation of the mechanism described in Section III-B.

Proposition 4. If $\gamma[n] > 0$ for $n = 0, \dots, N$, then Algorithm 2 is feasible, and the sampled PSD of the released time-domain data $\tilde{x}$ is $\tilde{\phi}$, and the cross-correlation function between x and $\tilde{x}$ is $R_{x\tilde{x}}[m] = R_{xx}[m] * f^*[-m]$ where f is the impulse response of $F(z)$.

Proof. The on-line steps are feasible. If the hypothesis about entry-wise positivity of the array γ is satisfied, then γ can be

Algorithm 2: Streaming implementation of SpDP

Input: Time-domain sensitive signal $x = \{x_k\}_{k \in \mathbb{N}}$, its PSD ϕ , its privatized form $\tilde{\phi}$, and the filter $F(z)$

Output: A time-domain signal $\tilde{x} := \{\tilde{x}_k\}_{k \in \mathbb{N}}$ related to x whose PSD is $\tilde{\phi}$

/* Offline: */

1 Compute $\gamma \in \mathbb{R}_+^{N+1}$, where

$$\gamma[n] := \tilde{\phi}[n] - |F(e^{j\omega_n})|^2 \phi[n], n = 0, \dots, N. \quad (11)$$

2 Determine $H(z)$ such that $|H(e^{j\omega_n})|^2 = \gamma[n]$, $n = 0, \dots, N$ where ω_n is defined in (5).

/* Online: */

3 **for** $k = 0, 1, \dots$ **do**

4 Generate white Gaussian noise w_k with 0 mean and unit variance.

 /* Apply $H(z)$ to w_k */

5 $c_k \leftarrow H(z)[w_k]$

 /* Apply $F(z)$ to x_k */

6 $d_k^F \leftarrow F(z)[x_k]$

7 $\tilde{x}_k \leftarrow x_k^F + c_k$

8 Release $\tilde{x}_k$

9 **end for**

Fig. 5: Illustration of streaming implementation

viewed as a sampled version of a PSD $\gamma(\omega)$ in the frequency range $[0, \pi)$. Thus, a stable spectral factor $H(z)$ exists so that $|H(e^{j\omega})|^2 = \gamma(\omega)$ [17]. There are many algorithms for computing such a spectral factor given a sampled version of the PSD, see [21] and references therein. One of them can be used to determine H which will by design satisfy the requirement in the algorithm, that $|H(e^{j\omega_n})|^2 = \gamma[n]$, $n = 0, \dots, N$. This proves feasibility. To show that the

released time-domain data $\tilde{x}$ has the desired PSD,

$$\begin{aligned} PSD(\tilde{x}) &= PSD(x^F + c) \\ &= PSD(x^F) + PSD(c) \end{aligned}$$

where the second equality follows from the fact that c is independent of x^F . Denoting by $\Phi_{xx}(\omega)$ the PSD of a stochastic process x , along with standard properties relating PSDs of inputs and outputs of filters, and the fact that the PSD of a i.i.d. unit variance Gaussian process is 1 at all frequencies [17], we get from the above that

$$\begin{aligned} \Phi_{\tilde{x}\tilde{x}}(\omega) &= |F(e^{j\omega})|^2 \Phi_{xx}(\omega) + |H(e^{j\omega})|^2 \\ \Rightarrow \Phi_{\tilde{x}\tilde{x}}(\omega_n) &= |F(e^{j\omega_n})|^2 \Phi_{xx}(\omega_n) + |H(e^{j\omega_n})|^2, \quad n = 0, \dots, N \\ &= |F(e^{j\omega_n})|^2 \Phi_{xx}(\omega_n) + \gamma[n], \quad n = 0, \dots, N \\ &= \tilde{\phi}[n], \quad n = 0, \dots, N \end{aligned}$$

where the last equality follows by substituting (11) in the previous equation. This shows that the PSD of the process $\tilde{x}$ at the frequencies ω_n have the desired value, and so the algorithm is a streaming implementation of the SpDP mechanism. The statement about the cross correlation follows from standard results on the cross-correlations between inputs and outputs of a linear filter (see [17, Chapter 9]) upon recognizing that c is independent of x . $\square$

To satisfy the positivity requirement of $\gamma[n]$, one has to choose the reduction filter F appropriately. A poorly designed filter can make $\gamma < 0$ at some n , in which case its spectral factorization into H is not theoretically possible. In that case, Algorithm 2 is not implementable. In practice, F is a low-pass filter, with adaptable cutoff frequency to ensure positivity of γ . Filter design can be performed after the mechanism has been applied and the private PSD $\tilde{\phi}$ is available, and so the positivity condition can always be maintained. Since the filter F is used in streaming implementation, it does not affect the privacy guarantees on the frequency-domain data ϕ . The positivity requirement is also maintained through use of the correlated Gaussian Mechanism. If there are large differences between the filtered version of ϕ and $\tilde{\phi}$, the spectral factorization of γ would result in time-domain noise with large standard deviation. This is due to the resulting large magnitude of γ . Instead, with correlated noise privatizing ϕ , there is less probability of large values of γ and thus the resulting noise from streaming implementation is of smaller standard deviation.

Apart from feasibility of streaming implementation, design of F along with the differentially private noise η determine the degree of correlation between released time-domain data $\tilde{x}$ and time-domain sensitive data x . Figure 5b illustrates this: if the filter F has low gain at some frequency, the gap γ between the PSDs of sensitive demand (i.e., ϕ) and filtered demand (i.e., $|F(e^{j\omega_n})|^2 \phi$) will be large at that frequency. Recall this gap is filled by the colored noise c , since the PSD of c is $|H(e^{j\omega_n})|^2 = \gamma[n]$. Thus, the released data will have a large noise variance c compared to time-domain sensitive data x at that frequency. Depending on the level of noise in the particular realization $\tilde{\phi}$ of the mechanism $\mathcal{M}$, the reduction filter may have to be designed with extremely

TABLE I: Parameters Used in SpDP

Parameter	Description	Value	Units
N	Number of NFFT Points	288	none
ϵ	Privacy Budget	$\log(2)$	none
δ	Privacy Failure Parameter	0.01	none
β	Degree of Noise Correlation	0.5	none
ω_N	Positive Filter Cutoff Frequency	0.39	Hz
K_p	Positive Filter Gain	0.8	None
ω_F	Cutoff Frequency of $F(z)$	2×10^{-4}	Hz
K_F	Gain of $F(z)$	0.8	None
B_{SpDP}	SpDP Adjacency Parameter	0.12	$\text{kW}^2 \cdot \text{Hour}$
$B_{TrDP}^{(1)}$	4-Hour Adjacency Parameter	0.2	kW
$B_{TrDP}^{(2)}$	1-Week Adjacency Parameter	39	kW

low gain at certain frequencies. In that case the time-domain privatized data $\tilde{x}$ produced by streaming implementation will have low correlation with the time-domain sensitive data x . Correspondingly, downstream analytics with the released time-domain data $\tilde{x}$ will then be less accurate than those done with the sensitive data x . The loss of accuracy will increase as the gain of F is reduced. In contrast, as the gain of F approaches 1, the loss of accuracy approaches 0 but streaming implementation may be infeasible.

IV. NUMERICAL EVALUATION

We evaluate the proposed paradigm on consumer demand from a single home in Pecan Street [16]. The privacy goal considered in our numerical evaluation is to protect demand data from a single home with a 5-minute sampling period. By first applying the SpDP mechanism $\mathcal{M}_{SpDP}$, and then performing the streaming implementation, we conduct a full application of SpDP. The result is that the privatized data $\tilde{d}$ is streamed by the smart meter in real time based on time-domain sensitive demand data d immediately as the sensitive data is measured.

The *Lisa Technology Package Data Analysis* (LTPDA) was used to generate correlated noise c noise from private PSD $\hat{\phi}$ [22]. Numerical values of the parameters used in the study are in Table I.

A. Choice of B

The adjacency parameter B depends on the choice of norm used to define distances between trajectories. Beyond this, B is a design choice and few guidelines exist on how to choose it. Since time-domain TrDP and SpDP use vastly different norms to define distances, the choice of B must differ in these two distinct privacy paradigms.

Recall that the privacy goal is to protect demand data from a single home. For SpDP we choose B as

$$B_{SpDP} = \max_{i,j \in \mathcal{N}} \|\phi_i - \phi_j\|_2 \quad (12)$$

where ϕ_i is the estimate of the sampled PSD computed from the i -th time-domain dataset $\mathbf{d}_i := d_{k_i}, d_{k_i+1}, \dots, d_{k_i+N_i}$, and $\mathcal{N}$ is a set of such time-domain data, with each dataset potentially of distinct duration N_i . Each numerical estimate of the PSD from a particular time-domain dataset can be

thought of a distinct PSD itself, which is a frequency-domain characterization of the potential behavior of a consumer.

The true (unknown) PSD of the consumer does not vary with time duration. Thus, to determine B_{SpDP} we want the time-domain datasets that produce similar PSD estimates, with differences among estimates attributable to estimation errors. It turned out that PSDs from 7 to 12 week durations produced the most similar PSD estimates and the result from (12) was $0.12 \text{ kW}^2 \times \text{Hour}$ and this value was chosen as B_{SpDP} . This adjacency parameter can be used in SpDP irrespective of time duration and protects events up to a frequency of 6 hour^{-1} , which is the Nyquist frequency corresponding to the sampling period of 5 minutes.

For time-domain TrDP with the same privacy goal, an appropriate choice of B will be

$$B_{TrDP}(K) = \max_{\ell, m} \|d_{1:K}^{(\ell)} - d_{1:K}^{(m)}\|_2 \quad (13)$$

which depends on the time interval K . Superscripts ℓ, m indicate two non-overlapping power demand trajectories of length K from a single consumer. As discussed in Section II-A, calibrating privacy in this way for TrDP requires an infinitely large adjacency parameter for data over an unbounded time interval, and so the choice of B necessarily depends on the time duration K over which privacy is to be provided.

For illustration, we consider two scenarios with two distinct K 's, denoted $K^{(1)}$ and $K^{(2)}$. The first is $K^{(1)} = 576$, corresponding to four hours of data. This yields $B_{TrDP}^{(1)} = 0.2 \text{ kW}$. The second is for $K^{(2)} = 2016$, corresponding to one week of data, which yields $B_{TrDP}^{(2)} = 39 \text{ kW}$.

The correlation coefficient between the time-domain sensitive data and the released time-domain data is taken as a measure of utility. A higher correlation coefficient corresponds to more useful data for downstream analytics. When the coefficient is 1 that means the time-domain sensitive data itself is released and there is no loss in accuracy of analytics, though it would also mean no privacy is afforded to the consumer.

B. Results

Results of the full implementation of SpDP and implementation of TrDP for the two scenarios, applied to the demand data from a house in Pecan Street are in Figure 6. Figure 6a shows time-domain sensitive data d and released data $\tilde{d}$ after applying SpDP mechanism (Algorithm 1) and streaming implementation (Algorithm 2). Figure 6b and 6c show the same sensitive data along with data privatized using TrDP and the Gaussian mechanism in Proposition 1. The additional parameters used in SpDP are in Table I.

TrDP shows superior performance over SpDP in the first scenario (TrDP designed for 4 hours) in terms of utility, but SpDP outperforms TrDP in utility for the second scenario (TrDP designed for a week). In fact, for any time duration higher than four hours SpDP outperformed TrDP. It should be emphasized that SpDP provides differential privacy for all time durations.

Table II shows the standard deviation of added time-domain noise and correlation coefficients between the time-domain sensitive data and released data for each result in Figure 6.

(a) SpDP Results

(b) TrDP Results using $B_{TrDP}^{(1)} = 0.2$ kW (c) TrDP Results using $B_{TrDP}^{(2)} = 39$ kW

Fig. 6: A 24-hour snapshot of time-domain sensitive data and privatized results from (a) SpDP, which are the same no matter the time duration considered, (b) TrDP, providing privacy for 4 hours, and (c) TrDP, providing privacy for one week. Data: Pecan Street Project [16].

TABLE II: Comparison of TrDP and SpDP Results

	Standard Deviation of Added Noise	Correlation Coefficient
SpDP	2.1 kW	0.34
TrDP - 4 Hours	0.92 kW	0.88
TrDP - 1 Week	0.18 MW	-0.015

From these results, TrDP only outperforms SpDP at extremely short time durations, illustrated by the higher correlation coefficient. Increasing the length of the signal for which privacy is sought results in poor performance of TrDP and little to no utility of the privatized signal.

V. CONCLUSIONS AND FUTURE WORK

A new notion of differential privacy, called Spectral Differential Privacy (SpDP), was presented. SpDP treats the power

spectral density (PSD) of the underlying process that generates the data as the sensitive information that needs protection. The frequency content of time-series data can reveal patterns as potential sources of exploitation by adversaries. Differential privacy guarantees on the frequency content of data are designed in this work to circumvent this issue. A key advantage of SpDP is the noise needed to provide a certain level of privacy to the PSD is independent of time duration. Moreover, the adjacency parameter – which determines the noise – in SpDP is selected from data and, to the best of our knowledge, this work is the first to justify this design choice for privatizing signals.

In contrast, the noise needed for privacy in time-domain TrDP increases without bound as the time duration increases. Numerical evaluations with a consumer's electrical demand data show that SpDP reduces the noise needed significantly compared to TrDP for equal levels of differential privacy as time duration increases.

Next steps in evaluating SpDP involve improvements to the reduction filter to hide specific data features and improve downstream analytics. Further, an assessment of analytics, such as billing and real-time monitoring, will be performed to quantify the error in the released data estimate. Additionally, the impact of the WSS assumption on mechanism design is of interest since this assumption may not hold for many smart grid processes due to seasonal or other cyclic phenomenon.

REFERENCES

- [1] S. Thorve, L. Kotut, and M. Semaan, "Privacy preserving smart meter data," *Proceedings of The 7th International Workshop on Urban Computing (UrbComp'18)*, August 2018.
- [2] Office of the General Counsel, "Data access and privacy issues related to smart grid technologies," United States Department of Energy, Tech. Rep., 2010.
- [3] The European Data Protection Supervisor, "Opinion of the european data protection supervisor on the commission recommendation on preparations for the roll-out of smart metering systems," European Union, Tech. Rep., June 2012.
- [4] V. Y. Pillitteri and T. L. Brewer, "Guidelines for smart grid cybersecurity," National Institute of Standards and Technology (NIST), Tech. Rep., September 2014.
- [5] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Foundations and Trends in Theoretical Computer Science*, vol. 9, no. 3&4, pp. 211–407, Aug. 2014. [Online]. Available: <http://dx.doi.org/10.1561/04000000042>
- [6] Z. Zheng, T. Wang, J. Wen, S. Mumtaz, A. K. Bashir, and S. H. Chauhdary, "Differentially private high-dimensional data publication in internet of things," *IEEE Internet of Things Journal*, vol. 7, no. 4, pp. 2640–2650, 2020.
- [7] J. Liu, C. Zhang, and Y. Fang, "Epic: A differential privacy framework to defend smart homes against internet traffic analysis," *IEEE Internet of Things Journal*, vol. 5, no. 2, pp. 1206–1217, 2018.

- [8] H. Bao and R. Lu, "A new differentially private data aggregation with fault tolerance for smart grid communications," *IEEE Internet of Things Journal*, vol. 2, no. 3, pp. 248–258, 2015.
- [9] J. Le Ny and G. J. Pappas, "Differentially private filtering," *IEEE Transactions on Automatic Control*, vol. 59, no. 2, pp. 341–354, 2013.
- [10] F. Farokhi, "Temporally discounted differential privacy for evolving datasets on an infinite horizon," *2020 ACM/IEEE 11th International Conference on Cyber-Physical Systems (ICCPS)*, pp. 1–8, 2020.
- [11] F. Siddiqui, S. Zeadally, C. Alcaraz, and S. Galvao, "Smart grid privacy: Issues and solutions," in *2012 21st International Conference on Computer Communications and Networks (ICCCN)*. IEEE, 2012, pp. 1–5.
- [12] S. R. Rajagopalan, L. Sankar, S. Mohajer, and H. V. Poor, "Smart meter privacy: A utility-privacy framework," in *2011 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, 2011, pp. 190–195.
- [13] L. Sankar, S. R. Rajagopalan, S. Mohajer, and H. V. Poor, "Smart meter privacy: A theoretical framework," *IEEE Transactions on Smart Grid*, vol. 4, no. 2, pp. 837–846, 2012.
- [14] L. Ou, Z. Qin, S. Liao, T. Li, and D. Zhang, "Singular spectrum analysis for local differential privacy of classifications in the smart grid," *IEEE Internet of Things Journal*, vol. 7, no. 6, pp. 5246–5255, 2020.
- [15] M. Savi, C. Rottondi, and G. Verticale, "Evaluation of the precision-privacy tradeoff of data perturbation for smart metering," *IEEE Transactions on Smart Grid*, vol. 6, pp. 2409–2416, 2015.
- [16] J. Rhodes, C. Upshaw, C. Harris, C. Meehan, D. Walling, P. Navrátil, A. Beck, K. Nagasawa, R. Fares, W. Cole, H. Kumar, R. Duncan, C. Holcomb, T. Edgar, A. Kwasinski, and M. Webber, "Experimental and data collection methods for a large-scale smart grid deployment: Methods and first results," *Energy*, vol. 65, pp. 462 – 471, 2014.
- [17] A. Papoulis and U. Pillai, *Probability, random variables and stochastic processes*, 4th ed. McGraw-Hill, 11 2001.
- [18] R. G. Brown and P. Y. Hwang, *Introduction to random signals and applied Kalman filtering*. Wiley Online Library, 1992, vol. 2.
- [19] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 2nd ed. Cambridge University Press, 2012.
- [20] L. Farina and S. Rinaldi, *Positive Linear Systems: Theory and Applications*, P. Hilton, H. Hochstadt, M. B. A. III, D. A. Cox, and J. T. Peter Lax, Eds. John Wiley & Sons, Inc., 2000.
- [21] P. Stoica and R. L. Moses, *Spectral Analysis of Signals*. Pearson Education, 2005.
- [22] E. M. Karsten Danzmann, "Lisa Technology Package Data Analysis (LTPDA)," Max Planck Society, 2014. [Online]. Available: <https://www.lisamission.org/ltlda/>

APPENDIX A PROOF OF PROPOSITION 2

Proof. Define $u = \mathcal{M}(x)$ and random variables $W \sim \mathcal{N}(0, \Sigma)$ with observations w and $U \sim \mathcal{N}(x, \Sigma)$ with observations u . For $S \in \ell_p$, denote $v := x - x'$ where x, x' are adjacent signals according to Proposition 1.

$$\mathbb{P}(U \in S) = \frac{1}{(2\pi)^{K/2} |\Sigma|^{1/2}} \int_S e^{-\frac{1}{2}(u-x)^T \Sigma^{-1}(u-x)} dw \quad (14)$$

$$= \frac{1}{(2\pi)^{K/2} |\Sigma|^{1/2}} \int_S \left[e^{-\frac{1}{2}(u-x')^T \Sigma^{-1}(u-x')} e^{(u-x')^T \Sigma^{-1}v - \frac{1}{2}v^T \Sigma^{-1}v} \right] du \quad (15)$$

$$\leq e^\epsilon \mathbb{P}(M(x') \in S) + \frac{1}{(2\pi)^{K/2} |\Sigma|^{1/2}} \int_{S \cap S_2} \left[e^{-\frac{1}{2}(u-x)^T \Sigma^{-1}(u-x)} \right. \quad (16)$$

$$\left. \mathbb{I}_{\{u|(u-x)^T \Sigma^{-1}v + \frac{1}{2}v^T \Sigma^{-1}v - \epsilon > 0\}} \right] du$$

Here, we completed the square and partitioned the sample space such that $\Omega = S_1 \cup S_2$ where $S_1 = \{u|(u-x')^T \Sigma^{-1}v - \frac{1}{2}v^T \Sigma^{-1}v - \epsilon \leq 0\}$ and $S_2 = \{u|(u-x')^T \Sigma^{-1}v - \frac{1}{2}v^T \Sigma^{-1}v - \epsilon > 0\}$. For S_2 note that, $(u-x')^T \Sigma^{-1}v = (u-x)^T \Sigma^{-1}v + v^T \Sigma^{-1}v$. This substitution is needed because the mean of U is defined by x , not x' .

Since Σ is a valid covariance matrix, there exists a positive definite, symmetric square root of its inverse Σ^{-1} that can be defined as $L := \Sigma^{-1/2}$. Let $y^T = (u-x)^T L$ and represent observations of random variable $Y \sim \mathcal{N}(0, I_K)$. From here we seek to bound the right-hand side of (16) by δ .

$$\mathbb{P}(U \in S_2) = \mathbb{P}\left(Y^T L v > \epsilon - \frac{1}{2}v^T L^2 v\right) \quad (17)$$

Defining an intermediate random variable $Z = \frac{Y^T L v}{\|Lv\|_2}$ with distribution $\mathcal{N}(0, 1)$, the integral of (16) can be rewritten as $\mathbb{P}\left(Z \geq \frac{\epsilon}{\|Lv\|_2} - \frac{\|Lv\|_2}{2}\right) \leq \delta$ which is equivalent to $\mathcal{Q}\left(\frac{\epsilon}{\|Lv\|_2} - \frac{\|Lv\|_2}{2}\right) \leq \delta$ for $\mathcal{Q}(t) = \frac{1}{\sqrt{2\pi}} \int_t^\infty e^{-\frac{u^2}{2}} du$. For ease of notation let $A = \mathcal{Q}^{-1}(\delta)$.

$$\epsilon - \frac{\|Lv\|_2^2}{2} \geq A\|Lv\|_2 \quad (18)$$

$$\|Lv\|_2 \leq -A + \sqrt{A^2 + 2\epsilon} \quad (19)$$

Note that $\|Lv\|_2 \leq \|L\|_2 \|v\|_2 \leq \|L\|_2 B$ since $\|v\|_2 = \|x - x'\|_2 \leq B$. Let the sorted eigenvalues of Σ be $\lambda_{max} \geq \dots \geq \lambda_{min}$ so the eigenvalues of L are $\frac{1}{\sqrt{\lambda_{min}}} \geq \dots \geq \frac{1}{\sqrt{\lambda_{max}}}$. Then,

$$\|L\|_2 B \leq -A + \sqrt{A^2 + 2\epsilon} \quad (20)$$

$$\frac{1}{\sqrt{\lambda_{min}}} \leq \frac{-A + \sqrt{A^2 + 2\epsilon}}{B} \quad (21)$$

$$\lambda_{min} \geq \frac{B^2}{(-A + \sqrt{A^2 + 2\epsilon})^2} \quad (22)$$

If $\lambda = \lambda_{min}$, then (ϵ, δ) -differential privacy follows. $\square$