Spectral Differential Privacy: Application to Smart Meter Data

Kendall Parker[1], Matthew Hale[1], Anand Radhakrishnan[2], Prabir Barooah[1]

***Abstract*— We present Spectral Differential Privacy (SpDP), a novel form of differential privacy designed to protect the frequency content of time series data. First, a notion of differential privacy on the space of power spectral densities is introduced. We then present a Gaussian-like mechanism for SpDP. Next, a novel streaming implementation is developed to enable real-time use of our 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 will require noise with infinite variance to provide privacy over an infinite time duration. We numerically evaluate our technique using smart meter data from a single home and compare the utility of SpDP to that of time-domain trajectory-level differential privacy. The noise added by spectral differential privacy is substantially smaller than that added by time-domain trajectory differential privacy, particularly when privacy over long time horizons is sought.**

I. Introduction

As data-driven technologies proliferate, there are increased privacy concerns associated with harvesting user data. For example, the smart power grid can use granular power usage data to ensure demand is met and prevent waste. However, such data is known to be quite revealing about users, potentially exposing their daily habits and activities [1], [2]. Similar privacy concerns exist for other cyber-physical systems [3], as well as other forms of critical infrastructure, such as autonomous transportation [4]. These applications therefore require protection of sensitive data that still accommodate system performance.

Differential privacy (DP) is a formal privacy framework that can balance the need for data and individuals' privacy concerns. As a statistical notion of privacy, DP protects data by carefully randomizing it or a function of it [5]. An appealing feature of DP is that it is immune to post-processing, in the sense that privacy is not weakened by arbitrary post-hoc operations on privatized data. In the smart grid, this means that arbitrary downstream analytics are permitted on privatized data, whether done by a utility company or a third-party partner.

While differential privacy originated in the context of databases, it has recently been extended to data that is in the form of trajectories or signals, termed trajectory-level differential privacy (TrDP) [6]. TrDP is implemented by adding noise to signals or functions thereof. It can require large noise variance, notably for data across long time horizons. In many applications, such as data analytics with smart meter data, an upper bound on data length for an analytic may not be known. Further, for many types of trajectory data the required noise to provide privacy over an arbitrary time duration can be arbitrarily large, making the privatized data useless for downstream analytics. This weakness of TrDP motivates the need for a privacy mechanism that is suitable for providing privacy to signals over arbitrary time durations. We provide examples in this work for smart meter data.

Existing literature has identified the frequency content of smart grid signals as highly sensitive data [7], [8], [9], and this is what we will privatize with a new notion of differential privacy. In this approach, which we term *Spectral Differential Privacy* (SpDP), we treat a signal's power spectral density (PSD), instead of the signal itself, as the sensitive information to protect. In this setting, a privacy mechanism is applied directly to a PSD to create a privatized PSD. If one can compute a PSD offline, then this mechanism can also be implemented offline. Since our motivation is privacy of smart meter data, smart meters must transmit time domain data in a streaming (as opposed to batch) fashion. We therefore 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.

Privacy for PSDs masks differences between PSDs within a specified distance of each other, and this distance corresponds to the magnitude of sensitive events that must be masked. One can choose this magnitude irrespective of the time duration of interest. Thus, the noise added to the time domain signal to provide SpDP is independent of the time duration of the data transmitted. In contrast, the noise added in Trajectory-level Differential Privacy increases with the duration of data. We illustrate the proposed method numerically using consumer data from the Pecan Street Project database [10], which illustrates the advantage of SpDP over TrDP for long duration data privacy.

The value of differential privacy to smart metering and related applications has been recognized by many researchers; a representative sample of existing works includes [11], [12], [13], [14]. Relative to those existing works, our contribution is developing the first differential privacy implementation for PSDs, together with numerical results benchmarking its performance on actual power usage data.

The structure of this paper is as follows: Section II summarizes formal definitions for trajectory-level differential privacy (TrDP) and reviews some of the mathematical preliminaries. Spectral Differential Privacy definition and

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

[2] AR is with Optym, Inc.

associated mechanism design is covered in Section III. Section IV then provides a streaming implementation design, and Section V provides a numerical example with smart meter data. Section VI provides conclusions and directions for future work.

II. Background, Motivation, and Problem Statement

We first briefly review some background material, and then formally state the two problems that are the subject of the remainder of the paper. The symbols $\mathbb{R}$ and $\mathbb{N}$ denote the sets of real and natural numbers. For a sequence $x : \mathbb{N} \to \mathbb{R}^n$, we use the notation $\|x\|_{\ell_2} = \left(\sum_{k=0}^{\infty} \|x(k)\|_2^2\right)^{1/2}$. We use the notation $\tilde{\ell}_2^n$ denote the set of all sequences $x : \mathbb{N} \to \mathbb{R}^n$ with $\|x(k)\|_2 < \infty$ for all k, i.e., sequences x whose entries are all finite.

A. Recap of TrDP

We state the essential notions from trajectory-level differential privacy (TrDP) in the following proposition; see [6] for a thorough exposition.

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
$$\mathbb{P}\big[M(x) \in A\big] \leq e^{\epsilon}\mathbb{P}\big[M(y) \in A\big] + \delta$$
for all measurable $A \subseteq \tilde{\ell}_2^n$
3) The Gaussian mechanism $M(x) = x + w$ is (ϵ, δ) differentially private with $w(k) \sim \mathcal{N}(0, \sigma^2 I)$, if
$$\sigma \geq \frac{B}{2\epsilon}\Big(Q^{-1}(\delta) + \sqrt{Q^{-1}(\delta)^2 + 2\epsilon}\Big)$$
and $Q(a) = \frac{1}{\sqrt{2\pi}}\int_a^{\infty} \exp(-u^2/2)du$ is the Gaussian tail integral.

The purpose of differential privacy 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.

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. For all data types, typical values are $\epsilon \in (0, \log 3)$ and $\delta \in [0, 0.05]$.

The adjacency parameter B is a design parameter: a larger B effectively declares two sequences with large ℓ_2 distance to be adjacent. An (ϵ, δ) privacy mechanism provides the privacy protection to all the signals within a distance B. The price one pays is the higher noise in the privatized data and thus, the accuracy of any analytics with the privatized data degrades. In particular, the variance of privacy noise is proportional to B^2.

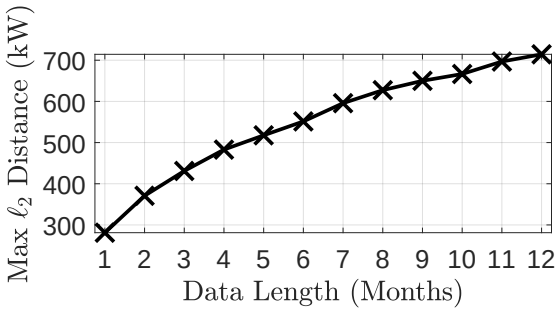

Fig. 1: The distance between $d_{1:N}$ and $d_{N+1:2N}$, as a function of data duration n illustrates the trouble with TrDP. Data used is a particular consumer's power demand from [10].

Consider two power demand (in kW) trajectories of length N from the *same customer*: $d_{1:N} := \{d_k\}_{k=1}^{N}$ and $d_{N+1:2N} := \{d_k\}_{k=N+1}^{2N}$[1] . The distance $\|d_{1:N} - d_{N+1:2N}\|_{\ell_2}$ keeps increasing without bound as N increases; see Figure 2 that uses electrical demand data collected at 1-minute sampling interval from a single home [10]. Since these two time series are from the same residence over two consecutive time intervals, any reasonable notion of adjacency should qualify them as adjacent. In a TrDP framework of differential privacy, a large B will be required to ensure that they do indeed count as adjacent. Larger B leads to larger noise to be added to the sensitive data. The utility of the privatized data is decreased commensurately. *If one wants to provide privacy that is independent of the time duration over which data is collected, B - and thus the noise added - must be infinitely large and the utility of the privatized data for analytics becomes zero.*

This discussion shows the difficulty of using standard TrDP to provide privacy to smart meter data, or any time series with a large number of samples. An exception is a time series x with decaying x_k's so that $x \in \ell_2$, but such time series are not relevant to smart grid applications.

B. Spectral Differential Privacy Problem Statement

We are interested in privatizing power demand trajectories from smart meters over arbitrarily long time horizons, which motivates our development of a new differential privacy framework. As noted in the introduction, the frequency content of such signals is sensitive, and it is therefore the frequency content that we privatize. We next state the problem that is the subject of the remainder of the paper; the ideas and terms contained in it - including the definition

[1] We 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.

of a power spectral density - are formalized in the rest of this section.

Problem 1: Given a time domain signal $d = \{d_k\}_{k\in\mathbb{N}}$ and its power spectral density (PSD) Φ_{dd}, do the following:

a. Design an (ϵ, δ) differential privacy mechanism $\mathcal{M}$ that privatizes Φ (in an appropriate sense).
b. Develop a streaming implementation of $\mathcal{M}$ that generates samples $\tilde{d}_k$ in real-time so that the PSD of $\tilde{d}$ is $\tilde{\Phi}$.

The privatized data that the mechanism and its streaming implementation produce must still be useful in downstream analytics. We use differential privacy because it can provide privacy while still allowing for reasonably accurate data to be released. We solve Problem 1.a. in Section III and Problem 1.b. in Section IV.

C. Mathematical Preliminaries

1) Power Spectral Density (PSD): A stochastic process x is called wide sense stationary (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. The power spectral density of a zero mean WSS process x is the Fourier transform of its autocorrelation,

$$\Phi(\omega) = \mathscr{F}[R_{xx}] = \sum_{m=-\infty}^{\infty} R_{xx}[m]e^{-j\omega m},$$

where $\mathscr{F}$ denotes the Fourier transform and ω is the (continuous) frequency variable [15]. The process x is assumed to be zero mean throughout for the sake for the sake of notational convenience; otherwise the mean must be subtracted in every definition that involves an expectation.

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

2) Reproducing Kernel Hilbert Space: We will need to limit PSDs to belong to a certain reproducing kernel Hilbert space (RKHS). A reproducing kernel Hilbert space $\mathcal{H}$ with kernel $K(\cdot, \cdot)$ (defined over some domain $\mathcal{T}$) is a Hilbert space that is generated by the closure of functions which can be represented as finite linear combinations of the kernel [16]. The RKHS used in the rest of the paper is the Sobolev space

$$H^1[0, \pi] = \left\{ f \in C^1[0, \pi] : \|f\|^2_{\mathcal{H}} < \infty \right\},$$

where the norm is

$$\|f\|^2_{\mathcal{H}} := \frac{f(0)^2 + f(\pi)^2}{2C} + \frac{1}{2\beta C}\int_0^{\pi} (f'(t)^2 + \beta^2 f(t)^2) d\lambda(t),$$

with $C, \beta > 0$. The space $H^1[0, \pi]$ is a RKHS with kernel $K(x, y) = C\exp(-\beta|y - x|)$ [16, Ch. 7].

In the sequel, $\mathcal{H}$ denotes the space $H^1[0, \pi]$ defined above.

III. Spectral Differential Privacy

This section develops Spectral Differential Privacy (SpDP), including appropriate notions of adjacency, and a specific mechanism for providing SpDP; this last development solves Problem 1.a.

A. Defining SpDP

Let x be a real-valued wide-sense stationary (WSS) stochastic process. In Spectral Differential Privacy we treat the PSD of x, $\Phi_{xx} : [0, \pi] \to \mathbb{R}^+$, as the sensitive data to be protected. We consider PSDs that are in the RKHS $\mathcal{H}$.

Remark 1: Differential privacy is often applied to collections of data from N users. Here, we simply set $N = 1$ and apply differential privacy to each individual user, an approach sometimes called "input perturbation" in the literature. See [6] for other uses of this approach in control systems. The interpretation of input perturbation privacy in our context is that each user's PSD is made approximately indistinguishable from nearby PSDs, which are produced by demand signals with similar frequency content. □

The approximate indistinguishability criterion is a hallmark of differential privacy and is enforced by a mechanism, identical to TrDP.

Definition 1: Fix a choice of adjacency parameter $B > 0$. Then adjacency relation Adj_B is defined for all $\Phi, \Xi \in \mathcal{H}$ as

$$\text{Adj}_B(\Phi, \Xi) = \begin{cases} 1 & ||\Phi - \Xi||_{\mathcal{H}} \leq B \\ 0 & \text{otherwise.} \end{cases}$$

In accordance with Remark 1, we will apply this definition to each user individually, thereby masking differences among all individuals' data. We emphasize that the boundedness defining Adj_B is not an assumption but a specification. 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. This definition does not restrict users' data in any way.

We next state differential privacy for PSDs, which is based on developments for functional data in [17]. In accordance with established differential privacy principles, this definition gives immunity to post-processing, in that post-hoc computations do not weaken privacy, and robustness to side information, in that learning additional information beyond a privatized PSD does not weaken differential privacy by much [5].

We use the notation Σ_∞ to denote the Borel σ-algebra over $\mathcal{H}$ under the topology induced by the ∞-norm, defined as $\|\Phi\|_\infty = \sup_{\omega\in[0,\pi]} |\Phi(\omega)|$. When the function $\Phi \in \mathcal{H}$ is also non-negative and integrable ($\int_0^\pi \Phi(\omega)d\omega$ exists), it is a valid PSD since the inverse DTFT $R[k] = \frac{1}{2\pi}\int_0^{2\pi} \Phi^{ext}(\omega)e^{j\omega k}d\omega$ — where Φ^{ext} is the extension of Φ to $[0, 2\pi]$ by mirroring so that it repeated tiling over the entire real line produces a periodic function — yields a valid autocorrelation sequence. We denote by $\mathcal{P}$ the set of real-valued non-negative functions in $[0, \pi]$.

Definition 2 (Spectral Differential Privacy (SpDP)): Consider functions in $\mathcal{H}$, which was defined in Sec. II-C.2, and fix $B > 0$, $\epsilon \in [0, \infty)$, and $\delta \in [0, 1)$. Then a

mechanism $\mathcal{M}$ provides (ϵ, δ)-*Spectral Differential Privacy* (SpDP) if the following two conditions are satisfied:

1) For all Φ, Ξ such that $\text{Adj}_B(\Phi, \Xi) = 1$, and all $A \in \Sigma_\infty$,

$$\mathbb{P}\big[\mathcal{M}(\Phi) \in A\big] \leq e^\epsilon \mathbb{P}\big[\mathcal{M}(\Xi) \in A\big] + \delta.$$

2) $\mathcal{M}(\Phi) \in \mathcal{H} \cap \mathcal{P}$.

If $\mathcal{M}$ only satisfies the first condition, we say $\mathcal{M}$ provides (ϵ, δ)-*functional differential privacy* (FnDP).

B. Advantage of SpDP Over TrDP

With the definition in hand, we now point out the advantage of SpDP over TrDP. Section II-A already presented evidence that the adjacency parameter B for TrDP needs to increase with time duration over which privacy protection is provided. This increases the amount of noise required for differential privacy.

SpDP does not suffer from this weakness because the distance between PSDs does not depend on time. Imagine a consumer Alice and let $\Phi_1^{(A)}$ and $\Phi_2^{(A)}$ be two possible PSDs of Alice's demand. Meaning, these two PSDs represent two possible, statistically distinct behaviors of Alice. To provide SpDP to Alice within the input perturbation framework referred to in Remark 1, the adjacency parameter B needs to be chosen so that $\|\Phi_1^{(A)} - \Phi_2^{(A)}\|_\mathcal{H} \leq B$ for all allowable $\Phi_1^{(A)}, \Phi_2^{(A)}$. *Even if the value of B needed is large, the value is independent of any time duration involved.*

In practice, the PSD — or the possible PSDs — of a consumer's demand is not known. It is possible to estimate the PSD of a stochastic process from sample paths [15]. Estimates with different data lengths can be thought of as possible PSDs that this consumer can produce. Figure 2 shows the distance $\|\Phi_{1:N} - \Phi_{N+1:2N}\|_\mathcal{H}$, where the two PSDs $\Phi_{1:N}$ and $\Phi_{N+1:2N}$ are distinct estimates of the consumer's PSD computed with distinct data sets $d_{1:N}$ and $d_{N+1:2N}$, respectively[2]. We see from the figure that the distance between the PSDs do not increase with increase in data length N. This is expected since the data is generated by the same consumer's behavior, so the statistics of the underlying stochastic process does not change with time. In other words, a fixed B can be used irrespective of the time duration involved.

C. A SpDP Mechanism

Similar to TrDP, additive Gaussian noise can be used to design a SpDP mechanism. The difference is that a continuous Gaussian process is used in SpDP since the sensitive data in this case is functional data.

Theorem 1: Let $\mathcal{H}$ be the RKHS $H^1[0, \pi]$ defined in Section II-C.2, with kernel K. Define the mechanism $\mathcal{M}$

$$\mathcal{M}(\Phi) = \Phi + B\frac{c(\delta)}{\epsilon}G, \tag{1}$$

[2]The norms are computed by using the definition of $\|\cdot\|_\mathcal{H}$ provided in Section II-C.2, after fitting a parametric model to the data-driven non-parametric PSD estimate to facilitate computation of the derivatives involved.

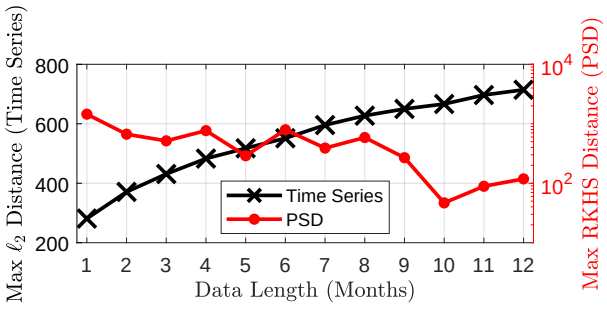

Fig. 2: A comparison of the distance between time-domain sensitive data and corresponding frequency-domain sensitive data (PSD) obtained with increasing data length. Data from Pecan Street project [10].

where G is a sample path of a zero mean Gaussian process whose auto-covariance is the kernel K, B is the adjacency parameter in Definition 1, and $c(\delta) \geq \big(2\log(\frac{2}{\delta})\big)^{\frac{1}{2}}$. Then $\mathcal{M}$ provides (ϵ, δ)-functional differential privacy in the sense of Definition 2.

Due to its technical nature, proof of the theorem is in the appendix.

The Gaussian-like mechanism in Theorem 1 does not provide Spectral Differential Privacy because it may not produce a valid PSD; the Gaussian process added can make the output function negative at some frequencies. Algorithm 1 presented next describes a mechanism to provide SpDP in the sense of Definition 2. We encapsulate all the steps of the algorithm into the notation $\mathcal{M}_{SpDP}$ so that $\tilde{\Phi} = \mathcal{M}_{SpDP}\big(\Phi\big)$ is the output of the algorithm.

The algorithm uses a positive filter. A positive filter is a dynamical system whose output is positive if the input is a non-negative signal and the initial state is non-negative [18]. A simple example is $P(s) = \frac{1}{s+1}$, where s is the Laplace variable.

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

Input: The sensitive PSD of a user's power usage signal, Φ, adjacency parameter B, privacy parameters ϵ and δ, sample path G of zero mean Gaussian process with auto-covariance K, positive filter $P(s)$

Output: $\tilde{\Phi}$: A differentially private form of Φ

`/* Apply Gaussian mechanism */`

1 Set $\Phi_{inter} \leftarrow \Phi + B\frac{c(\delta)}{\epsilon}G$ as in Equation (1)

`/* Make values non-negative */`

2 Set $\Phi_{inter}(\omega) \leftarrow \Phi_{inter}(\omega)^+$ for all $\omega \in [0, \pi]$

`/* Apply` $P(s)$ `non-causally */`

3 Set $\tilde{\Phi} \leftarrow P(s)[\Phi_{inter}]$

We next show through Theorem 2 that Algorithm 1 solves

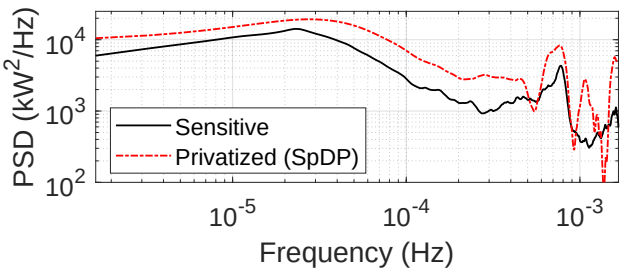

Fig. 3: Comparison of power spectral densities for true, sensitive demand and private demand. The sensitive PSD was smoothed with a low-pass filter to remove numerical artifacts. Private demand contains added noise from (1). Data from Pecan Street project [10].

Problem 1.a.

Theorem 2: Algorithm 1 provides (ϵ, δ)-spectral differential privacy.

Proof: The intermediate function $\Phi_{inter} = \Phi + B\frac{c(\delta)}{\epsilon}G$ satisfies (ϵ, δ)-functional differential privacy by Theorem 1. However, Φ_{inter} is negative at some frequencies with non-zero probability because of Gaussian process path G. Thresholding negative values to 0 (in step 2) makes the function non-negative at all frequencies, but the resulting function can have points of non-differentiability. Filtering with a positive filter makes the resulting function sufficiently smooth and retains non-negativity. The filter is applied non-causally to avoid phase change. Since all operations are done on $[0, \pi]$, symmetry is maintained by mirror imaging, which ensures that the output of the mechanism, $\tilde{\Phi}$, is a valid PSD. That $\tilde{\Phi}$ is (ϵ, δ)-differentially private follows from immunity to post-processing of differential privacy: Φ_{inter} is (ϵ, δ)-functional differentially private, and all subsequent operations are merely post-processing, which means their outputs have the same level of privacy. ■

Figure 3 shows a numerical example of the sensitive data Φ_{dd} (the PSD of a customer's demand d) and the privatized data $\tilde{\Phi}$ obtained by applying the mechanism $\mathcal{M}_{SpDP}$. The privacy parameters used are $\delta = 0.001$, $\epsilon = \log 2$, and $B = 14.5\frac{kW^2}{Hz}$. Choice of B will be discussed in Section V.

A note on numerical computation is in order. Algorithm 1 is for PSDs, which are functions of the continuous frequency parameter ω. All numerical calculations must be done on discrete data, so the PSD is first sampled, and instead of the filter $P(s)$, its sampled-data counterpart $P(z)$ is applied on the sampled PSD.

A related practical issue is obtaining the PSD Φ_{dd} of a consumer's demand $\{d_k\}_{k\in\mathbb{N}}$. One can *estimate* a non-parametric form of the sampled version of the PSD, from the samples d_k by using a standard estimation algorithm, such as averaging of periodograms [15]. To avoid technicalities regarding estimation accuracy, the difference between the sampled version of the PSD (by using the DFT) versus the true PSD (which is the DTFT), we assume that the customer's smart meter has access to the PSD $\Phi_{dd}(\omega)$, whether by a non-parametric estimation method, or by fitting a parametric model to the data, or by a combination thereof.

Strictly speaking, the filter $P(s)$ in Algorithm 1 is not needed. However, the filter makes the resulting PSD sufficiently smooth so that satisfies the differentiability conditions needed for the output $\tilde{\Phi}$ belongs to $\mathcal{H}$. Apart from theoretical convenience, this smoothness helps with the spectral factorization involved in streaming implementation, which is discussed in Section IV.

IV. A Streaming Implementation of SpDP

Thus far, we have identified a method to privatize the PSD of a power demand signal which occurs offline. Since smart meters cannot transmit PSDs, a time-domain implementation of SpDP is needed that we call streaming implementation due to its real-time requirement. We assume for development of this privacy mechanism that smart meters are tamper resistant, trusted, and have the ability to perform filtering.

Recall the streaming implementation problem specified in Prob. 1.b.: for a demand signal $d := \{d_k\}_{k\in\mathbb{N}}$, the streaming implementation must produce in real time a sequence of samples $\tilde{d}_k$ so that the PSD of the process $\tilde{d} := \{\tilde{d}_k\}_{k\in\mathbb{N}}$ is the privatized PSD $\tilde{\Phi} := \mathcal{M}_{SpDP}(\Phi_{dd})$. The streaming implementation will occur at the smart meter.

For the algorithm performing streaming implementation, Φ_{dd} and $\tilde{\Phi}$ are thus known *a priori*. However, d_k is available only in real-time. Our streaming implementation is shown in Algorithm 2.

Algorithm 2: Streaming implementation of SpDP

Input: The sensitive signal $\{d_k\}$, its PSD Φ_{dd}, its privatized form $\tilde{\Phi}$, and the filter $F(z)$

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

`/* Offline: */`

1 Compute

$$\eta(\omega) := \tilde{\Phi}_{dd}(\omega) - |F(e^{j\omega})|^2\Phi_{dd}(\omega)$$

2 Compute $H(z)$ such that $|H(e^{j\omega})|^2 = \eta(\omega)$

`/* Online: */`

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

4 Generate white 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 d_k */`

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

7 Release $\tilde{d}_k = d_k^F + c_k$

8 **end for**

To explain the algorithm and discuss the role of the filter F that must be provided to the algorithm, as well as its design, we first recall the following basic facts.

Proposition 2 ([15]): 1) If w is a white, WSS, and zero-mean stochastic process with variance σ^2, then $\Phi_{ww}(\omega) = \sigma^2$ for all ω.

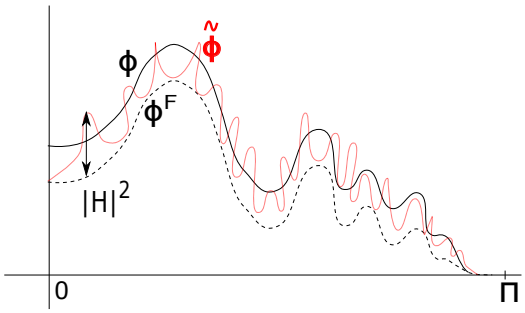

(a) Offline

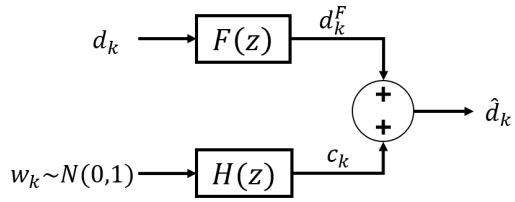

(b) Online

Fig. 4: Visualizing streaming implementation.

2) If x and y are mutually uncorrelated processes, and $z = x + y$, then $\Phi_{zz}(\omega) = \Phi_{xx}(\omega) + \Phi_{yy}(\omega)$ for all ω.
3) If $G(z)$ is a stable linear filter with WSS input x, its output z is also WSS and $\Phi_{zz}(\omega) = |G(e^{j\omega})|^2 \Phi_{xx}(\omega)$.
4) (Paley-Wiener condition) If a function $\Phi(\omega)$, $\omega \in \mathbb{R}$, is real-valued, non-negative, even and 2π-periodic, then there exists an $H(z)$ and $H(z^{-1})$ with both analytic in $|z| > 1$ satisfying $|H(e^{j\omega})|^2 = \Phi(\omega)$ if and only if η and $\log(\eta)$ are integrable functions over $[0, 2\pi)$.

The offline and online operations involved are illustrated in Figure 4. In it, we use d^F to denote the signal resulting from applying a filter F to the process d, and Φ^F to denote the PSD of d^F. As seen from Figure 4a, use of the filter F on the sensitive data d reduces the PSD of the filtered process d^F compared to the sensitive PSD Φ. The difference between the "target" PSD $\tilde{\Phi}$ and that of the PSD of the filtered data is $\eta(\omega)$. A process with this PSD is generated by driving the filter H with white noise, so that their sum has the desired PSD $\tilde{\Phi}$.

The next theorem establishes the conditions under which Algorithm 2 is a valid streaming implementation of the privacy mechanism $\mathcal{M}_{SpDP}$ in Algorithm 1.

Theorem 3: If $\eta(\omega)$ defined in Algorithm 2 satisfies the Paley-Wiener criterion described in Proposition 2, then Algorithm 2 produces a private demand signal $\tilde{d}_k$ in real-time with PSD $\tilde{\Phi}_{dd}$.

Proof: By hypothesis, η is a valid PSD and it can be spectrally factorized to obtain $H(z)$ so the offline steps are feasible. To show that the samples $\tilde{d}_k$ have the desired PSD,

$$\begin{aligned} \text{PSD}(\{\tilde{d}_k\}) &= \text{PSD}(\{d^F\}) + \text{PSD}(\{c_k\}) \\ \Rightarrow \Phi_{\tilde{d}\tilde{d}}(\omega) &= |F(e^{j\omega})|^2 \Phi_{dd}(\omega) + |H(e^{j\omega})|^2 \\ &= \tilde{\Phi}(\omega). \end{aligned}$$

The first equality follows from colored noise c_k and filtered demand d_k^F being uncorrelated, since c_k is obtained by filtering a white noise sequence that is independent on the demand. The second equality follows from Proposition 2. The third equality results from plugging in the definition of η — since it is equal to $|H|^2$ — in the right-hand side of the second equation. ■

Remark 2: The filter F is a design choice and plays a key role. From the hypothesis of Theorem 3, we see that the feasibility of Algorithm depends on a proper design of F. A poorly designed F can make $\eta(\omega) < 0$ at some ω, in which its spectral factorization into H is not possible. In that case, Algorithm 2 is not implementable.

The use of the positive filter $P(s)$ during mechanism design is helpful in ensuring smoothness of $\tilde{\Phi}$. We suspect the Paley-Wiener conditions are satisfied if F is designed so that η is positive.

Apart from feasibility of streaming implementation, design of F determines the degree of correlation between the released time domain data $\tilde{d}$ and the sensitive time domain data d. This can be seen from Figure 4a: if the filter F has low gain at some frequency, the gap η between the PSDs of sensitive demand (i.e., Φ_{dd}) and filtered demand (i.e., Φ^F) will be large at that frequency. Recall that this gap is filled by the noise c_k, since the PSD of c_k is $|H|^2 = \eta$. Thus, the released data will have a high noise (c_k) compared to the signal (sensitive data d_k) at that frequency. The released data $\tilde{d}$ will thus have a smaller correlation with the sensitive data d. Analytics performed with the released data $\tilde{d}$ will be less accurate than those done with the sensitive data d, and 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 then the streaming implementation may be infeasible because of negative values of $\eta(\omega)$. □

V. Numerical Evaluation of SpDP and comparison with TrDP

The SpDP mechanism was already described in Section III; so here we only present the results of the streaming implementation of that mechanism. The numerical results in this section are for privacy parameters $\delta = 0.001$, $\epsilon = \log 2$ (for both TrDP and SpDP). For SpDP, the parameters C and β were chosen as 1 and $0.2\frac{1}{Hz}$, respectively.

In the differential privacy literature, the adjacency parameter B is a design choice and relatively few guidelines exist on how to choose it. It also depends on the choice of norm used to define distance. Since time domain Trajectory-level Differential Privacyand Spectral Differential Privacy use vastly different norms, the choice of B must differ in these two distinct privacy paradigms.

Many distinct estimates of the PSD of a particular residence's demand are performed using time domain data of varying length N, from one-month ($N = 30 \times 24 \times 60$), going up to a year ($N = 365 \times 24 \times 60$). The maximum distance (the $\mathcal{H}$-norm) between these distinct PSD estimates turns out to be 1450 kW^2/Hz; see Figure 2. We choose 1/100 of this distance as the adjacency parameter for SpDP, i.e., $B_{SpDP} = 14.5$ kW^2/Hz.

The streaming implementation in Algorithm 2 was performed using one-month long consumer demand data from a single home in Pecan Street [10]. Figure 5a shows the sensitive time-domain data and the data privatized with the streaming implementation of SpDP.

To compare with TrDP, we consider two distinct data durations: A month and a year. For the first, the distances between pairs of time series of demand data, each series corresponding to one-month (i.e., $N = 30 \times 24 \times 60$) are computed. The maximum distance between these month-long datasets turns out to be 281 kW; see Figure 2. We choose 1/100 of this distance as the adjacency parameter for TrDP, i.e., $B^{(1)}_{TrDP} = 2.81$ kW. Repeating the process for a year ($N = 365 \times 24 \times 60$), the maximum distance is seen to be 714 kW; see Figure 2. The B value is again chosen as 1/100 of this maximum distance, leading to $B^{(2)}_{TrDP} = 7.14$ kW.

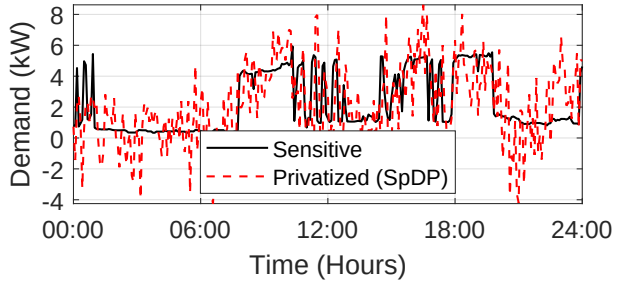

(a) SpDP Results

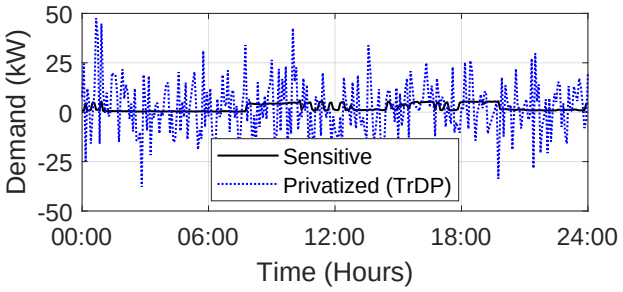

(b) TrDP Results using $B^{(1)}_{TrDP} = 2.81$ kW

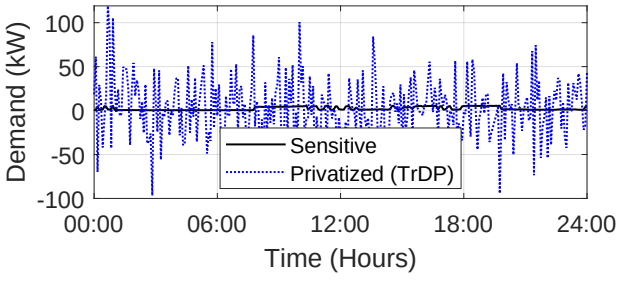

(c) TrDP Results using $B^{(2)}_{TrDP} = 7.14$ kW

Fig. 5: A 24-hour snapshot comparison of sensitive data and privatized results from (a) SpDP, which are the same no matter the time duration considered, (b) TrDP, providing privacy for month-long data, and (c) TrDP, providing privacy for year-long data. Data used is individual consumer demand from [10].

TABLE I: Results Comparison for TrDP and SpDP

	B	Standard Deviation	Correlation Coefficient	SNR
SpDP	$14.5 \frac{kW^2}{Hz}$	$81 \frac{kW^2}{Hz}$	0.39	2.03
TrDP – 1 Month	2.81 kW	13 kW	0.15	0.17
TrDP – 12 Months	7.14 kW	33 kW	0.07	0.03

Time domain data privatized with TrDP is also shown in Figure 5b for $B = B^{(1)}_{TrDP}$ and 5c for $B = B^{(2)}_{TrDP}$. The first can be thought of as providing privacy for month-long data and the second for year-long data, both within the traditional TrDP framework. Note that the sensitive demand data is the same in all three figures. We can see from the figures, and from result in Table I that signal to noise ratio (SNR) for SpDP implementation is much better than the SNR for TrDP. This difference is most prominent when TrDP is used to provide year-long data privacy. *Moreover, the SNR for SpDP is not only better than that of TrDPbut is also independent of the time duration over which data is collected or shared.* Thus, SpDP is likely to enable higher accuracy downstream analytics with privatized data than TrDP, especially when long duration data is involved. This is supported by the improved correlation coefficients of SpDP between sensitive and privatized data in Table I. Precise quantification of the level of accuracy needed and its benefits will depend on the type of analytics, and is outside the scope of this paper.

VI. Conclusions & Future Work

A new notion of differential privacy — Spectral Differential Privacy (SpDP) — was presented that is better suited for providing privacy to long duration data than the existing Trajectory-level Differential Privacy (TrDP). In SpDP the power spectral density is considered as the sensitive data to privatize. The main advantage is that the noise needed to provide a certain level of privacy is independent of time duration. 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 TrDPfor equal levels of differential privacy.

This paper only takes the first step in developing mechanisms and their streaming implementations for SpDP. The Gaussian-like mechanism provided here is only one possibility. We believe the noise added during streaming implementation can be further reduced by alternate mechanism designs. Future development also needs to investigate the impact of the WSS assumption on mechanism design. This assumption may not hold for power demand signals due to seasonal or other cyclic phenomenon. Additionally, future work will include refinement of filters used in streaming implementation to hide specific features of power demand, assessment of the impact on downstream data analytics such as billing, etc.

References

[1] Office of the General Counsel, "Data access and privacy issues related to smart grid technologies," United States Department of Energy, Tech. Rep., 2010.

[2] 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.

[3] J. Giraldo, E. Sarkar, A. A. Cardenas, M. Maniatakos, and M. Kantarcioglu, "Security and privacy in cyber-physical systems: A survey of surveys," *IEEE Design & Test*, vol. 34, no. 4, pp. 7–17, 2017.

[4] C. Bloom, J. Tan, J. Ramjohn, and L. Bauer, "Self-driving cars and data collection: Privacy perceptions of networked autonomous vehicles," in *Proceedings of the Thirteenth USENIX Conference on Usable Privacy and Security*, ser. SOUPS '17. USA: USENIX Association, 2017, p. 357–375.

[5] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Found. Trends Theor. Comput. Sci.*, vol. 9, no. 3–4, pp. 211–407, Aug. 2014.

[6] J. Le Ny and G. J. Pappas, "Differentially private filtering," *IEEE Transactions on Automatic Control*, vol. 59, no. 2, pp. 341–354, 2013.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

[11] G. Ács and C. Castelluccia, "I have a dream! (differentially private smart metering)," in *Information Hiding*, T. Filler, T. Pevný, S. Craver, and A. Ker, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 118–132.

[12] M. Backes and S. Meiser, "Differentially private smart metering with battery recharging," in *Data Privacy Management and Autonomous Spontaneous Security*, J. Garcia-Alfaro, G. Lioudakis, N. Cuppens-Boulahia, S. Foley, and W. M. Fitzgerald, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 194–212.

[13] G. Barthe, G. Danezis, B. Grégoire, C. Kunz, and S. Zanella-Beguelin, "Verified computational differential privacy with applications to smart metering," in *2013 IEEE 26th Computer Security Foundations Symposium*. IEEE, 2013, pp. 287–301.

[14] 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.

[15] A. Papoulis and U. Pillai, *Probability, random variables and stochastic processes*, 4th ed. McGraw-Hill, 11 2001.

[16] A. Berlinet and C. Thomas-Agnan, *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Kluwer Academic Publishers, 2004.

[17] R. Hall, A. Rinaldo, and L. Wasserman, "Differential privacy for functions and functional data," *J. Mach. Learn. Res.*, vol. 14, no. 1, p. 703–727, Feb. 2013.

[18] 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.

VII. Appendix

A. Proof for Theorem 1

A Gaussian mechanism for general functions of a continuous variable was first presented in [17, Corollary 9]. Our spectral differential privacy framework differs because it considers individual functions rather than functions derived from a database, e.g., due to interpolating among points. Our framework also differs because it considers a parameterized adjacency relationship on the space of PSDs themselves, rather than adjacency for a set of databases from which functions are derived. Accordingly, our proof of Theorem 1 must re-establish the Gaussian mechanism under the conditions needed for spectral differential privacy.

We do so by showing that, for a fixed $B > 0$ and for all PSDs Φ and Ξ such that $\mathrm{Adj}_B(\Phi, \Xi) = 1$, all finitely sampled releases of a privatized PSD are (ϵ, δ)-differentially private. First, consider a set of n points $(\omega_1, \ldots, \omega_n)$, where $\omega_i \in [0, \pi]$ for all i. Then, for K the kernel of $H^1[0, \pi]$, recall that G has autocovariance function K. For point $\omega, \nu \in [0, \pi]$, define

$$K(\omega, \nu) = \mathrm{cov}\big(G(\omega), G(\nu)\big).$$

Then, for some $n \in \mathbb{N}$, we define the Gram matrix

$$M\big(\omega_1, \ldots, \omega_n\big) = \begin{pmatrix} K(\omega_1, \omega_1) & \cdots & K(\omega_1, \omega_n) \\ \vdots & \ddots & \vdots \\ K(\omega_n, \omega_1) & \cdots & K(\omega_n, \omega_n) \end{pmatrix}.$$

By [17, Proposition 8], we know that, for all fixed $n \in \mathbb{N}$, and all adjacent Φ and Ξ,

$$\left\| M^{-1/2}\big(\omega_1, \ldots, \omega_n\big) \begin{pmatrix} \Phi(\omega_1) - \Xi(\omega_1) \\ \vdots \\ \Phi(\omega_n) - \Xi(\omega_n) \end{pmatrix} \right\|_2 \leq \|\Phi - \Xi\|_{\mathcal{H}} \leq B,$$

where $\| \cdot \|_2$ denotes the Euclidean norm on $\mathbb{R}^n$. Define $\Delta(\Phi, \Xi; \omega_1, \ldots, \omega_n)$ to be the argument of the first norm above. Then, in particular,

$$\sup_{\substack{\Phi, \Xi \text{ s.t.} \\ \mathrm{Adj}_B(\Phi, \Xi) = 1}} \sup_{n < \infty} \sup_{(\omega_1, \ldots, \omega_n) \in [0, \pi]^n} \|\Delta(\Phi, \Xi; \omega_1, \ldots, \omega_n)\|_2 \leq B.$$

Then, the mechanism

$$\tilde{\Phi} = \Phi + B \frac{c(\delta)}{\epsilon} G$$

provides (ϵ, δ)-differential privacy to any n-point sampling of Φ with respect to a finite-dimensional σ-algebra adapted to that dimension. Using the argument of Proposition 6 in [17], we conclude that privacy holds in the space $H^1[0, \pi]$ with respect to the σ-algebra Σ_∞.