



Titre: Differentially private MIMO filtering for event streams
Title:

Auteurs: Jérôme Le Ny, & Meisam Mohammady
Authors:

Date: 2018

Type: Article de revue / Article

Référence: Le Ny, J., & Mohammady, M. (2018). Differentially private MIMO filtering for event streams. IEEE Transactions on Automatic Control, 63(1), 144-157.
Citation: <https://doi.org/10.1109/tac.2017.2713643>

 **Document en libre accès dans PolyPublie**
Open Access document in PolyPublie

URL de PolyPublie: <https://publications.polymtl.ca/2867/>
PolyPublie URL:

Version: Version finale avant publication / Accepted version
Révisé par les pairs / Refereed

Conditions d'utilisation: Tous droits réservés / All rights reserved
Terms of Use:

 **Document publié chez l'éditeur officiel**
Document issued by the official publisher

Titre de la revue: IEEE Transactions on Automatic Control (vol. 63, no. 1)
Journal Title:

Maison d'édition: IEEE
Publisher:

URL officiel: <https://doi.org/10.1109/tac.2017.2713643>
Official URL:

Mention légale: ©2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.
Legal notice:

Differentially Private MIMO Filtering for Event Streams

Jerome Le Ny, *Senior Member, IEEE*, and Meisam Mohammady

Abstract—Rigorous privacy-preserving mechanisms that can process and analyze dynamic data streams in real-time are required to encourage a wider adoption of many large-scale monitoring and control systems recording the detailed activities of their users, such as intelligent transportation systems, smart grids or smart buildings. Motivated by scenarios where signals originate from many sensors capturing privacy-sensitive events about individuals and several statistics of interest need to be continuously published in real-time, we consider the problem of designing multi-input multi-output (MIMO) systems processing event streams while providing certain differential privacy guarantees on the input signals. We show how to construct and optimize MIMO extensions of the zero-forcing mechanism, which we previously proposed for single-input single-output systems. Some of these extensions can take a statistical model of the input signals into account. We illustrate our privacy-preserving filter design methodology in two examples: privately monitoring and forecasting occupancy in a building equipped with multiple motion detection sensors; and analyzing the activity of a Markov chain model of a simple shared processing server.

Index Terms—Privacy, Filtering, Estimation, Multidimensional systems

I. INTRODUCTION

Privacy issues associated with social networking applications or monitoring and decision systems collecting personal data to operate are receiving an increasing amount of attention [3], [4]. Indeed, privacy concerns are causing delays or cancellations in the deployment of some smart power grids, location-based services, and various “Internet of Things” applications [5]. In order to encourage the adoption of these systems, which can provide important societal benefits, new tools are needed to provide clear privacy protection guarantees and allow users to balance utility with privacy rigorously [6].

Offering privacy guarantees for a system generally involves sacrificing some level of performance, and evaluating the resulting trade-offs rigorously requires a quantitative definition of privacy. Various such definitions have been proposed, such as disclosure risk [7] in statistics, k -anonymity [8], information-theoretic privacy [9], or conditions based on observability [10], [11]. However, in the last few years the notion of differential privacy has emerged essentially as a

standard specification [12], [13]. Intuitively, a system processing privacy-sensitive inputs from individuals is differentially private if its published outputs are not too sensitive to the data provided by any single participant. This definition is naturally linked to the notion of system gain for dynamical systems, see [14], [15]. A general concern in privacy-related research is how to provably limit the risk of privacy breaches caused by an adversary managing to link information published based on a sensitive dataset with other (often publicly) available information, as in [8], [16], [17] for example. One operational advantage of differential privacy compared to some other definitions is that it provides strong guarantees without involving the difficult task of modeling all the potentially available auxiliary information.

Differential privacy is a strong notion of privacy, but might require large perturbations to the published results of an analysis in order to hide individual data. This is especially true for applications where users continuously contribute data over time, and it is thus important to carefully design real-time mechanisms that can limit the impact on system performance of differential privacy requirements. Previous work on designing differentially private mechanisms for the publication of time series include [18], [19], but these mechanisms are not causal and hence not suited for real-time applications. The mechanism described in Section IV of this paper could also be interpreted as a dynamic, causal version of the matrix mechanism introduced in [20] for static databases. The papers [21]–[23] describe real-time differentially private mechanisms to approximate a few specific filters processing a stream of 0/1 variables, representing the occurrence of events attributed to individuals. For example, [21], [22] consider a private accumulator providing at each time the total number of events that occurred in the past. This paper is inspired by this scenario, and builds on our previous work on this problem in [14, Section IV] [15, Section VI]. Here we extend our analysis in particular to multi-input multi-output (MIMO) linear time-invariant (LTI) systems, which considerably broadens the applicability of the scheme to more common situations where multiple sensors monitor an environment and we wish to concurrently publish several statistics of interest. An application example is that of analyzing spatio-temporal records provided by networks of simple counting sensors, e.g., motion detectors in buildings or inductive-loop detectors in traffic information systems [24]. The literature on the differentially private processing of multi-dimensional time series is still very limited, but includes [25], which considers a single-input multiple-output filter where each output channel corresponds to a moving average filter with a different size for the averaging window, as well as

Manuscript received March 21, 2016; revised March 25, 2017; accepted May 22, 2017. Recommended by Associate Editor S. S. Saab.

This work was supported by NSERC under Grant RGPIN-435905-13. J. Le Ny is with the department of Electrical Engineering, Polytechnique Montreal, and with GERAD, Montreal, QC H3T 1J4, Canada. M. Mohammady was with the department of Electrical Engineering, Polytechnique Montreal. {jerome.le-ny, meisam.mohammady}@polymtl.ca

Preliminary versions of some of the results contained in this paper were presented at CDC 2013 [1] and CDC 2014 [2].

[26], which discusses an application to traffic monitoring.

To summarize, the contributions and organization of this paper are as follows. In Section II, we present a new generic scenario where we need to approximate a general MIMO LTI system by a mechanism offering differential privacy guarantees for the input signals. The formal definitions necessary to state the problem are also provided in that section. In Section III we perform some preliminary system sensitivity calculations that are necessary in the rest of the paper. Section IV presents a general approximation scheme for MIMO systems that provides differential privacy guarantees for the input signals. The design methodology and performance of the privacy-preserving filter are illustrated in Section V in the context of a building occupancy estimation problem. Note that Sections III-V provide a more detailed presentation of the theoretical and simulation results contained in our conference paper [2]. Finally, Section VI presents an additional privacy-preserving mechanism that can approximate the desired outputs more closely but requires the second-order statistics about the input signals to be publicly available. It extends to the MIMO case some of the results presented in our conference paper [1]. This section also illustrates the relationship between our problem and certain joint transmitter-receiver optimization problems arising in the communication systems literature [27], [28].

Notation: Throughout the paper we use the following standard abbreviations: LTI for Linear Time-Invariant, SISO for Single-Input Single-Output, SIMO for Single-Input Multiple-Output, and MIMO for Multiple-Input Multiple-Output. Unless specified otherwise, we consider discrete-time signals starting at time 0, dynamical systems or filters are assumed causal and (for simplicity of exposition) transfer functions or transfer matrices $G(z) = \sum_{t=0}^{\infty} G_t z^{-t}$ have real-valued coefficients. We fix a base probability space $(\Omega, \mathcal{F}, \mathbb{P})$. We denote the multivariate normal distribution with mean vector μ and covariance matrix Σ by $\mathcal{N}(\mu, \Sigma)$. For m an integer with $m \geq 1$, we write $[m] := \{1, \dots, m\}$. The notations $\|x\|_1 = \sum_{k=1}^p |x_k|$ and $\|x\|_2 = (\sum_{k=1}^p |x_k|^2)^{1/2}$ are used to denote the 1- and 2-norms in \mathbb{R}^p or \mathbb{C}^p , and we reserve the notation $\|\cdot\|$ for norms on signal and system spaces. $\text{col}(x_1, \dots, x_p)$ denotes a column vector or signal with components x_i , $i = 1, \dots, p$, and $\text{diag}(x_1, \dots, x_m)$ denotes a diagonal $m \times m$ matrix with the x_i 's on the diagonal. Finally, for H a Hermitian matrix, $H \succ 0$ means that it is positive definite and $H \succeq 0$ that it is positive semi-definite.

II. PROBLEM STATEMENT

A. Generic Scenario

We consider m sensors detecting events, with sensor i producing a discrete-time scalar signal $\{u_{i,t}\}_{t \geq 0} \in \mathbb{R}$, for $i \in [m]$. In a building monitoring scenario for example, the sensors could be motion detectors distributed at various locations and polled at regular intervals, with $u_{i,t} \in \mathbb{N}$ the number of detected events reported for period t . We denote u the resulting vector-valued signal, i.e., $u_t \in \mathbb{R}^m$. An LTI filter F , with m inputs and p outputs, takes input signals u from the sensors and publishes output signals $y = Fu$ of interest, with $y_t \in \mathbb{R}^p$. In our example, we might be

interested in continuously updating real-time estimates of the number of people in various parts of the building, as well as short- and medium-term occupancy forecasts, in order to optimize the operations of the Heating, Ventilation, and Air Conditioning (HVAC) system. The problem considered in this paper consists in replacing the filter F by a system processing the input u and producing a signal \hat{y} as close as possible to the desired output y (minimizing for example the mean square error $\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|y_t - \hat{y}_t\|_2^2]$), while providing some privacy guarantees to the individuals whose activities are captured by the input signals u . The privacy constraint is explained and quantified in the next subsection.

B. Differential Privacy

As mentioned in the introduction, a differentially private mechanism publishes data in a way that is not too sensitive to the presence or absence of a single individual. A formal definition of differential privacy is provided in Definition 1 below. In the previous building monitoring example, one goal of a privacy constraint could be to provide guarantees that an individual cannot be tracked too precisely from the published (typically aggregate) data. Indeed, Wilson and Atkeson [29] for example demonstrate how to track individual users in a building using a network of simple binary sensors such as motion detectors.

1) *Adjacency Relation:* Formally, we start by defining a symmetric binary relation, denoted Adj , on the space \mathcal{D} of datasets of interest, which captures what it means for two datasets to differ by the data of a single individual. Essentially, it is hard to determine from a differentially private output which of any two adjacent input datasets was used. Here, $\mathcal{D} := \{u : \mathbb{N} \mapsto \mathbb{R}^m\}$ is the set of input signals, and we have $\text{Adj}(u, u')$ if and only if we can obtain the signal u' from u by adding or subtracting the events corresponding to just one user. Motivated again by applications to spatial monitoring, we consider in this paper the following adjacency relation

$$\text{Adj}(u, u') \text{ iff} \tag{1}$$

$$\forall i \in [m], \exists t_i \in \mathbb{N}, \alpha_i \in \mathbb{R}, \text{ s.t. } u'_i - u_i = \alpha_i \delta_{t_i}, |\alpha_i| \leq k_i,$$

parametrized by a vector $k \in \mathbb{R}^m$ with components $k_i > 0$. According to (1), we wish to make it hard to detect deviations on each input signal component at any *single time period* (here δ_{t_i} denotes the discrete impulse signal with impulse at t_i), and by at most k_i . Let $e_i \in \mathbb{R}^m$ be the i^{th} standard basis vector, i.e., with coordinates $e_{ij} = \delta_{ij}$, $j = 1, \dots, m$. Then for two adjacent signals u, u' , we have with the notation in (1)

$$u' - u = \sum_{i=1}^m \alpha_i \delta_{t_i} e_i. \tag{2}$$

Note in passing that we could place additional constraints on k to capture additional knowledge about the problem, which can help design mechanisms with better performance, as we discuss later. For example, if we know that a given person can activate at most $l < m$ sensor and each k_i is 1, we can add the constraint $\|k\|_1 \leq l$.

The adjacency relation (1) extends the one considered in [15], [21]–[23] to the case of multiple input signals. It

places two constraints on the influence that an individual can have on the input data in order for our differentially private mechanisms to offer him guarantees. First, any given sensor can report an event due to the presence of the individual only once over the time interval of interest for our analysis. This is a sensible constraint in applications such as traffic monitoring with fixed motion detectors activated only once by each car traveling along a road, certain location-based services where a customer would check-in say at most once per day at each visited store, or certain health-monitoring applications where an individual would report a sickness or visit the emergency room only once. For a building monitoring scenario however, a single user could trigger the same motion detector several times over a relatively short period. A first solution to this issue consists in splitting the data stream of problematic sensors into several successive intervals, each considered as the signal from a new virtual sensor, so that an individual's data is present only once in each interval. A MIMO mechanism can then process such data and offer guarantees, addressing one of the main issues for the applicability of the model proposed in [21], [22]. However, increasing the number of inputs degrades the privacy guarantees or the output quality that we can provide. Hence in general no privacy guarantee will be offered to users who activate the same sensor too frequently. The second constraint imposed by (1) is that we bound the magnitude of an individual's contribution by k_i , but this is not really problematic in applications such as motion detection, where we can typically take $k_i = 1$.

2) *Definition of Differential Privacy*: Mechanisms that are differentially private [12], [30], [31, Definition 2.4] necessarily randomize their outputs, in such a way that they satisfy the property of Definition 1 below. Following these references, the term "mechanism" here simply refers to a randomized map M from some input space D of datasets to some output space R of published results. So for $d \in D$, $M(d, \cdot)$ is a measurable map from our sample space Ω to R . We follow however the standard practice of omitting the argument $\omega \in \Omega$ when we denote the random variable $M(d)$.

Definition 1. Let D be a space equipped with a symmetric binary relation denoted Adj , and let (R, \mathcal{M}) be a measurable space. Let $\epsilon, \delta \geq 0$. A mechanism $M : D \times \Omega \rightarrow R$ is (ϵ, δ) -differentially private for Adj (and \mathcal{M}) if for all $d, d' \in D$ such that $\text{Adj}(d, d')$, we have

$$\mathbb{P}(M(d) \in S) \leq e^\epsilon \mathbb{P}(M(d') \in S) + \delta, \quad \forall S \in \mathcal{M}. \quad (3)$$

If $\delta = 0$, the mechanism is said to be ϵ -differentially private.

This definition quantifies the allowed deviation for the output distribution of a differentially private mechanism, given any two adjacent input datasets d and d' . One can also show that it is impossible to design a statistical test with small error to decide if d or d' was used by a differentially private mechanism to produce its output [32], [33]. In this paper, the space D was defined as the space of input signals, and the adjacency relation considered is (1). The output space R is the space of output signals $R := \{y : \mathbb{N} \rightarrow \mathbb{R}^p\}$. Finally, a differentially private mechanism will consist of a system approximating our MIMO filter of interest F , as well as a

source of noise necessary to randomize the outputs and satisfy (3). We also refer the reader to [14], [15] for a technical discussion on the (standard) σ -algebra \mathcal{M} used on the output signal space to offer useful guarantees.

3) *Sensitivity*: Enforcing differential privacy can be done by randomly perturbing the published output of a system, at the price of reducing its utility or quality. Hence, we are interested in evaluating as precisely as possible the amount of noise necessary to make a mechanism differentially private. For this purpose, the following quantity plays an important role.

Definition 2. The ℓ_2 -sensitivity of a system G with m inputs and p outputs with respect to the adjacency relation Adj is defined by $\Delta_2^{m,p}G := \sup_{\text{Adj}(u,u')} \|Gu - Gu'\|_2$, where $\|Gv\|_2 = \sqrt{\sum_{t=-\infty}^{\infty} |(Gv)_t|_2^2}$.

4) *A Basic Differentially Private Mechanism*: The basic mechanism of Theorem 1 below (see [15]), extending [30], can be used to answer queries in a differentially private way. To present the result, we recall first the definition of the Q -function $Q(x) := \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-\frac{u^2}{2}} du$. Now for $\epsilon, \delta > 0$, let $\xi = Q^{-1}(\delta)$ and define $\kappa_{\delta,\epsilon} = \frac{1}{2\epsilon}(\xi + \sqrt{\xi^2 + 2\epsilon})$.

Theorem 1. Let G be a system with m inputs, p outputs, and with ℓ_2 -sensitivity $\Delta_2^{m,p}G$ with respect to an adjacency relation Adj . Then the mechanism $M(u) = Gu + w$ is (ϵ, δ) -differentially private with respect to Adj if $w = \{w_t\}_{t \geq 0}$ is a p -dimensional Gaussian white noise, with $w_t \sim \mathcal{N}(0, \sigma^2 I_p)$ for $\sigma \geq \kappa_{\delta,\epsilon} \times \Delta_2^{m,p}G$.

Proof: Let u and u' be adjacent signals. For all $T \geq 0$, we have

$$\sqrt{\sum_{t=0}^T |(Gu)_t - (Gu')_t|_2^2} \leq \Delta_2^{m,p}G,$$

which implies by [15, Theorem 3] that releasing the sequence $(M(u))_{0:T}$ is (ϵ, δ) -differentially private, for any $T \geq 0$. The result is then a consequence of [15, Lemma 2]. ■

The mechanism M described in Theorem 1, which provides a differentially-private version of a system G , is called an output-perturbation mechanism. We see that the amount of noise sufficient for differential privacy with this mechanism is proportional to the ℓ_2 -sensitivity of the filter and to $\kappa_{\delta,\epsilon}$, which can be shown to behave roughly as $O(\ln(1/\delta))^{1/2}/\epsilon$. Note that we add noise proportional to the sensitivity of the whole filter G independently on *each* output channel, even if G were diagonal say, otherwise trivial attacks that simply average a sufficient number of outputs could potentially detect the presence of an individual with high probability [15].

In conclusion we could obtain a differentially private mechanism for our original problem by simply adding a sufficient amount of noise to the output of our desired filter F , provided we can compute its sensitivity, which is the topic of the next section. However, it is possible in general to design mechanisms with much less overall noise than this output-perturbation scheme, as discussed in Sections IV and VI.

III. SENSITIVITY CALCULATIONS

For the following sensitivity calculations (see Definition 2), the \mathcal{H}_2 norm of an LTI system plays an important role. We recall its definition for a system G with m inputs

$$\|G\|_2^2 = \sum_{i=1}^m \|G\delta_0 e_i\|_2^2 = \frac{1}{2\pi} \int_{-\pi}^{\pi} \text{Tr}(G(e^{j\omega})^* G(e^{j\omega})) d\omega.$$

Writing $G(z) = [G_{ij}(z)]_{i,j}$ for the $p \times m$ transfer matrix, we also note that $\|G\|_2^2 = \sum_{i,j} \|G_{ij}\|_2^2$.

A. Exact solutions for the SIMO and Diagonal Cases

Generalizing the SISO scenario considered in [14], [15] to the case of a SIMO system, we have immediately the following theorem.

Theorem 2 (SIMO LTI system). *Let G be an LTI system with one input, p outputs and such that $\|G\|_2 < \infty$. For the adjacency relation (1), we have $\Delta_2^{1,p} G = k_1 \|G\|_2$.*

Proof: For u and u' adjacent

$$\|G(u - u')\|_2^2 = |\alpha_1|^2 \|G\delta_{t_1}\|_2^2 \leq k_1^2 \|G\|_2^2,$$

and the bound is attained if $|\alpha_1| = k_1$. ■

For a system G with multiple inputs, the special case where G is diagonal, i.e., its transfer matrix is $G(z) = \text{diag}(G_{11}(z), \dots, G_{mm}(z))$, also leads to a simple sensitivity result. Note that in this case, we have $\|G\|_2^2 = \sum_{i=1}^m \|G_{ii}\|_2^2$.

Theorem 3 (Diagonal LTI system). *Let G be a diagonal LTI system with m inputs and outputs, such that $\|G\|_2 < \infty$. For the adjacency relation (1), denoting $K = \text{diag}(k_1, \dots, k_m)$, we have*

$$\Delta_2^{m,m} G = \|GK\|_2 = \left(\sum_{i=1}^m \|k_i G_{ii}\|_2^2 \right)^{1/2}.$$

Proof: If G is diagonal, then for u and u' adjacent, we have from (2)

$$\begin{aligned} \|G(u - u')\|_2^2 &= \left\| \sum_{i=1}^m \alpha_i G\delta_{t_i} e_i \right\|_2^2 \\ &= \|\text{col}(\alpha_1 g_{11} * \delta_{t_1}, \dots, \alpha_m g_{mm} * \delta_{t_m})\|_2^2, \end{aligned}$$

where g_{ii} denotes the impulse response of G_{ii} . Hence

$$\begin{aligned} \|G(u - u')\|_2^2 &= \sum_{i=1}^m \|\alpha_i g_{ii} * \delta_{t_i}\|_2^2 \\ &= \sum_{i=1}^m |\alpha_i|^2 \|G_{ii}\|_2^2, \end{aligned}$$

and $|\alpha_i| \leq k_i$, for all i . Again the bound is attained if $|\alpha_i| = k_i$ for all i . ■

B. Upper and Lower Bound for the general MIMO Case

For MISO or general MIMO systems, the sensitivity calculations are no longer so straightforward, because the impulses on the various input channels, obtained from the difference of two adjacent signals u, u' , all possibly influence any given output. Still, the following result provides simple bounds on the sensitivity.

Theorem 4. *Let G be an LTI system with m inputs, p outputs and such that $\|G\|_2 < \infty$. For the adjacency relation (1), denoting $K = \text{diag}(k_1, \dots, k_m)$ and $|k|_2 = (\sum_{i=1}^m k_i^2)^{1/2}$, we have*

$$\|GK\|_2 \leq \Delta_2^{m,p} G \leq |k|_2 \|G\|_2. \quad (4)$$

Proof: We have $G(u - u') = \sum_{i=1}^m \alpha_i G\delta_{t_i} e_i$, and moreover $\|G\|_2^2 = \sum_{i=1}^m \|G\delta_{t_i} e_i\|_2^2$ by definition. For the upper bound, we can write

$$\begin{aligned} \|G(u - u')\|_2 &= \left\| \sum_{i=1}^m \alpha_i G\delta_{t_i} e_i \right\|_2 \leq \sum_{i=1}^m |\alpha_i| \|G\delta_{t_i} e_i\|_2 \\ &\leq |k|_2 \left(\sum_{i=1}^m \|G\delta_{t_i} e_i\|_2^2 \right)^{1/2}, \end{aligned}$$

where the last inequality results from the Cauchy-Schwarz inequality.

For the lower bound, let us first take $u' \equiv 0$. Then consider an adjacent signal u with a single discrete impulse of height k_i at time t_i on each input channel i , for $i = 1, \dots, m$, with $t_1 < t_2 < \dots < t_m$. Let $\eta > 0$. Denote the ‘‘columns’’ of G as G_i for $i = 1, \dots, m$, i.e., $G u = \sum_{i=1}^m G_i u_i$. Since $\|G\|_2 < \infty$, $\|G_i u_i\|_2 < \infty$, and hence $|(G_i u_i)_t| \rightarrow 0$ as $t \rightarrow \infty$. Hence by taking $t_{i+1} - t_i$ large enough for each $1 \leq i \leq m - 1$, i.e., waiting for the effect of impulse i on the output to be sufficiently small, we can choose the signal u such that

$$\|G u\|_2^2 = \left\| \sum_{i=1}^m G_i u_i \right\|_2^2 \geq \sum_{i=1}^m k_i^2 \|G\delta_{t_i} e_i\|_2^2 - \eta.$$

Since this is true for any $\eta > 0$ and $\|G\delta_{t_i} e_i\|_2^2 = \|G_i\|_2^2$, we get $(\Delta_2^{m,p} G)^2 \geq \|GK\|_2^2 = \sum_{i=1}^m k_i^2 \|G_i\|_2^2$. ■

Note that if $k_1 = \dots = k_m$, the upper bound on the sensitivity is $k_1 \|G\|_2 \sqrt{m}$. We can compare this bound to the situation where G is diagonal, in which case the sensitivity is exactly $k_1 \|G\|_2$ from Theorem 3. The following example shows that the upper bound of Theorem 4 cannot be improved for the general MISO or MIMO case.

Example 1. Consider the MISO system $G(z) = [G_{11}(z), \dots, G_{1m}(z)]$, with $g_{1i} = \delta_{\tau_i}$ the impulse response of G_{1i} , for some times τ_1, \dots, τ_m . Then $\|G\|_2^2 = m$. Now let $u' \equiv 0$ and $u = \sum_{i=1}^m \delta_{t_i} e_i$, so that u and u' are adjacent, with $k_1 = \dots = k_m = 1$, and moreover let us choose the times t_i such that $\tau_i + t_i$ is a constant, i.e., take $t_i = \kappa - \tau_i$ for some $\kappa \geq \max_i \{\tau_i\}$. Then $G u = \sum_{i=1}^m g_{1i} * u_i = m \delta_\kappa$, and so $\|G u\|_2^2 = m^2$. This shows that the upper bound of Theorem 4 is tight in this case. Note that this happens because all the events of the signal u influence the output at

the same time. Indeed, if the times $\tau_i + t_i$ are all distinct, then we get $\|Gu\|_2^2 = m$.

C. Exact solution for the MIMO Case

For completeness, we give in this subsection an exact expression for the sensitivity of a general finite-dimensional MIMO LTI filter. Let G be a stable finite-dimensional LTI system with m inputs, p outputs and state space representation

$$\begin{aligned} x_{t+1} &= Ax_t + Bu_t \\ y_t &= Cx_t + Du_t, \end{aligned} \quad (5)$$

with $x_0 = 0$. Recall the definition of the observability Gramian P_0 , which is the unique positive semi-definite solution of the equation

$$A^T P_0 A - P_0 + C^T C = 0.$$

Let B_i, D_i be the i^{th} column of the matrix B and D respectively, for $i = 1, \dots, m$. Finally, define for $i, j \in \{1, \dots, m\}$, $i \neq j$, and $\tau \in \mathbb{Z}$

$$S_{ij}^\tau = \begin{cases} B_i^T (A^{\tau-1})^T C^T D_j + B_i^T (A^\tau)^T P_0 B_j, & \text{if } \tau > 0 \\ D_i^T D_j + B_i^T P_0 B_j, & \text{if } \tau = 0 \\ D_i^T C A^{|\tau|-1} B_j + B_i^T P_0 A^{|\tau|} B_j, & \text{if } \tau < 0. \end{cases} \quad (6)$$

Theorem 5. *Let G be a stable finite-dimensional LTI system with m inputs, p outputs and state space representation (5). Then, for the adjacency relation (1), we have*

$$(\Delta_2^{m,p} G)^2 = \|GK\|_2^2 + \sum_{\substack{i,j=1 \\ i \neq j}}^m k_i k_j \left(\sup_{\tau \in \mathbb{Z}} |S_{ij}^\tau| \right). \quad (7)$$

Proof: In view of (2), we have

$$\Delta_2^{m,p} G = \sup_{|\alpha_i| \leq k_i, t_i \geq 0} \left\| \sum_{i=1}^m \alpha_i G \delta_{t_i} e_i \right\|_2.$$

For $y_i = G \delta_{t_i} e_i$ and $y = \sum_{i=1}^m \alpha_i y_i$, we have

$$\begin{aligned} \|y\|_2^2 &= \sum_{t=0}^{\infty} \left| \sum_{i=1}^m \alpha_i y_{i,t} \right|^2 \\ &= \sum_{t=0}^{\infty} \sum_{i=1}^m \alpha_i^2 |y_{i,t}|^2 + \sum_{t=0}^{\infty} \sum_{\substack{i,j=1 \\ i \neq j}}^m \alpha_i \alpha_j y_{i,t}^T y_{j,t} \\ &\leq \|GK\|_2^2 + \sum_{\substack{i,j=1 \\ i \neq j}}^m k_i k_j \left| \sum_{t=0}^{\infty} y_{i,t}^T y_{j,t} \right|, \end{aligned}$$

where $K = \text{diag}(k_1, \dots, k_m)$ and the bound can be attained by taking $\alpha_i \in \{-k_i, k_i\}$, depending on the sign of $S_{ij}^{t_i, t_j} := \sum_{t=0}^{\infty} y_{i,t}^T y_{j,t}$.

Next, we derive the more explicit expression for the second term $S_{ij}^{t_i, t_j}$, given in the theorem. First,

$$y_{i,t} = \begin{cases} 0, & t < t_i, \\ D_i, & t = t_i \\ C A^{t-t_i-1} B_i, & t > t_i. \end{cases}$$

Then if $t_i = t_j$, we find that

$$S_{ij}^{t_i, t_j} = D_i^T D_j + B_i^T P_0 B_j,$$

with $P_0 = \sum_{t=0}^{\infty} (A^t)^T C^T C A^t$ the observability Gramian. If $t_i < t_j$, then

$$S_{ij}^{t_i, t_j} = B_i^T (A^{t_j-t_i-1})^T C^T D_j + B_i^T (A^{t_j-t_i})^T P_0 B_j.$$

The case $t_i > t_j$ is symmetric. Hence we find that $S_{ij}^{t_i, t_j}$ depends only on the difference $\tau = t_i - t_j$, and our notation (6) corresponds to $S_{ij}^\tau := S_{ij}^{t_i, t_i+\tau}$. ■

In (7), the maximization over inter-event times τ still needs to be performed and depends on the parameters of the specific system G . This result could be used to evaluate carefully the amount of noise necessary in an output perturbation mechanism, but unfortunately it seems too unwieldy at this point to be used in more advanced mechanism optimization schemes, such as the ones discussed in the next sections.

Still, the expression (7) provides some intuition about the way the system dynamics influence its sensitivity. In particular, the second term in (7) can give insight into the gap between the sensitivity and the lower bound in (4). Note from the expression of S_{ij}^τ in (6) that one way to decrease the sensitivity of G is to increase sufficiently the required time $|t_i - t_j|$ between the events contributed by a single user, in order for $\|A^{t_i-t_j}\|_2$ to be small enough. Hence, a lower bound on inter-event times in different streams could be introduced in the adjacency relation to reduce a system's sensitivity. This would weaken the differential privacy guarantee but help in the design of mechanisms with better performance. Another possibility would be to have a privacy-preserving mechanism simply ignore events from a given user as long as the lower bound on inter-event times is not reached.

IV. ZERO-FORCING MIMO MECHANISMS

Using the sensitivity calculations of Section III, we can now design differentially private mechanisms to approximate a given filter F , as discussed in Section II-A. The mechanisms described below generalize to the MIMO case some ideas introduced in [14]. The general approximation architecture considered is described on Fig. 1, with the filters G and H to design. Following Theorem 1, by introducing a Gaussian white noise signal w with variance $(\kappa_{\delta, \epsilon} \Delta_2 G)^2 I$, the signal v is made differentially private. The signal \hat{y} computed from v is differentially private no matter what the system H is, see [15, Theorem 1]. We then construct G and H to minimize the Mean Square Error (MSE) between y and \hat{y}

$$e_{mse}(G, H) = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [|y_t - \hat{y}_t|_2^2]. \quad (8)$$

In this section, we fix the system H to be of the form $H = FL$, with L a left inverse of the pre-filter G , i.e., such that $L(z)G(z) = I_m$, assuming such an L exists. We call the resulting mechanisms Zero-Forcing Equalization (ZFE) mechanisms, and we denote $e_{mse}^{ZFE}(G, L) := e_{mse}(G, H)$ in this case. It was shown in [14] for the SISO case that this setup allows significant performance improvements compared to the output-perturbation mechanism. The latter is recovered when $G = F$ and H is the identity.

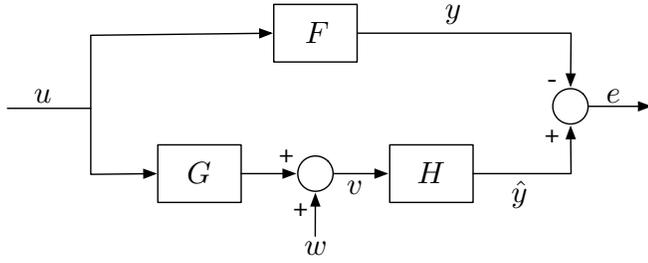


Fig. 1. Approximation setup for differentially private filtering. w is a noise signal guaranteeing that v and hence \hat{y} are differentially private signals.

A. SIMO system approximation

First, let us assume that F on Fig. 1 is a SIMO filter, with p outputs. Hence for all ω , $F(e^{j\omega})$ is a complex-valued p -dimensional vector. The next theorem gives a condition on the optimal filters G and L minimizing (8) for the ZFE architecture described above, which allows us to construct them, and shows that one can take these optimal ZFE filters to be SISO.

Theorem 6. *Let F be a SIMO LTI system with $\|F\|_2 < \infty$. For any SIMO LTI system G such that $\|G\|_2 < \infty$ and with a left-inverse L , we have*

$$e_{mse}^{ZFE}(G, L) \geq k_1^2 \kappa_{\delta, \epsilon}^2 \left(\frac{1}{2\pi} \int_{-\pi}^{\pi} |F(e^{j\omega})|_2 d\omega \right)^2. \quad (9)$$

Suppose moreover that F satisfies the Paley-Wiener condition $\frac{1}{2\pi} \int_{-\pi}^{\pi} \ln |F(e^{j\omega})|_2 d\omega > -\infty$. Then there exists a SISO LTI system G with causal inverse L such that $\|G\|_2 < \infty$,

$$|G(e^{j\omega})|^2 = |F(e^{j\omega})|_2 \text{ for almost every } \omega \in [-\pi, \pi), \quad (10)$$

and any such G, L achieves the lower bound (9) and satisfies $\|G\|_2 = \|FL\|_2$.

Theorem 6 generalizes [15, Theorem 8] to the SIMO case. Finding G SISO satisfying (10), $\|G\|_2 < \infty$ and with causal inverse $L = G^{-1}$ is a classical spectral factorization problem, see [34] for example for more details on how to construct G from this frequency-domain condition.

Proof: Consider a first stage $G(z) = \text{col}(G_1(z), \dots, G_q(z))$ with $\|G\|_2 < \infty$, taking the input signal u and producing q intermediate outputs that must be perturbed, for some $q \geq 1$. The second stage is $H = FL$, with $L(z) = [L_1(z), \dots, L_q(z)]$ satisfying $\sum_{i=1}^q L_i(z)G_i(z) = 1$. Let us also define the transfer functions M_i , $i = 1, \dots, q$, such that $M_i(z) = L_i(z^{-1})$, hence $M_i(e^{j\omega}) = L_i(e^{j\omega})^*$, and thus in particular

$$|M_i(e^{j\omega})|^2 = |L_i(e^{j\omega})|^2, i = 1, \dots, q, \quad (11)$$

$$\text{and } \sum_{i=1}^q M_i(e^{j\omega})^* G_i(e^{j\omega}) = 1. \quad (12)$$

From Theorem 2, the sensitivity of the first stage for input signals that are adjacent according to (1) is $k_1 \|G\|_2$. Hence, according to Theorem 1, adding a white Gaussian noise w to the output of G with covariance matrix $k_1^2 \kappa_{\delta, \epsilon}^2 I_q$ is

sufficient to ensure that the signal v on Fig. 1 is differentially private. The MSE for this mechanism can be expressed as

$$e_{mse}^{ZFE}(G, L) = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [|(Fu)_t - (FL(Gu + w))_t|^2]$$

$$e_{mse}^{ZFE}(G, L) = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [|(FLw)_t|^2]$$

$$e_{mse}^{ZFE}(G, L) = k_1^2 \kappa_{\delta, \epsilon}^2 \|G\|_2^2 \|FL\|_2^2,$$

using the fact that w is white Gaussian noise and a property of the \mathcal{H}_2 norm. We are thus led to consider the minimization of $\|FL\|_2^2 \|G\|_2^2$ over the pre-filters G such that $\|G\|_2 < \infty$ and $\|FL\|_2 < \infty$. We have

$$\begin{aligned} & \|FL\|_2^2 \|G\|_2^2 \\ &= \frac{1}{2\pi} \int_{-\pi}^{\pi} \text{Tr}(L^*(e^{j\omega}) F^*(e^{j\omega}) F(e^{j\omega}) L(e^{j\omega})) d\omega \times \\ & \quad \frac{1}{2\pi} \int_{-\pi}^{\pi} \text{Tr}(G^*(e^{j\omega}) G(e^{j\omega})) d\omega \\ &= \frac{1}{2\pi} \int_{-\pi}^{\pi} |F(e^{j\omega})|_2^2 |L(e^{j\omega})|_2^2 d\omega \times \frac{1}{2\pi} \int_{-\pi}^{\pi} |G(e^{j\omega})|_2^2 d\omega \\ &= \frac{1}{2\pi} \int_{-\pi}^{\pi} |F(e^{j\omega})|_2^2 |M(e^{j\omega})|_2^2 d\omega \times \frac{1}{2\pi} \int_{-\pi}^{\pi} |G(e^{j\omega})|_2^2 d\omega, \end{aligned}$$

where in the last equality we used (11).

Now consider the inner product $\langle f, g \rangle = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(e^{j\omega})^* g(e^{j\omega}) d\omega$ on the space of 2π -periodic functions with values in \mathbb{C}^q . By the Cauchy-Schwarz inequality for this inner product applied to the functions $\omega \mapsto |F(e^{j\omega})|_2 M(e^{j\omega})$ and $\omega \mapsto G(e^{j\omega})$, we obtain the following bound

$$\|FL\|_2^2 \|G\|_2^2 \geq \left(\frac{1}{2\pi} \int_{-\pi}^{\pi} |F(e^{j\omega})|_2 \sum_{i=1}^q M_i^*(e^{j\omega}) G_i(e^{j\omega}) d\omega \right)^2,$$

i.e., using (12),

$$\|FL\|_2^2 \|G\|_2^2 \geq \left(\frac{1}{2\pi} \int_{-\pi}^{\pi} |F(e^{j\omega})|_2 d\omega \right)^2.$$

Also, the two sides in the Cauchy-Schwarz inequality are equal, i.e., the bound is attained, if

$$G(e^{j\omega}) = |F(e^{j\omega})|_2 M(e^{j\omega}). \quad (13)$$

Note that this condition does not depend on q . Hence we can simply take $q = 1$ and multiply (13) by $G(e^{j\omega})^*$ to get $|G(e^{j\omega})|^2 = |F(e^{j\omega})|_2$, which is (10). Since $|F(e^{j\omega})|_2$ is a nonnegative function on the unit circle, integrable on $[-\pi, \pi]$ (consequence of $\|F\|_2 < \infty$), if it satisfies the Paley-Wiener condition, it has indeed a (minimum phase) spectral factor G with causal inverse satisfying (10) almost everywhere [34, p. 199]. Finally, $\|FL\|_2 = \|G\|_2$ is a straightforward consequence of the stronger condition (10). ■

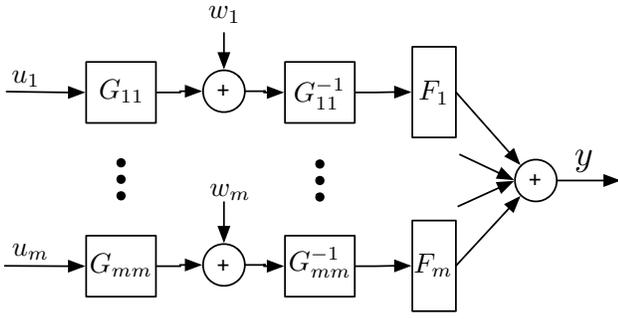


Fig. 2. (Suboptimal) ZFE mechanism for a MIMO system $Fu = \sum_{i=1}^m F_i u_i$, and a diagonal pre-filter $G(z) = \text{diag}(G_{11}(z), \dots, G_{mm}(z))$. Here $F_i(z)$ is a $p \times 1$ transfer matrix, for $i = 1, \dots, m$.

B. MIMO system approximation

Let us now assume that F has $m > 1$ inputs. We write $F(z) = [F_1(z), \dots, F_m(z)]$, with F_i a $p \times 1$ transfer matrix. In this case, in view of the complicated expression (7) for the sensitivity of a MIMO filter, we only provide a suboptimal ZFE mechanism, together with an estimate of the gap between the performance of our mechanism and that of the optimal ZFE mechanism.

1) *Diagonal Pre-filter Optimization*: The idea is to restrict our design to pre-filters G that are $m \times m$ and diagonal, with squared sensitivity equal to $(\Delta_2^{m,m} G)^2 = \|GK\|_2^2 = \sum_{i=1}^m \|k_i G_{ii}\|_2^2$, where $K = \text{diag}(k_1, \dots, k_m)$, according to Theorem 3. We then consider the architecture shown on Fig. 2, where the signal w is a white Gaussian noise with covariance matrix $\kappa_{\delta,\epsilon}^2 \|GK\|_2^2 I_m$, making the signals $G_{ii}u_i + w_i, i = 1, \dots, m$, differentially private. The problem of optimizing the pre-filters G_{ii} can then be seen as designing m SIMO mechanisms.

Theorem 7. Let $F = [F_1, \dots, F_m]$ be a MIMO LTI system with $\|F\|_2 < \infty$. For any diagonal LTI system $G(z) = \text{diag}(G_{11}(z), \dots, G_{mm}(z))$ such that $\|G\|_2 < \infty$, with inverse L , we have

$$e_{mse}^{ZFE}(G, L) \geq \kappa_{\delta,\epsilon}^2 \left(\frac{1}{2\pi} \int_{-\pi}^{\pi} \sum_{i=1}^m k_i |F_i(e^{j\omega})|_2 d\omega \right)^2. \quad (14)$$

If moreover each F_i satisfies the Paley-Wiener condition $\frac{1}{2\pi} \int_{-\pi}^{\pi} \ln |F_i(e^{j\omega})|_2 d\omega > -\infty$, this lower bound is attained by some systems G_{ii} with causal inverses G_{ii}^{-1} , such that $\|k_i G_{ii}\|_2 = \|F_i G_{ii}^{-1}\|_2 < \infty$ and $|G_{ii}(e^{j\omega})|^2 = \frac{1}{k_i} |F_i(e^{j\omega})|_2$, for almost every $\omega \in [-\pi, \pi)$.

In other words, the best diagonal pre-filter for the MIMO ZFE mechanism can be obtained from m spectral factorizations of the functions $\omega \mapsto \frac{1}{k_i} |F_i(e^{j\omega})|_2, i = 1, \dots, m$.

Proof: Following the same reasoning as in the proof of Theorem 6, the MSE for the mechanism of Fig. 2 can be expressed as

$$e_{mse}^{ZFE}(G) = \kappa_{\delta,\epsilon}^2 \|GK\|_2^2 \|FG^{-1}\|_2^2, \quad (15)$$

with $G^{-1}(z) = \text{diag}(G_{11}(z)^{-1}, \dots, G_{mm}(z)^{-1})$, assuming that the inverses exist. Now remark that $\|FG^{-1}\|_2^2 =$

$\frac{1}{2\pi} \int_{-\pi}^{\pi} \sum_{i=1}^m \frac{|F_i(e^{j\omega})|_2^2}{|G_{ii}(e^{j\omega})|^2} d\omega$. Hence from the Cauchy-Schwarz inequality again, we obtain the lower bound

$$e_{mse}^{ZFE}(G) \geq \kappa_{\delta,\epsilon}^2 \left(\frac{1}{2\pi} \int_{-\pi}^{\pi} \sum_{i=1}^m \frac{|F_i(e^{j\omega})|_2}{|G_{ii}(e^{j\omega})|} |k_i G_{ii}(e^{j\omega})| d\omega \right)^2$$

$$e_{mse}^{ZFE}(G) \geq \kappa_{\delta,\epsilon}^2 \left(\frac{1}{2\pi} \int_{-\pi}^{\pi} \sum_{i=1}^m k_i |F_i(e^{j\omega})|_2 d\omega \right)^2,$$

and this bound is attained if $k_i |G_{ii}(e^{j\omega})| = \frac{|F_i(e^{j\omega})|_2}{|G_{ii}(e^{j\omega})|}$, i.e., $k_i |G_{ii}(e^{j\omega})|^2 = |F_i(e^{j\omega})|_2$, for $i = 1, \dots, m$. ■

Remark 1. Note that the integrand on the right-hand side of (14) can be written

$$\sum_{i=1}^m k_i |F_i(e^{j\omega})|_2 := \|F(e^{j\omega})K\|_{2,1},$$

where $\|\cdot\|_{2,1}$ is the so-called $L_{2,1}$ or R_1 matrix norm, and appears in [35] for example.

2) *Comparison with Non-Diagonal Pre-filters*: For F a general MIMO system, it is possible that we could achieve a better performance with a ZFE mechanism where G is not diagonal, i.e., by combining the inputs before adding the privacy-preserving noise. To provide a better understanding of how much could potentially be gained by carrying out this more involved optimization over general pre-filters G rather than just diagonal pre-filters, the following theorem provides a lower bound on the MSE achievable by any ZFE mechanism.

Theorem 8. Let $F = [F_1, \dots, F_m]$ be a MIMO LTI system with $\|F\|_2 < \infty$. For any $m \times m$ LTI system G such that $\|G\|_2 < \infty$, with left inverse L such that $\|FL\|_2 < \infty$, we have

$$e_{mse}^{ZFE}(G, L) \geq \kappa_{\delta,\epsilon}^2 \left(\frac{1}{2\pi} \int_{-\pi}^{\pi} \|F(e^{j\omega})K\|_* d\omega \right)^2, \quad (16)$$

where $\|F(e^{j\omega})K\|_*$ denotes the nuclear norm (i.e., sum of singular values) of the matrix $F(e^{j\omega})K$.

The lower bound (16) on the MSE achievable with a general pre-filter in a ZFE mechanism should be compared to the performance (14) that can actually be achieved with diagonal pre-filters. Note that these bounds coincide for $m = 1$. For $m > 1$, the gap depends on the difference between the integrals of the $L_{2,1}$ norm and the nuclear norm of $F(e^{j\omega})K$, see Fig. 7 for an illustration.

Proof: With $K = \text{diag}(k_1, \dots, k_m)$ as usual, we define $\check{G} = GK$ and $\check{L} = K^{-1}L$, so that $\check{L}\check{G} = I$. Let $\check{F} = FK$. With the lower bound of Theorem 4, designing a ZFE mechanism based on sensitivity as above would require adding a noise with variance at least $\kappa_{\delta,\epsilon}^2 \|\check{G}\|_2^2$. This would lead to an

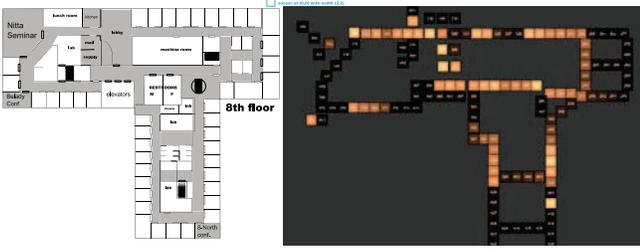


Fig. 3. Left: plan of one of the two floors in the MERL building used for the sensor network experiment [38]. The shaded areas are hallways, lobbies and meeting rooms equipped with binary motion detection sensors, placed a few meters apart and recording events roughly every second. Right: a figure taken from [37] shows a visualisation of a crowd movement during a fire drill.

MSE at least equal to $\kappa_{\delta,\epsilon}^2 \|\check{G}\|_2^2 \|\check{F}\check{L}\|_2^2$. Now note that

$$\begin{aligned} \|\check{F}\check{L}\|_2^2 &= \frac{1}{2\pi} \int_{-\pi}^{\pi} \text{Tr}(\check{F}(e^{j\omega})\check{L}(e^{j\omega})\check{L}(e^{j\omega})^*\check{F}(e^{j\omega})^*)d\omega \\ &= \frac{1}{2\pi} \int_{-\pi}^{\pi} \text{Tr}(\check{F}(e^{j\omega})^*\check{F}(e^{j\omega})\check{L}(e^{j\omega})\check{L}(e^{j\omega})^*)d\omega \\ &= \frac{1}{2\pi} \int_{-\pi}^{\pi} \text{Tr}(A(e^{j\omega})^2\check{L}(e^{j\omega})\check{L}(e^{j\omega})^*)d\omega \\ &= \frac{1}{2\pi} \int_{-\pi}^{\pi} \text{Tr}(A(e^{j\omega})\check{L}(e^{j\omega})\check{L}(e^{j\omega})^*A(e^{j\omega}))d\omega, \end{aligned}$$

where for all ω , $A(e^{j\omega})$ is the unique Hermitian positive-semidefinite square root of $\check{F}(e^{j\omega})^*\check{F}(e^{j\omega})$, i.e., $A(e^{j\omega})^2 = \check{F}(e^{j\omega})^*\check{F}(e^{j\omega})$ [36, Theorem 7.2.6]. Then, once again from the Cauchy-Schwarz inequality, now for the inner product $\langle M, N \rangle = \frac{1}{2\pi} \int_{-\pi}^{\pi} \text{Tr}(M(e^{j\omega})^*N(e^{j\omega}))d\omega$,

$$\begin{aligned} \|GK\|_2^2 \|FL\|_2^2 &= \|\check{G}\|_2^2 \|\check{F}\check{L}\|_2^2 \\ &= \left[\frac{1}{2\pi} \int_{-\pi}^{\pi} \text{Tr}(\check{G}(e^{j\omega})^*\check{G}(e^{j\omega}))d\omega \right] \\ &\quad \times \left[\frac{1}{2\pi} \int_{-\pi}^{\pi} \text{Tr}(A(e^{j\omega})\check{L}(e^{j\omega})\check{L}(e^{j\omega})^*A(e^{j\omega}))d\omega \right] \\ &\geq \left(\frac{1}{2\pi} \int_{-\pi}^{\pi} \text{Tr}(A(e^{j\omega})\check{L}(e^{j\omega})\check{G}(e^{j\omega}))d\omega \right)^2 \end{aligned}$$

$$\text{and so } e_{mse}^{ZFE}(G) \geq \kappa_{\delta,\epsilon}^2 \left(\frac{1}{2\pi} \int_{-\pi}^{\pi} \|F(e^{j\omega})K\|_* d\omega \right)^2,$$

where $\|F(e^{j\omega})K\|_* = \text{Tr}(A(e^{j\omega}))$ denotes the nuclear norm of the matrix $F(e^{j\omega})K$. ■

V. APPLICATION TO PRIVACY-PRESERVING ESTIMATION OF BUILDING OCCUPANCY

In this section we illustrate the design process and the performance of the ZFE mechanism in the context of an application to filtering and forecasting occupancy-related events in an office building equipped with motion detection sensors. As mentioned in Section II-B, such sensor networks raise privacy concerns since some occupants could potentially be tracked from the published information, especially when it is correlated with public information such as the location of their office. Since the amount of private information leakage

depends on the output signals the system aims to generate, we adjust the privacy-preserving noise level based on the filter specification using the ZFE mechanism. As an example, we simulate the outputs of a 3×15 MIMO filter processing input signals collected during a sensor network experiment carried out at the Mitsubishi Electric Research Laboratories (MERL), described in [38] and on Fig. 3. We refer the reader to [37] for examples of identification of individual trajectories from this dataset.

The original dataset contains the traces of 213 sensors placed a few meters apart and spread over two floors of a building, where each sensor recorded with millisecond accuracy over a year the exact times at which they detected some motion. For illustration purposes we downsampled the dataset in space and time, summing all the events recorded by several sufficiently close sensors over 3 minute intervals. From this step, we obtained 15 input signals $u_i, i = 1, \dots, 15$, corresponding to 15 spatial zones (each zone covered by a cluster of about 14 sensors), with a discrete-time period corresponding to 3 minutes and where $u_{i,t}$ is the number of events detected by all the sensors in zone i during period t . Let us assume say that during a given discrete-time period, a single individual can activate at most 4 sensors in any zone, hence $k_i = 4$ for $1 \leq i \leq 15$. Moreover, we need to assume that a single individual only activates the sensors in a given zone once over the time interval for which we wish to provide differential privacy. Section II-B discussed how to relax this requirement by splitting the input data into successive time windows and creating additional input channels.

Suppose that we want to compute simultaneously and in real-time the following three outputs from the 15 input signals

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} = \begin{bmatrix} f_1(z)\mathbf{1}_{1 \times 5} & \mathbf{0}_{1 \times 10} \\ \mathbf{0}_{1 \times 4} & f_2(z)\mathbf{1}_{1 \times 8} & \mathbf{0}_{1 \times 3} \\ & & f_3(z) \end{bmatrix} u, \quad (17)$$

where

- y_1 is the sum of the simple moving averages over the past 60 min for zones 1 to 5, i.e., $f_1(z) = \frac{1}{20} \sum_{k=0}^{19} z^{-k}$,
- y_2 is $\sum_{i=5}^{12} f_2 u_i$, with f_2 a low-pass filter with Gaussian shaped finite impulse response of length 20, obtained using Matlab's function `gaussdesign(0.5, 2, 10)`.
- y_3 is the scalar output of a 1×15 MISO filter f_3 designed to forecast at each period t the average total number of events per time-period that will occur in the whole building during the window $[t+60 \text{ min}, t+90 \text{ min}]$. This filter was constructed by identifying an ARMAX model [39] between the 15 inputs (plus a scalar white noise) and the desired output, with the calibration done using one part of the dataset. The model chosen takes the form

$$y_{3,t} = \sum_{i=1}^4 a_i y_{3,t-i} + b_0 u_t + b_1 u_{t-1} + e_t + c_1 e_{t-1},$$

where a_1, \dots, a_4 and b_0, b_1 are scalar and row vectors respectively forming the filter f_3 , c_1 is a scalar and e_t is a zero-mean standard white noise input postulated by the ARMAX model for system identification purposes.

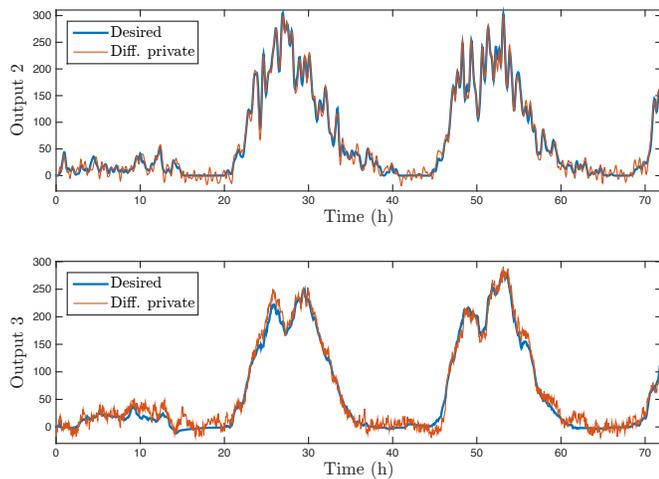


Fig. 4. Sample paths over 72 hours (the sampling period is 3 min) for the outputs 2 and 3 of our differentially private approximation of filter (17), shown together with the desired outputs. The privacy-related parameters are $\epsilon = \ln 5$, $\delta = 0.05$, $k_i = 4$ for $1 \leq i \leq 15$. The average ℓ_2 norm per period for the error signal corresponding to each figure is 9.3 and 14.5 respectively.

Fig. 4 shows sample paths over a 72h period of the 2nd and 3rd outputs of the desired filter and of its $(\ln 5, 0.05)$ -differentially private approximation obtained using the ZFE MIMO mechanism with diagonal pre-filter. The 15 optimal pre-filters were obtained approximately via least-squares fit of $\sqrt{|F_i(e^{j\omega})|_2}$ with negligible approximation error (using Matlab's function `yulewalk` implementing the Yule-Walker method [40]), rather than true spectral factorization as mentioned in Theorem 7. One apparent feature of the privacy-preserving outputs seen on Fig. 4 is that the noise level is independent of the size of the desired output signal, hence low signal values tend to be easily buried in the noise. This is one drawback of mechanisms relying on global sensitivity measures and additive noise. Another noticeable element is the fact that the noise remaining on each output can have quite different characteristics depending on the desired filter F , with the post-filter FG^{-1} removing more high-frequency components on the second output than on the third.

Referring to (14) for the performance achieved with diagonal pre-filters as here and the lower bound (16), we obtain for the filter F designed in our example

$$\frac{1}{2\pi} \int_{-\pi}^{\pi} \sum_{i=1}^{15} k_i |F_i(e^{j\omega})|_2 d\omega \approx 2.2 \times \frac{1}{2\pi} \int_{-\pi}^{\pi} \|F(e^{j\omega})K\|_* d\omega.$$

Hence, compared to the performance potentially achievable with general pre-filters, using diagonal pre-filters degrades the root-mean-square error (RMSE) by a multiplicative factor of at most 2.2. Note also from (14), (16) that the dependency of the ZFE mechanism performance with respect to the parameters ϵ, δ is captured simply by the factor $\kappa_{\delta, \epsilon}$.

VI. EXPLOITING SECOND ORDER STATISTICS ON THE INPUT SIGNALS

One issue with zero-forcing equalizers is the noise amplification at frequencies where $|G(e^{j\omega})|$ is small, due to

the inversion in $H = FG^{-1}$ [41]. This issue is not as problematic for the optimal ZFE mechanism, since in this case the amplification is compensated by the fact that $|F(e^{j\omega})|$ and $|G(e^{j\omega})|$ given in Theorem 7 are both small at the same frequencies. Nonetheless, in general one can improve on the ZFE mechanism by using more advanced equalization schemes in the design of the post-filter H of Fig. 1, and we discuss one such scheme in this section. Note however that the mechanisms presented below require that the input signals satisfy certain properties, such as wide-sense stationarity, and that some publicly available information on these signals be available, e.g., their second order statistics. Hence, the ZFE mechanism remains generally useful due to its broad applicability, even in the absence of any model of the input signal dynamics.

A. LMMSE Mechanisms

One can improve on the ZFE mechanism if some information on the statistics of the input signal u is publicly available. In general, constructing the optimum maximum-likelihood estimate of $\{(Fu)_k\}_{k \geq 0}$ from $\{v_k\}_{k \geq 0}$ on Fig. 1 is computationally intensive and requires the knowledge of the full joint probability distribution of $\{u_k\}_{k \geq 0}$ [41]. Hence, we explore in this subsection a family of simpler schemes based on Linear Minimum Mean Square Error (LMMSE) estimation, which we call LMMSE mechanisms.

Assume that it is publicly known that u is wide-sense stationary (WSS) with known mean vector μ and matrix-valued autocorrelation sequence $R_u[k] = \mathbb{E}[u_t u_{t-k}^T] = R_u[-k]^T, \forall k$. Without loss of generality, we can assume μ to be zero, since the output y is also WSS with known mean equal to $F(1)\mu$. The z -spectrum matrix of u is denoted $P_u(z) = \sum_{k=-\infty}^{\infty} R[k]z^{-k}$. For simplicity of exposition, $P_u(z)$ is assumed to have rational entries and be positive definite on the unit circle, i.e., $P_u(e^{j\omega}) \succ 0$, for all $\omega \in [-\pi, \pi)$. More generally, given two vector-valued WSS signals u and v , we denote the cross-correlation matrix $R_{uv}[k] = \mathbb{E}[u_t v_{t-k}^T]$, the cross power spectral density matrix $P_{uv}(e^{j\omega}) = \sum_{k=-\infty}^{\infty} R_{uv}[k]e^{-j\omega k}$ and the z -cross-spectrum $P_{uv}(z) = \sum_{k=-\infty}^{\infty} R_{uv}[k]z^{-k}$. All z -spectra are assumed to be rational and in particular, we do not discuss the case of spectra with impulses [42].

Referring again to the architecture of Fig. 1, an LMMSE mechanism is characterized by a choice of pre-filter G and the use of a Wiener filter H to estimate y from v [34], [43]. Recall that the Wiener filter produces an estimate \hat{y} minimizing the MSE between y and \hat{y} over linear filters, assuming that the signals y, v are WSS with known second-order joint statistics. Here, these statistics can be expressed in terms of those of u, w , and of the transfer function G . The main contribution of this section is to derive a numerically computable lower bound on the MSE achievable by LMMSE mechanisms for which the pre-filter G is diagonal. The lower bound provides an estimate of how far a specific choice of such diagonal G is from the optimum choice. Our procedure involves the following steps. For the derivation of the lower bound, we assume for tractability reasons that H is an infinite impulse response (IIR) Wiener smoother, i.e., not necessarily

causal. The reason is that we can then express the estimation performance analytically as a function of G . We then optimize this performance measure over diagonal pre-filters G .

For specific implementations of causal LMMSE mechanisms, two natural choices are discussed. First, we can take the pre-filter G obtained from the computation of the lower bound, and design for H a *causal* Wiener filter, or perhaps a slightly non-causal filter if it can be implemented by introducing a delay that is tolerable for a specific application. This mechanism will not attain the lower bound in general, since the bound and corresponding G were obtained by removing the causality constraint on H . Another possibility is to take G as in the ZFE mechanism, and simply replace $H = FG^{-1}$ by a Wiener filter. This choice has the advantage of always improving on the ZFE mechanism, and tends to perform well in practice.

1) *Lower Bound on MSE*: The (non-causal) Wiener smoother H has the transfer function $H(z) = P_{yv}(z)P_v(z)^{-1}$ [34, Section 7.8]. According to Theorem 3, for G diagonal we can take the privacy-preserving noise w to be white and Gaussian with covariance $\sigma^2 I_m$ with $\sigma^2 = \kappa_{\delta,\epsilon}^2 \|GK\|_2^2$. Since u and w are uncorrelated, we have

$$\begin{aligned} P_{yv}(z) &= F(z)P_u(z)G(z^{-1})^T, \\ P_v(z) &= G(z)P_u(z)G(z^{-1})^T + \sigma^2 I_m. \end{aligned} \quad (18)$$

Hence

$$\begin{aligned} H(z) &= F(z)P_u(z)G(z^{-1})^T \\ &\times (G(z)P_u(z)G(z^{-1})^T + \kappa_{\delta,\epsilon}^2 \|G\|_2^2 I_m)^{-1}. \end{aligned} \quad (19)$$

The MSE can then be expressed as $e_{mse}^{LMMSE}(G) = \frac{1}{2\pi} \int_{-\pi}^{\pi} \text{Tr}(P_y(e^{j\omega}) - P_{\hat{y}}(e^{j\omega}))d\omega$ [34, Chapter 7]. In our case,

$$\begin{aligned} P_{\hat{y}}(e^{j\omega}) &= H(e^{j\omega})P_v(e^{j\omega})H(e^{j\omega})^* \\ &= P_{yv}(e^{j\omega})P_v(e^{j\omega})^{-1}P_{yv}(e^{j\omega})^* \\ P_{\hat{y}} &= FP_uG^*(\sigma^2 I_m + GP_uG^*)^{-1}GP_uF^*, \end{aligned}$$

where on the last line and in the following we often omit the argument $e^{j\omega}$ next to matrices, to simplify the notation. Let us denote $P_u(e^{j\omega}) = \Delta(e^{j\omega})\Delta(e^{j\omega})$, where $\Delta(e^{j\omega}) \succeq 0$ is the principal square root of $P_u(e^{j\omega})$. We have then

$$\begin{aligned} P_y - P_{\hat{y}} &= FP_uF^* - P_{\hat{y}} \\ &= F\Delta(I_m - \Delta G^*(\sigma^2 I_m + G\Delta\Delta G^*)^{-1}G\Delta)\Delta F^* \\ &= F\Delta \left(I_m + \frac{1}{\sigma^2} \Delta G^* G \Delta \right)^{-1} \Delta F^*, \end{aligned}$$

where the last expression was obtained using the matrix inversion lemma. Finally, defining $\tilde{G}(e^{j\omega}) := \frac{1}{\|GK\|_2} G(e^{j\omega})K$, we obtain the expression

$$e_{mse}^{LMMSE}(\tilde{G}) = \frac{1}{2\pi} \int_{-\pi}^{\pi} \text{Tr} [F\Delta Z^{-1}\Delta F^*] d\omega \quad (20)$$

$$\text{with } Z = \left(I_m + \frac{1}{\kappa_{\delta,\epsilon}^2} \Delta K^{-1} \tilde{G}^* \tilde{G} K^{-1} \Delta \right),$$

where again the arguments $e^{j\omega}$ were omitted. The objective (20) should be minimized over all transfer functions \tilde{G} that

by definition must satisfy the constraint

$$\|\tilde{G}\|_2^2 = \frac{1}{2\pi} \int_{-\pi}^{\pi} \text{Tr}(\tilde{G}(e^{j\omega})^* \tilde{G}(e^{j\omega}))d\omega = 1. \quad (21)$$

Note that in (20) we can show that we recover the expression (15) of the performance of the ZFE mechanism in the limit $P_u(e^{j\omega}) \rightarrow \infty$.

To obtain a lower bound on performance (for diagonal pre-filters), we now minimize the performance measure (20) over the choice of diagonal pre-filters G satisfying (21). First, in the case where $P_u(e^{j\omega})$ is positive definite and diagonal for all ω , i.e., the different input signals are uncorrelated, we have in fact an allocation problem whose solution is of the ‘‘waterfilling type’’ [44]. Namely, denote $P_u(e^{j\omega}) = \text{diag}(p_1(e^{j\omega}), \dots, p_m(e^{j\omega}))$ and $X(e^{j\omega}) = \tilde{G}(e^{j\omega})^* \tilde{G}(e^{j\omega}) = \text{diag}(x_1(e^{j\omega}), \dots, x_m(e^{j\omega}))$, with $x_i(e^{j\omega}) = |\tilde{G}_{ii}(e^{j\omega})|^2$. Omitting the expression $e^{j\omega}$ in the integrals for clarity, (20) and (21) read

$$\begin{aligned} \min_{x(\cdot)} \quad & \frac{1}{2\pi} \int_{-\pi}^{\pi} \sum_{i=1}^m \frac{1}{\frac{1}{p_i} + \frac{x_i}{\kappa_{\delta,\epsilon}^2 k_i^2}} |F_i|_2^2 d\omega \\ \text{s.t.} \quad & \frac{1}{2\pi} \int_{-\pi}^{\pi} \sum_{i=1}^m x_i d\omega = 1, \quad x_i(e^{j\omega}) \geq 0, \forall \omega, i, \end{aligned}$$

and the solution to this convex problem is

$$x_i(e^{j\omega}) = \max \left\{ 0, \frac{\kappa_{\delta,\epsilon} k_i |F_i(e^{j\omega})|_2}{\lambda} - \frac{\kappa_{\delta,\epsilon}^2 k_i^2}{p_i(e^{j\omega})} \right\}, \quad (22)$$

where $\lambda > 0$ is adjusted so that the solution satisfies the equality constraint (21) (note that since (22) is not smooth we would have to approximate this solution if we wanted to implement the corresponding G with a finite-dimensional filter). Problems of this type are discussed in the communication literature on joint transmitter-receiver optimization [27], [28], which is not too surprising in view of our approximation setup on Fig. 1. When $P_u(e^{j\omega})$ is not diagonal, it is shown in [28] that the problem can be reduced to the diagonal case if G can be arbitrary, however the argument does not appear to carry through under our constraint that G must also be diagonal. Nonetheless, one can obtain a solution arbitrarily close to the minimum one using semidefinite programming. First, we discretize the optimization problem at the set of frequencies $\omega_q = \frac{q\pi}{N}$, $q = 0, \dots, N$. Note that all functions are even functions of ω , hence we can restrict our attention to the interval $[0, \pi]$. Then, we define the $m(N+1)$ variables $x_{iq} = x_i(e^{j\omega_q})$, with $x_{iq} \geq 0$, and $X_q = \text{diag}(x_{1q}, \dots, x_{mq})$. Using the trapezoidal rule to approximate the integrals, we obtain the following optimization problem

$$\min_{\{X_q, M_q\}_{0 \leq q \leq N}} \frac{1}{2N} \sum_{q=0}^{N-1} \text{Tr} [M_q + M_{q+1}] \quad (23)$$

$$\text{s.t.} \quad \begin{bmatrix} M_q & F_q \Delta_q \\ \Delta_q F_q^* & I_m + \Delta_q X_q \tilde{\Delta}_q^* \end{bmatrix} \succeq 0, \quad 0 \leq q \leq N, \quad (24)$$

$$\frac{1}{2N} \sum_{q=0}^{N-1} \text{Tr} [X_q + X_{q+1}] = 1, \quad \text{and } X_q \succeq 0, \quad 0 \leq q \leq N,$$

where $F_q := F(e^{j\omega_q})$, $\Delta_q := \Delta(e^{j\omega_q})$ and $\tilde{\Delta}_q = \frac{1}{\kappa_{\delta,\epsilon}} \Delta_q K^{-1}$. Note that (24) is equivalent to $M_q \succeq F_q \Delta_q \left(I_m + \tilde{\Delta}_q X_q \tilde{\Delta}_q^* \right)^{-1} \Delta_q F_q^*$ by taking the Schur complement. The optimization problem (23), (24) is a semidefinite program (SDP), and can thus be solved efficiently even for a relatively fine discretization of the interval $[0, \pi]$. Assuming such a sufficiently fine discretization, the value obtained from solving the SDP provides a lower bound on the performance achievable by any LMMSE mechanism using a diagonal pre-filter G and a deconvolution Wiener filter H (or even smoother).

2) *Pre- and Post-filter Design*: We single out here two pre-filters G of interest to construct practical LMMSE mechanisms. The first is simply the square root pre-filter of Theorem 7. By using it we improve on the the ZFE mechanism, assuming the statistical assumptions on u are satisfied. Another possible pre-filter G is the one obtained from solving the SDP above. The transfer functions \tilde{G}_{ii} (and hence G_{ii}) of the filter \tilde{G} can then be obtained by interpolation of their squared magnitude $x_i(e^{j\omega})$ from the variables x_{iq} and m spectral factorizations. Note that the resulting filter G would only be optimal in the LMMSE family if we were using a Wiener smoother H , i.e., if we relaxed the causality constraint on H .

Assuming that $G(z)$ is chosen to be rational and with our standing assumption that P_u is rational, then P_v is rational as well, and note from (18) that $P_v(e^{j\omega}) \succ 0$ for all $-\pi \leq \omega < \pi$. Hence $P_v(z)$ has the canonical spectral factorization $P_v(z) = L(z)P_e L(z^{-1})^T$ where $P_e \succ 0$ and L and L^{-1} are analytic in the region $|z| \geq 1$ and $L(\infty) = I_m$ [34, Section 7.8]. Then, the causal Wiener filter is $H(z) = [P_{yv}(z)L(z^{-1})^{-T}]_+ P_e^{-1} L(z)^{-1}$, where for a linear filter $M(z)$ with (matrix-valued) impulse response $\{m_t\}_{-\infty \leq t \leq \infty}$, $[M(z)]_+$ denotes the causal filter with impulse response $\{m_t \mathbf{1}_{\{t \geq 0\}}\}_t$.

B. Example: Events Generated by Markovian Dynamics

We illustrate through a simple example that performance improvement over the ZFE mechanism is possible when we take into account additional public information about the input signals, as described in the previous subsection. Consider a server (shared computing resource, clerk in an administrative office, etc.) processing jobs submitted by customers, and which can be either idle or busy. An event is recorded when the server switches from one state to the other. One way of representing this system is via the Markov chain model with 4 states shown on Fig. 5, with an event being generated when the server enters the intermediary states s_1 and s_2 , in which it can only stay for one time period. The events recorded at the transition times should be kept private. For example, in a shared computing environment, processing times and patterns could provide some information about the jobs submitted by a user that he or she does not want to disclose [45]. On the other hand, we assume that the parameters α, β governing the transition probabilities and shown on Fig. 5 are public information representing known aggregate statistics about the server dynamics.

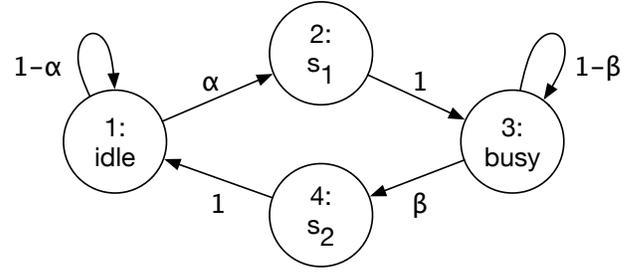


Fig. 5. Transition probabilities for the Markov chain example. The numbering of states 1 to 4 is the one used to express the z -spectrum (25). We assume $\alpha, \beta \notin \{0, 1\}$, in which case the chain is ergodic. The stationary probabilities are then $p_1 = \beta/q$, $p_2 = p_4 = \alpha\beta/q$, $p_3 = \alpha/q$, with $q = \alpha + \beta + 2\alpha\beta$.

For a general time-homogeneous finite-state Markov chain $\{x_t\}_t$ with N states, we can assume without loss of generality that the state space is the set of N -dimensional basis vectors denoted $\{e_1, \dots, e_N\}$, with $e_{ij} = \delta_{ij}$. Denote the N -by- N transition probability matrix Π , with elements $\Pi_{ij} = \mathbb{P}(x_{t+1} = e_i | x_t = e_j)$. Assume moreover that the chain is ergodic (irreducible and aperiodic) and stationary, and introduce the column vector p of stationary probabilities, i.e., $p_i = \mathbb{P}(x_t = e_i)$, so that $\Pi p = p$. Note that with this encoding of the states we have $\mathbb{E}[x_t] = p$. For convenience we also introduce the 0-mean sequence $\tilde{x}_t = x_t - p$, taking values in $\{\tilde{e}_1, \dots, \tilde{e}_N\}$ with $\tilde{e}_i = e_i - p$. Finally, let $D = \text{diag}(p_1, \dots, p_N)$. Then one can show [46] that the z -spectrum of the sequence \tilde{x}_t is

$$P_{\tilde{x}}(z) = (zI - H)^{-1}(D - \Pi D \Pi^T)(z^{-1}I - H^T)^{-1}, \quad (25)$$

where $H = \Pi - p1^T$. We use this formalism to represent the server dynamics above via the Markov chain of Fig. 5, assumed to be stationary. Hence $N = 4$, and the two stationary signals generated by the server entering the intermediate states s_1 and s_2 encoded here as e_2 and e_4 are $u_{1,t} = e_2^T x_t = e_2^T \tilde{x}_t + p_2$, $u_{2,t} = e_4^T x_t = e_4^T \tilde{x}_t + p_4$. Denoting $C = [e_2 \ e_4]^T$ and $\tilde{u} = C\tilde{x}$ the centered input signals, we have $P_{\tilde{u}}(z) = CP_{\tilde{x}}(z)C^T$.

Suppose now for example that we want to release a filtered version of the two-dimensional input signal u , for the MISO filter $F(z) = [F_1(z) \ F_2(z)]$, where F_1 and F_2 are low-pass FIR filters with triangular impulse responses of different lengths

$$F_i(z) = \frac{2}{N_i} \left(\sum_{k=0}^{N_i-1} (k+1)z^{-k} + \sum_{k=N_i}^{2N_i-1} (2N_i - k)z^{-k} \right), \quad (26)$$

with $N_1 = 50$, $N_2 = 25$. We can take $k_1 = k_2 = 1$ in the adjacency relation since the signals $u_{i,t}$ switch between values separated by 1. Fig. 6 shows an example of sample paths obtained using the ZFE mechanism and the LMMSE mechanism replacing the inverse post-filter of the ZFE mechanism by a Wiener filter. In this case choosing the pre-filter obtained from solving the optimization problem (23) (optimal for the Wiener smoother) gives a worse error. Then, for the same two mechanisms providing (ϵ, δ) -differential privacy, Fig. 7 shows the RMSE obtained for $\delta = 0.05$ and different

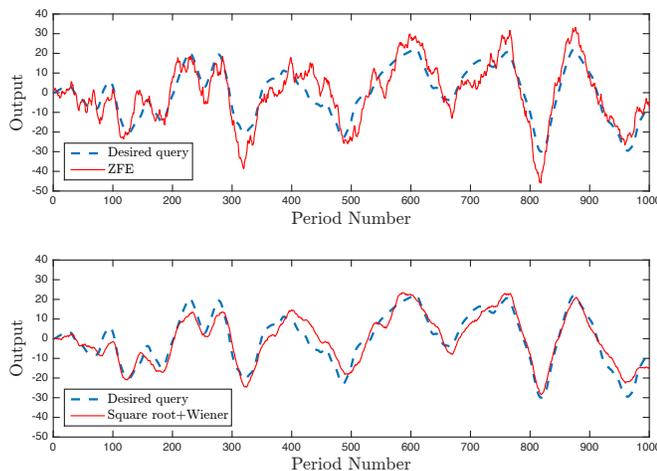


Fig. 6. Sample paths obtained for the query (26) and two $(\ln(3), 0.05)$ differentially private approximations: the ZFE mechanism (top, empirical RMSE = 7.1), and the same mechanism with the inverse post-filter replaced by a (causal) Wiener filter (bottom, empirical RMSE = 4.5). The input signals u_1 and u_2 are generated by the Markov chain of Fig. 5 as explained in the text, with parameters $\alpha = 0.05$, $\beta = 0.15$.

values of ϵ . We also plot the performance lower bounds for ZFE mechanisms with general pre-filters (Theorem 8) and for LMMSE mechanisms for any diagonal pre-filter (solution of (23)). We can see that ZFE mechanisms become quickly unusable as ϵ become small, i.e., when a significant amount of privacy preserving noise is introduced, even if we could optimize over general pre-filters rather than just diagonal ones. Recall in particular from (16), (14) that the dependance on the privacy parameters ϵ, δ of both the diagonal ZFE mechanism performance and of the lower bound is captured completely by the multiplicative factor $\kappa_{\delta, \epsilon}$. This is not the case for the LMMSE mechanism however, since the Wiener filter takes the privacy noise variance into account. Moreover, for the LMMSE mechanism, we see that there is in this case little performance to gain by further optimizing the diagonal pre-filter, or even by replacing the post-filter by a non-causal smoother.

VII. CONCLUSION

In this paper we have developed several approximations of MIMO filters that enforce differential privacy guarantees, while attempting to minimize the impact on performance of this privacy specification. An optimal ZFE mechanism extending the mechanism in [14], [15] was obtained for the approximation of SIMO filters, and a suboptimal one considering only diagonal pre-filters was obtained for general MIMO filters, together with a bound on the performance achievable with non-diagonal pre-filters. We also illustrated the significant performance gain that can be expected by leveraging a model of the input signals, in this case on the second order statistics, especially for high privacy levels (i.e., high noise levels, small δ, ϵ).

The architecture of Fig. 1 considered here for the privacy mechanisms appears to be in fact quite generic. It decomposes the problem into a standard estimation problem, for which

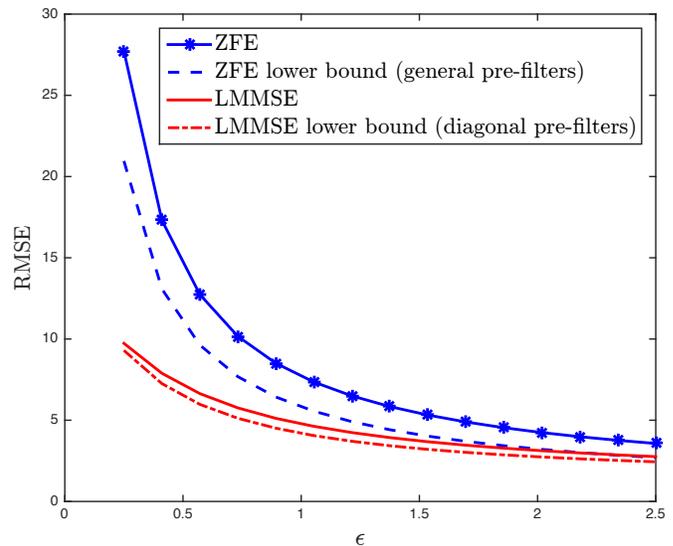


Fig. 7. Performance comparison (RMSE) for the $(\epsilon, 0.05)$ -differentially private ZFE mechanism and LMMSE mechanism (using the same pre-filter as ZFE but a causal Wiener post-filter), as a function of ϵ . We also show the lower bounds on performance from Theorem 8 and from (23), which give an indication of the maximum performance improvement that one could get by further optimizing the mechanisms.

many alternative techniques could be used depending on the scenario considered, and a first stage privacy-preserving filter optimization problem. Natural extensions of this work include considering higher performance equalizers such as decision-feedback equalizers (see [1] for the SISO case) and designing mechanisms that work more naturally with relaxations of the adjacency relation allowing users to activate the same sensor multiple times.

REFERENCES

- [1] J. Le Ny, "On differentially private filtering for event streams," in *Proceedings of the 52nd Conference on Decision and Control*, Florence, Italy, December 2013.
- [2] J. Le Ny and M. Mohammady, "Differentially private MIMO filtering for event streams and spatio-temporal monitoring," in *Proceedings of the 53rd Conference on Decision and Control*, Los Angeles, CA, December 2014.
- [3] R. H. Weber, "Internet of things - new security and privacy challenges," *Computer Law and Security Review*, vol. 26, pp. 23–30, 2010.
- [4] President's Council of Advisors on Science and Technology, "Big data and privacy: A technological perspective," Report to the President, Executive Office of the President of the United States, Tech. Rep., May 2014.
- [5] Electronic Privacy Information Center (epic). Online: <http://epic.org/>.
- [6] H. Chan and A. Perrig, "Security and privacy in sensor networks," *Computer*, vol. 36, no. 10, pp. 103–105, Oct 2003.
- [7] G. Duncan and D. Lambert, "Disclosure-limited data dissemination," *Journal of the American Statistical Association*, vol. 81, no. 393, pp. 10–28, March 1986.
- [8] L. Sweeney, "k-anonymity: A model for protecting privacy," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, no. 05, pp. 557–570, 2002.
- [9] L. Sankar, S. R. Rajagopalan, and H. V. Poor, "Utility-privacy trade-offs in databases: An information-theoretic approach," *IEEE Trans. on Information Forensics and Security*, vol. 8, no. 6, pp. 838–852, June 2013.
- [10] M. Xue, W. Wang, and S. Roy, "Security concepts for the dynamics of autonomous vehicle networks," *Automatica*, vol. 50, pp. 852–857, 2014.
- [11] N. E. Manitaras and C. N. Hadjicostis, "Privacy-preserving asymptotic average consensus," in *Proceedings of the European Control Conference*, Zurich, Switzerland, 2013.

- [12] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Proceedings of the Third Theory of Cryptography Conference*, New York, NY, March 2006, pp. 265–284.
- [13] C. Dwork, "Differential privacy," in *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP)*, ser. Lecture Notes in Computer Science, vol. 4052, Venice, Italy, July 2006.
- [14] J. Le Ny and G. J. Pappas, "Differentially private filtering," in *Proceedings of the 51st Conference on Decision and Control*, Maui, HI, December 2012.
- [15] —, "Differentially private filtering," *IEEE Transactions on Automatic Control*, vol. 59, no. 2, pp. 341–354, February 2014.
- [16] A. Narayanan and V. Shmatikov, "Robust de-anonymization of large sparse datasets (how to break anonymity of the Netflix Prize dataset)," in *Proceedings of the IEEE Symposium on Security and Privacy*, Oakland, CA, May 2008.
- [17] J. A. Calandrino, A. Kilzer, A. Narayanan, E. W. Felten, and V. Shmatikov, "'You might also like': Privacy risks of collaborative filtering," in *Proceedings of the IEEE Symposium on Security and Privacy*, Berkeley, CA, May 2011.
- [18] V. Rastogi and S. Nath, "Differentially private aggregation of distributed time-series with transformation and encryption," in *Proceedings of the ACM Conference on Management of Data (SIGMOD)*, Indianapolis, IN, June 2010.
- [19] Y. D. Li, Z. Zhang, M. Winslett, and Y. Yang, "Compressive mechanism: Utilizing sparse representation in differential privacy," in *Proceedings of the 10th annual ACM workshop on Privacy in the electronic society*, Chicago, IL, October 2011.
- [20] C. Li and G. Miklau, "An adaptive mechanism for accurate query answering under differential privacy," in *Proceedings of the Conference on Very Large Databases (VLDB)*, Istanbul, Turkey, 2012.
- [21] C. Dwork, M. Naor, T. Pitassi, and G. N. Rothblum, "Differential privacy under continual observations," in *Proceedings of the ACM Symposium on the Theory of Computing (STOC)*, Cambridge, MA, June 2010.
- [22] T.-H. H. Chan, E. Shi, and D. Song, "Private and continual release of statistics," *ACM Transactions on Information and System Security*, vol. 14, no. 3, pp. 26:1–26:24, November 2011.
- [23] J. Bolot, N. Fawaz, S. Muthukrishnan, A. Nikolov, and N. Taft, "Private decayed predicate sums on streams," in *Proceedings of the 16th International Conference on Database Theory*, Genoa, Italy, 2013, pp. 284–295.
- [24] J. Le Ny, A. Touati, and G. J. Pappas, "Real-time privacy-preserving model-based estimation of traffic flows," in *Proceedings of the Fifth International Conference on Cyber-Physical Systems (ICCPs)*, Berlin, Germany, April 2014.
- [25] J. Cao, Q. Xiao, G. Ghinita, N. Li, E. Bertino, and K.-L. Tan, "Efficient and accurate strategies for differentially-private sliding window queries," in *Proceedings of the International Conference on Extending Database Technology*, Genoa, Italy, 2013, pp. 191–202.
- [26] L. Fan, L. Xiong, and V. Sunderam, "Differentially private multi-dimensional time series release for traffic monitoring," in *27th Conference on Data and Applications Security and Privacy*, ser. Lecture Notes in Computer Science, vol. 7964. Springer, 2013, pp. pp 33–48.
- [27] J. Salz, "Digital transmission over cross-coupled linear channels," *AT&T Technical Journal*, vol. 64, no. 6, pp. 1147–1159, July-August 1985.
- [28] J. Yang and S. Roy, "On joint transmitter and receiver optimization for multiple-input-multiple-output (MIMO) transmission systems," *IEEE Journal on Communications*, vol. 42, no. 12, pp. 3221–3231, December 1994.
- [29] D. H. Wilson and C. Atkeson, "Simultaneous tracking and activity recognition (STAR) using many anonymous, binary sensors," in *Pervasive Computing*, ser. Lecture Notes in Computer Science, H.-W. Gellersen, R. Want, and A. Schmidt, Eds. Springer Berlin Heidelberg, 2005, vol. 3468, pp. 62–79.
- [30] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, "Our data, ourselves: Privacy via distributed noise generation," in *Proceedings of the 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, St. Petersburg, Russia, May 2006, pp. 486–503.
- [31] 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, 2014.
- [32] L. Wasserman and S. Zhou, "A statistical framework for differential privacy," *Journal of the American Statistical Association*, vol. 105, no. 489, pp. 375–389, March 2010.
- [33] P. Kairouz, S. Oh, and P. Viswanath, "The composition theorem for differential privacy," *IEEE Transactions on Information Theory*, 2017, in Press.
- [34] T. Kailath, A. H. Sayed, and B. Hassibi, *Linear Estimation*. Prentice Hall, 2000.
- [35] C. Ding, D. Zhou, X. He, and H. Zha, "R1-PCA: rotational invariant L_1 -norm principal component analysis for robust subspace factorization," in *Proceedings of the 23rd International Conference on Machine Learning (ICML)*, Pittsburgh, PA, June 2006, pp. 281–288.
- [36] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 2nd ed. Cambridge University Press, 2012.
- [37] Y. A. Ivanov, C. R. Wren, A. Sorokin, and I. Kaur, "Visualizing the history of living spaces," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1153–1160, November-December 2007.
- [38] C. Wren, Y. Ivanov, D. Leigh, and J. Westhues, "The MERL motion detector dataset: 2007 workshop on massive datasets," Mitsubishi Electric Research Laboratories, Tech. Rep. TR2007-069, November 2007.
- [39] L. Ljung, *System Identification: Theory for the User*, ser. Information and System Sciences. Prentice Hall, 1998.
- [40] P. Stoica and R. L. Moses, *Spectral Analysis of Signals*. Prentice Hall, 2005.
- [41] J. Proakis, *Digital Communications*, 4th ed. McGraw-Hill, 2000.
- [42] A. Papoulis, "Predictable processes and Wold's decomposition: A review," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 33, no. 4, pp. 933–938, August 1985.
- [43] H. V. Poor, *An Introduction to Signal Detection and Estimation*, 2nd ed. Springer, 1994.
- [44] D. P. Palomar and J. R. Fonollosa, "Practical algorithms for a family of waterfilling solutions," *IEEE Transactions on Signal Processing*, vol. 53, no. 2, pp. 686–695, February 2005.
- [45] S. Kadloor, X. Gong, N. Kiyavash, and P. Venkatasubramanian, "Designing router scheduling policies: A privacy perspective," *IEEE Transactions on Signal Processing*, vol. 60, no. 4, pp. 2001–2012, April 2012.
- [46] G. Bilardi, R. Padovani, and G. L. Pierobon, "Spectral analysis of functions of Markov chains with applications," *IEEE Transactions on Communications*, vol. 31, no. 7, pp. 853–861, July 1983.



Jerome Le Ny (S'05-M'09-SM'16) received the Engineering Degree from the École Polytechnique, France, in 2001, the M.Sc. degree in Electrical Engineering from the University of Michigan, Ann Arbor, in 2003, and the Ph.D. degree in Aeronautics and Astronautics from the Massachusetts Institute of Technology, Cambridge, in 2008. He is currently an Associate Professor with the Department of Electrical Engineering, Polytechnique Montreal, Canada. He is a member of GERAD, a multi-university research center on decision analysis. From 2008

to 2012 he was a Postdoctoral Researcher with the GRASP Laboratory at the University of Pennsylvania. His research interests include robust and stochastic control, scheduling and dynamic resource allocation problems, mean-field control, and enforcing privacy in sensor and actuator networks, with applications to networked control systems, autonomous and multi-robot systems, and transportation systems.



Meisam Mohammady received the B.S. degree from Sharif University, Iran, in 2012, and the M.S. degree in Electrical Engineering from Polytechnique Montreal, Canada, in June 2015. He is currently a Ph.D. student at Concordia University, Montreal. His research interests are in privacy-preserving information processing.