A Model Randomization Approach to Statistical Parameter Privacy

Ehsan Nekouei, *Member, IEEE*, Henrik Sandberg, *Senior Member, IEEE*,
Mikael Skoglund, *Fellow, IEEE*, and Karl Henrik Johansson, *Fellow, IEEE*

***Abstract*—In this article, we study a privacy filter design problem for a sequence of sensor measurements whose joint probability density function (p.d.f.) depends on a private parameter. To ensure parameter privacy, we propose a filter design framework which consists of two components: a randomizer and a nonlinear transformation. The randomizer takes the private parameter as input and randomly generates a pseudo parameter. The nonlinear mapping transforms the measurements such that the joint p.d.f. of the filter's output depends on the pseudo parameter rather than the private parameter. It also ensures that the joint p.d.f. of the filter's output belongs to the same family of distributions as that of the measurements. The design of the randomizer is formulated as an optimization problem subject to a privacy constraint, in terms of mutual information, and it is shown that the optimal randomizer is the solution of a convex optimization problem. Using information-theoretic inequalities, we show that the performance of any estimator of the private parameter, based on the output of the privacy filter, is limited by the privacy constraint. The structure of the nonlinear transformation is studied in the special cases of independent and identically distributed, Markovian, and Gauss-Markov measurements. Our results show that the privacy filter in the Gauss-Markov case can be implemented as two one-step ahead Kalman predictors and a set of minimum mean square error predictors. A numerical example on occupancy privacy in a building automation system illustrates the approach.**

***Index Terms*—Information theory, Kalman filtering, privacy in networked control systems.**

I. Introduction

A. Motivation

NETWORKED control systems are omnipresent in our daily lives by providing essential services such as intelligent transportation, smart grid, and intelligent buildings. Sensors are crucial components of any system as they provide critical information that can be used for control, diagnosis, and monitoring purposes. However, sensor measurements in a networked system typically contain information about private variables. Thus, directly revealing the measurements to untrusted parties may expose the system to the risk of privacy loss. For example, the occupancy level of a building, which is a highly private variable, can be inferred from the CO_2 and temperature measurements of the building [1]. Privacy breaches may have negative consequences such as reputation damage and financial losses due to lawsuits. The sheer importance of privacy has motivated numerous research efforts to develop privacy-preserving solutions for networked control systems.

B. Related Work

A variety of problems related to privacy and networked estimation and control have been studied in the literature. The privacy aspect of the hypothesis testing problem as well as various solutions for privacy-aware hypothesis testing have been studied in the literature, e.g., [2]–[6]. The authors in [7] considered a multisensor hypothesis testing problem wherein a fusion center receives the decisions of a set of sensors and an adversary overhears the local decisions of a subset of sensors. They studied the optimal privacy-aware hypothesis testing rule that minimizes the Bayes risk subject to a privacy constraint at the adversary. In [8], the authors considered a similar setup and investigated the optimal privacy-aware design of the Neyman–Pearson test.

The authors in [9] studied the optimal privacy filter design problem for a public Markov chain that is correlated with a private Markov chain. Tanaka *et al.* [10] considered a linear Gaussian plant in which the sensor measurements are communicated with an untrusted cloud controller. They studied the optimal privacy filter design problem subject to a constraint on the privacy of the states of the plant. The authors of [11] studied the optimal privacy-aware control law for a Markov decision process subject to a privacy constraint on an adversary which has access to the input and output of the Markov chain. We note that the information-theoretic approach to privacy has

Manuscript received 12 April 2021; revised 19 August 2021; accepted 16 January 2022. Date of publication 25 January 2022; date of current version 30 January 2023. This work was supported in part by the Research Grants Council of Hong Kong under the Project CityU 9048210, in part by the Chow Sang Sang Group Research Fund sponsored by Chow Sang Sang Holdings International Ltd., in part by the Knut and Alice Wallenberg Foundation, in part by the Swedish Foundation for Strategic Research, and in part by the Swedish Research Council. Recommended by Associate Editor G. Gu. *(Corresponding author: Ehsan Nekouei.)*

Ehsan Nekouei is with the Department of Electrical Engineering, City University of Hong Kong, Hong Kong, China (e-mail: enekouei@cityu.edu.hk).

Henrik Sandberg, Mikael Skoglund, and Karl Henrik Johansson are with the School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden, and also with Digital Futures, SE-100 44 Stockholm, Sweden (e-mail: hsan@kth.se; skoglund@kth.se; kallej@kth.se).

Color versions of one or more figures in this article are available at https://doi.org/10.1109/TAC.2022.3145664.

Digital Object Identifier 10.1109/TAC.2022.3145664

been extensively studied in the literature, e.g., [12]–[15]. In such an approach, one observes a public random variable which is correlated with a private variable. Here, the objective is to generate a degraded version of the public variable such that the distortion due to the privacy filter is minimized subject to an information-theoretic constraint on the privacy. The interested reader is referred to [16] for an overview of information-theoretic approaches to privacy in estimation and control.

Privacy-aware solutions to estimation, filtering, and average consensus problems have been developed in the literature based on the notion of differential privacy. The authors in [17] developed a framework for privacy-aware filtering of the measurements of a dynamical system using the concept of differential privacy. Sandberg *et al.* [18] studied the state estimation in an electricity distribution network subject to a constraint on the consumers' privacy. Privacy-aware average consensus algorithms were developed in [19] and [20] to guarantee the privacy of agents' initial states. The authors of [21] developed a framework based on differential privacy to address the privacy of the initial state as well as the way-points of each agent in a distributed control problem.

The statistical parameter privacy problem has been studied in [22] and [23]. Bassi *et al.* [22] studied the statistical parameter privacy of an independent and identically distributed (i.i.d.) sequence of random variables where their common pdf depends on a private parameter. In their setup, the privacy filter consists of a randomly selected stochastic kernel which generates a random output based on each observed random variable. They characterized the leakage level of private information under the Bayes statistical risk as the privacy measure. The authors in [23] studied the optimal design of controller and privacy filter for a linear Gaussian plant. The system dynamics should be kept private from an adversary interested in inferring the system dynamics based on the state measurements and control inputs. Assuming Fisher information as privacy metric, they showed that the optimal privacy filter is in the form of a state-dependent Gaussian stochastic kernel.

C. Contributions

In this article, we consider a sequence of measurements, observed by a sensor, whose joint pdf depends on a private parameter. To ensure parameter privacy of the measurements, we propose a filter design framework which consists of two parts: a randomizer and a nonlinear transformation. The randomizer takes the private parameter as input and randomly generates a pseudoparameter. The nonlinear transformation alters the measurements such that the joint pdf of the filter's output is characterized by the pseudoparameter rather than the private parameter. The nonlinear transformation also ensures that the joint pdf of the filter's output belongs to the same family of distributions as the measurements. The nonlinear transformation has a feedforward–feedback structure which enables real-time and causal computation of the disguised measurements with low complexity.

In our setup, the randomizer is designed by minimizing the average distortion due to the privacy filter subject to a privacy constraint in terms of the mutual information between the private and the pseudoparameters. Using information-theoretic inequalities, we show that the performance of any estimator of the private parameter based on the output of the privacy filter is limited by the privacy constraint. We investigate the structure of the nonlinear transformation for i.i.d., Markovian, and Gauss–Markov measurements. Our results show that the structure of the nonlinear transformation involves two one-step ahead Kalman predictors in the Gauss–Markov case. The Kalman predictors significantly reduce the complexity of generating the disguised measurements. The result for i.i.d. measurements has appeared in [24].

Different from [22] and [24], the current article develops a privacy filter design framework without imposing the i.i.d. assumption on the joint pdf of the measurements. In our setup, the sensor measurements are processed by a nonlinear transformation, rather than a stochastic kernel, which ensures that the distribution of the output of privacy filter belongs to the same family of distributions as the measurements. This requirement is not guaranteed by the frameworks in [22] and [23] since stochastic kernels significantly alter the distribution of the measurements.

Our mutual information privacy metric is fundamentally different from the Fisher information metric in [23], in the sense that the mutual information provides a lower bound on the error probability of any arbitrary estimator of private variables, whereas Fisher information provides a lower bound on the mean square error of any unbiased estimator of private parameters.

D. Outline

This article is organized as follows. Section II describes our system model and standing assumptions. Section III introduces the model randomization approach to parameter privacy. Section IV investigates the structure of the nonlinear transformation in special cases. Section V presents the numerical results followed by the concluding remarks in Section VI.

II. Problem Formulation

A. System Model and Objectives

Consider a sensor that measures the stochastic process Y_k over the time-horizon $k = 1, \dots, T$ where $Y_k = [Y_k^1, \dots, Y_k^d]^\top$ is a d-dimensional random vector. We assume that the joint pdf of the measurements over the time-horizon $1, \dots, T$ is parameterized by Θ which takes values in $\boldsymbol{\Theta} = \{\theta_1, \dots, \theta_m\}$ with probability mass function $\mathsf{Pr}(\Theta = \theta_i) = p_i$. The value of the parameter Θ is fixed during the time-horizon $k = 1, \dots, T$. The joint pdf of $\{Y_k\}_{k=1}^T$ is denoted by $p_{\theta_i}(y_1, \dots, y_T)$ when Θ is equal to θ_i. We refer to $p_\theta(y_1, \dots, y_T)$ as the *statistical model* of the measurements which belongs to the family of pdfs

$$\mathcal{M} = \{p_\theta\left(y_1, \dots, y_T\right)\}_\theta .$$

We assume that Θ carries private information and the sensor sequentially communicates its measurements with an untrusted party, hereafter named the "user," for monitoring, control, or storage purposes. Directly revealing the measurements to the user will result in the loss of private information since the

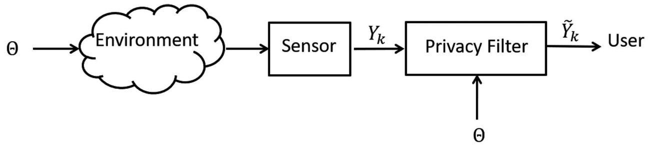

Fig. 1. Information sharing with an untrusted user via a privacy filter.

true value of Θ can be inferred from the measurements. A common approach for ensuring privacy is to use a privacy filter that mediates the information sharing between the sensor and user as shown in Fig. 1. In this article, we develop a privacy filter design framework that achieves the following three objectives.

1) It ensures that the output of the privacy filter accurately represents the sensor measurements.
2) It guarantees that an adversary with access to the output of the privacy filter cannot reliably infer the value of Θ.
3) The joint pdf of the output of the privacy filter belongs to the family of distributions $\mathcal{M}$.

In general, the output distribution of a privacy filter might be arbitrarily complex. Objective 3 avoids this situation by ensuring that the joint pdf of the privacy filter's output belongs to the same family of distributions as that of the measurements. This is especially important when the filter's output is used to perform computations, such as filtering, whose complexity depends on the underlying distribution of the data.

Note that adding noise to the measurements may not ensure the privacy of the parameter Θ. To highlight this point, assume that the Y_ks are i.i.d. according to a Gaussian distribution with mean Θ and unit variance. Let $\tilde{Y}_k = Y_k + N_k$ denote the shared information with the user where $\{N_k\}_k$ is a sequence of zero mean i.i.d. distributed random variables. Using the law of large numbers, we have

$$\frac{1}{T}\sum_{k=1}^{T}\tilde{Y}_k \to \Theta$$

almost surely as T tends to infinity, which indicates that the user can reliably estimate the private parameter when T is large. Hence, the noise addition mechanism does not ensure the privacy of Θ. In this example, the statistical model of the measurements belongs to the family of distributions

$$\mathcal{M} = \left\{\frac{1}{\sqrt{2\pi}}\mathrm{e}^{\frac{1}{2}(x-\theta)^2}\right\}_{\theta\in\mathbb{R}}$$

where θ takes values in the set of real numbers. If the additive noise terms N_ks are Laplacian distributed [25], objective 3 is not satisfied and the pdf of the shared information with the user becomes intractable.

Remark 1: Although we assume that Θ is fixed, our results can be easily extended to the case that Θ is time-varying, e.g., when the private parameter is given by $\Theta = [\Theta_1, \ldots, \Theta_T]$ where Θ_k is the parameter that characterizes the conditional distribution of Y_k, given $Y_1, \ldots, Y_{k-1}$.

B. Motivating Example: Building Automation

Consider a building automation application in which a sensor measures the CO_2 level of a room over the time horizon $k = 1, \ldots, T$. The response of CO_2 concentration to the presence of humans can be modeled as

$$X_{k+1} = aX_k + W_k + b\Theta$$
$$Y_k = X_k + V_k$$

where $a \in (0,1)$ and b are constants, X_k denotes the CO_2 level at time-step k, Y_k represents the sensor measurement, W_k denotes the external disturbance, V_k denotes the measurement noise, and $\Theta \in \{0, \ldots, L\}$ denotes the occupancy level of the room, i.e., $\Theta = i$ indicates that there are i persons in the room during the horizon $k = 1, \ldots, T$.

In building automation, it is common to transmit the sensor measurements over communication networks for control or monitoring purposes. Note that Θ is the private statistical parameter of the shared information $\{Y_k\}_{k=1}^{T}$. When the CO_2 measurements are accessible by untrusted parties, e.g., a hacker or a "cloud," the occupancy level, which carries private information, can be inferred from the measurements.

We examine the privacy of the occupancy information in two cases: 1) when CO_2 measurements are directly shared with the user; 2) when a noise addition mechanism is employed to ensure occupancy privacy. To this end, we first construct an estimator of the occupancy level based on the shared information with the user. Let $\tilde{Y}_k$ denote the shared information at time-step k and define the averages $\bar{Y}_k$ and $\bar{\bar{Y}}_k$ as

$$\bar{Y}_k = \frac{1}{T_h}\sum_{i=k-(T_h-1)}^{k}\tilde{Y}_i$$
$$\bar{\bar{Y}}_k = \frac{1}{T_h}\sum_{i=k-T_h}^{k-1}\tilde{Y}_i \tag{1}$$

for $k \geq T_h + 1$ where T_h is the window size. Let $\hat{\Theta}_k$ denote the estimator of occupancy based on the shared information up to time-step k, $\tilde{Y}_1, \ldots, \tilde{Y}_k$, which is defined as

$$\hat{\Theta}_k = \frac{\bar{Y}_k - \bar{\bar{Y}}_k}{a}. \tag{2}$$

Fig. 2 shows the output of the occupancy estimator, as a function of time, when the true occupancy level is equal to 1 or 2. The solid lines in Fig. 2 show the occupancy estimates when the estimator has access to the CO_2 measurements, i.e., $\tilde{Y}_k = Y_k$. The dashed lines in this figure represent the occupancy estimates under a noise addition privacy mechanism wherein Gaussian noise is added to the measurements. In this case, the estimator has access to $\tilde{Y}_k = Y_k + N_k$ where $\{N_k\}_k$ is a sequence of i.i.d. Gaussian random variables with zero mean and unit variance. In our simulations, $\{W_k, V_k\}_k$ is assumed to be a sequence of i.i.d. Gaussian random variables with zero mean and variance 0.1.

According to Fig. 2, the occupancy estimator in (2) can reliably infer the occupancy level even if the noise addition privacy mechanism is employed. This observation confirms that directly sharing sensor measurements with an entrusted party might result in the loss of private information. Moreover, based

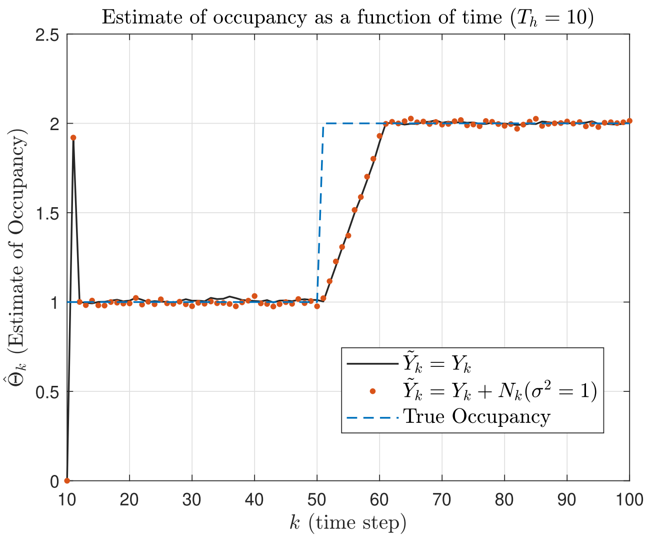

Fig. 2. Estimates of occupancy (Θ) as a function of time for different levels of occupancy.

on Fig. 2, noise addition mechanisms may not be capable of ensuring statistical parameter privacy.

C. Notations and Assumptions

The shorthand notation $Y_{1:k}$ is used to represent the sequence of random variables $Y_1, \ldots, Y_k$. The realization of $Y_{1:k}$ is denoted by $y_{1:k}$. The shorthand notation $Y_k^{1:l}$ denotes the collection of the first l components of Y_k. The lth component of Y_k is denoted by Y_k^l. When Θ is equal to θ, the joint pdf of $Y_{1:k}$ is denoted by $p_\theta(y_1, \ldots, y_k)$ which is assumed to be nonzero almost everywhere in $\mathbb{R}^{k\times d}$.

The conditional pdf of Y_k^l given $\{Y_k^{1:l-1} = y_k^{1:l-1}, Y_{1:k-1} = y_{1:k-1}, \Theta = \theta_i\}$ is represented by

$$p_{\theta_i}\left(x \left| y_k^{1:l-1}, y_{1:k-1}\right.\right)$$

where its corresponding cumulative distribution function (cdf) is assumed to be absolutely continuous with respect to the Lebesgue measure. The conditional cdf of Y_k^l given $\{Y_k^{1:l-1} = y_k^{1:l-1}, Y_{1:k-1} = y_{1:k-1}, \Theta = \theta_i\}$ is defined as

$$F_{l,k,\theta_i}\left(z \left| y_k^{1:l-1}, y_{1:k-1}\right.\right) = \int_{-\infty}^{z} p_{\theta_i}\left(x \left| y_k^{1:l-1}, y_{1:k-1}\right.\right) dx. \quad (3)$$

The inverse function of $F_{l,k,\theta_i}(\cdot|y_k^{1:l-1}, y_{1:k-1})$ is denoted by $F^{-1}_{l,k,\theta_i}(\cdot|y_k^{1:l-1}, y_{1:k-1})$. Note that, for all θ_i, the transformation $F_{l,k,\theta_i}(\cdot|\cdot)$ is a mapping from $\mathbb{R}^{l+(k-1)d}$ to $[0,1]$ which is increasing in its first argument when its second argument is fixed.

The following conventions are adopted in the rest of this article:

$$\begin{aligned}
F_{1,1,\theta_i}\left(z \left| y_k^{1:0}, y_{1:0}\right.\right) &= F_{1,1,\theta_i}(z),\\
F_{l,1,\theta_i}\left(z \left| y_1^{1:l-1}, y_{1:0}\right.\right) &= F_{l,1,\theta_i}\left(z \left| y_1^{1:l-1}\right.\right),\\
F_{1,k,\theta_i}\left(z \left| y_k^{1:0}, y_{1:k-1}\right.\right) &= F_{1,k,\theta_i}\left(z \left| y_{1:k-1}\right.\right)
\end{aligned}$$

where $F_{1,1,\theta_i}(z)$ is the conditional cdf of Y_1^1 given $\Theta = \theta_i$, $F_{l,1,\theta_i}(z|y_1^{1:l-1})$ is the conditional cdf of Y_1^l given $\{Y_1^{1:l-1} = y_1^{1:l-1}, \Theta = \theta_i\}$, and $F_{1,k,\theta_i}(z|y_{1:k-1})$ is the conditional cdf of Y_k^1 given $\{Y_{1:k-1} = y_{1:k-1}, \Theta = \theta_i\}$.

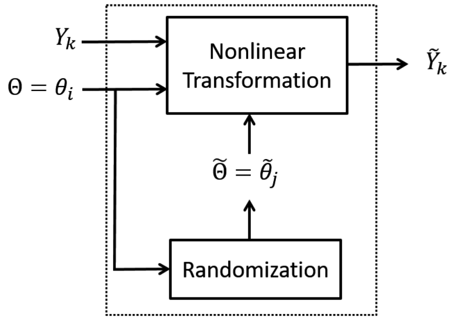

Fig. 3. Structure of the proposed privacy filter.

III. Model Randomization Approach

In this section, we discuss the model randomization framework for ensuring the statistical parameter privacy of the sensor measurements. In this framework, the privacy filter consists of two components, a randomizer and a nonlinear transformation, as shown in Fig. 3, which collectively attain the objectives 1), 2), and 3) in Section II. The randomizer takes the value of the private variable Θ as input and generates a realization of the *pseudoparameter* $\tilde{\Theta}$ which remains constant during the horizon. The pseudoparameter takes values in the set $\tilde{\boldsymbol{\Theta}} = \{\tilde{\theta}_1, \ldots, \tilde{\theta}_{\tilde{m}}\}$. At each time-step k, the nonlinear transformation generates $\tilde{Y}_k$ based on the sensor measurement at time-step k and the values of Θ and $\tilde{\Theta}$. Then, $\tilde{Y}_k$ is revealed to the user.

In this section, the design of the privacy filter is studied without imposing any special structure on the joint pdf of the sensor measurements. We start by discussing the structure of the nonlinear transformation in the next subsection. Then, the optimal design of the randomizer is discussed, followed by the privacy analysis of Θ under the proposed framework.

A. Nonlinear Transformation

Let $\tilde{y}_k = [\tilde{y}_k^1, \ldots, \tilde{y}_k^d]^\top$ denote the realization of the output of the privacy filter at time k. Assuming $\Theta = \theta_i$ and $\tilde{\Theta} = \tilde{\theta}_j$, the nonlinear transformation at time-step k generates $\tilde{y}_k^l$, i.e., the lth entry of $\tilde{y}_k$, according to

$$\tilde{y}_k^l = F^{-1}_{l,k,\tilde{\theta}_j}\left(u_k^l \left| \tilde{y}_k^{1:l-1}, \tilde{y}_{1:k-1}\right.\right) \quad (4)$$

where u_k^l is given by

$$u_k^l = F_{l,k,\theta_i}\left(y_k^l \left| y_k^{1:l-1}, y_{1:k-1}\right.\right) \quad (5)$$

and $F_{l,k,\theta_i}(\cdot|y_k^{1:l-1}, y_{1:k-1})$ is the cdf of Y_k^l given $\{Y_k^{1:l-1} = y_k^{1:l-1}, Y_{1:k-1} = y_{1:k-1}, \Theta = \theta_i\}$ defined in (3). According to (4) and (5), the output of the privacy filter at time-step k is generated based on the history of the measurements up to time-step k. Moreover, the lth entry of $\tilde{y}_k$ is generated using its first $l-1$ entries, i.e., $\tilde{y}_k^{1:l-1}$, which implies that the entries of the filter's output are generated sequentially. The structure of the nonlinear transformation at time-step k is illustrated in Fig. 4.

The *feedforward–feedback* structure of the nonlinear transformation, in Fig. 4, is a unique aspect of the proposed privacy

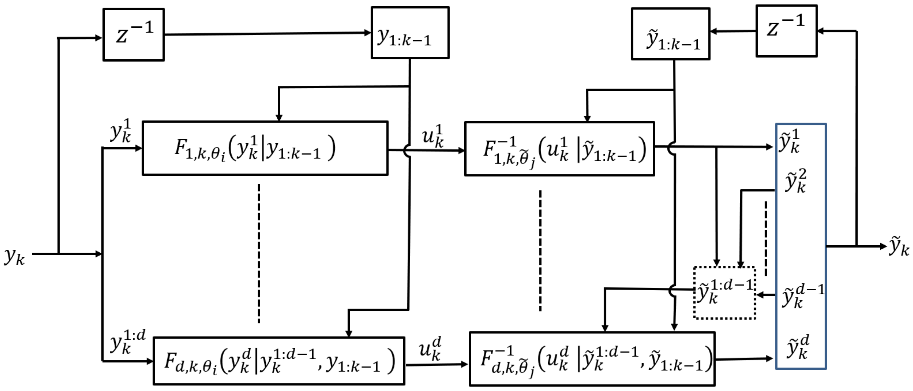
z^{-1}
$y_{1:k-1}$
$\tilde{y}_{1:k-1}$
z^{-1}
y_k^1
$F_{1,k,\theta_i}(y_k^1|y_{1:k-1})$
u_k^1
$F^{-1}_{1,k,\tilde{\theta}_j}(u_k^1|\tilde{y}_{1:k-1})$
$\tilde{y}_k^1$
$\tilde{y}_k^2$
y_k
$\tilde{y}_k$
$\tilde{y}_k^{1:d-1}$
$\tilde{y}_k^{d-1}$
$y_k^{1:d}$
$F_{d,k,\theta_i}(y_k^d|y_k^{1:d-1}, y_{1:k-1})$
u_k^d
$F^{-1}_{d,k,\tilde{\theta}_j}(u_k^d|\tilde{y}_k^{1:d-1}, \tilde{y}_{1:k-1})$
$\tilde{y}_k^d$

Fig. 4. Structure of the nonlinear transformation at time-step k.

filter. The feedforward component computes u_k^is, whereas the feedback component computes $\tilde{y}_k^i$s based on the past outputs of the filter. This structure allows us to cast the proposed privacy filter as a *dynamical system* enabling recursive computation of the pdfs $p_{\theta_i}(y_k|y_{1:k-1})$ and $p_{\theta_j}(\tilde{y}_k|\tilde{y}_{1:k-1})$. Under the proposed scheme, $p_{\theta_i}(y_k|y_{1:k-1})$ and $p_{\theta_j}(\tilde{y}_k|\tilde{y}_{1:k-1})$ can be computed recursively over time to reduce the computational cost of generating the disguised measurements. The recursive structure significantly reduces the computational cost when the measurements are generated by a linear Gaussian system (see Section IV-C for more details).

In certain applications, it is important to generate the disguised measurements in real time, i.e., $\tilde{y}_k$ must be causally generated using the sensor measurements up to time k rather than all the sensor measurements $Y_1, \ldots, Y_T$. The recursive structure of the proposed privacy filter enables the causal (real-time) generation of the disguised measurements.

The next theorem studies the joint pdf of the output of the nonlinear transformation.

Theorem 1: Consider the nonlinear transformation in (4) and (5) and assume that the joint pdf of the measurements belongs to the family of distributions $\mathcal{M} = \{p_\theta(y_1, \ldots, y_T)\}_\theta$. Given $\Theta = \theta_i$ and $\tilde{\Theta} = \tilde{\theta}_j$, the joint pdf of the filter's output is $p_{\tilde{\theta}_j}(\tilde{y}_1, \ldots, \tilde{y}_T)$ for all i, j.

Proof: See Appendix A. ■

According to Theorem 1, the joint pdf of the output of the privacy filter over the horizon $1, \ldots, T$ is characterized by the value of the pseudoparameter $\tilde{\Theta}$ rather than the value of the private parameter Θ. Moreover, the nonlinear transformation ensures that the joint pdf of the filter's output also belongs to the family of probability distributions $\mathcal{M}$. Under the proposed framework, the statistical model of the filter's output, i.e., $p_{\tilde{\theta}}(\tilde{y}_1, \ldots, \tilde{y}_T)$, is randomly chosen from the set of probability distributions $\{p_{\tilde{\theta}}(\tilde{y}_1, \ldots, \tilde{y}_T)\}_{\tilde{\theta} \in \tilde{\boldsymbol{\Theta}}}$ where $\tilde{\boldsymbol{\Theta}} = \{\tilde{\theta}_1, \ldots, \tilde{\theta}_{\tilde{m}}\}$ is the support set of the randomizer's output. Hence, we refer to this framework as the model randomization approach to the parameter privacy.

B. Randomizer

The parameter $\tilde{\Theta}$ is selected from the set $\tilde{\boldsymbol{\Theta}} = \{\tilde{\theta}_1, \ldots, \tilde{\theta}_{\tilde{m}}\}$ using a randomized mapping. More precisely, given $\Theta = \theta_i$, the value of $\tilde{\Theta}$ is randomly generated according to

$$\tilde{\Theta} = \begin{cases} \tilde{\theta}_1 & \text{w.p.} \quad P_{1i}, \\ \vdots & \qquad \vdots \qquad\qquad \text{if} \quad \Theta = \theta_i \\ \tilde{\theta}_{\tilde{m}} & \text{w.p.} \quad P_{\tilde{m}i}, \end{cases}$$

where $\sum_j P_{ji} = 1$ for all i. Thus, the randomizer with probability P_{ji} selects $\tilde{\theta}_j$ as the value of $\tilde{\Theta}$ when Θ is equal to θ_i. The set of randomization probabilities $\{P_{ji}\}_{ji}$ are designed such that the accuracy of the output of the privacy filter is maximized and, simultaneously, a desired privacy level is achieved.

Due to the nonlinear transformation, the output of the privacy filter might be different from its input. To quantify the difference between the input and the output of the filter, we define the average distortion between $Y_{1:T}$ and $\tilde{Y}_{1:T}$ as

$$\frac{1}{T}\sum_{k=1}^{T} \mathsf{E}\left[d\left(Y_k, \tilde{Y}_k\right)\right]$$

where the distortion function $d(\cdot, \cdot) : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}_+$ captures the deviation of the output of the privacy filter from its input.

We consider the mutual information between the private parameter and pseudoparameter as the privacy metric. Let $\mathsf{I}[\Theta; \tilde{\Theta}]$ denote the mutual information between Θ and $\tilde{\Theta}$ which can be written as

$$\mathsf{I}\left[\Theta; \tilde{\Theta}\right]$$
$$= \sum_{i,j} \mathsf{Pr}\left(\Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j\right) \log \frac{\mathsf{Pr}\left(\Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j\right)}{\mathsf{Pr}\left(\Theta = \theta_i\right) \mathsf{Pr}\left(\tilde{\Theta} = \tilde{\theta}_j\right)}.$$

Note that $\mathsf{I}[\Theta; \tilde{\Theta}]$ captures the amount of information that can be inferred about the private parameter by observing the pseudoparameter. When the mutual information between Θ and $\tilde{\Theta}$ is zero, the pseudoparameter has no information about the private parameter. Also, the mutual information between Θ and $\tilde{\Theta}$ achieves its maximum level when Θ can be correctly inferred by observing $\tilde{\Theta}$. Thus, a relatively large level of $\mathsf{I}[\Theta; \tilde{\Theta}]$ indicates that Θ can be reliably inferred from $\tilde{\Theta}$.

The optimal randomization probabilities are obtained by minimizing the average distortion subject to a privacy constraint. More precisely, the optimal randomization probabilities are the solution of the following optimization problem:

$$\begin{aligned} \underset{\{P_{ji}\}_{j,i}}{\text{minimize}} \quad & \frac{1}{T}\sum_{k=1}^{T} \mathsf{E}\left[d\left(Y_k,\tilde{Y}_k\right)\right] \\ & P_{ji} \geq 0, \forall i,j \\ & \sum_j P_{ji} = 1, \quad \forall i \\ & \mathsf{I}\left[\Theta;\tilde{\Theta}\right] \leq I_0 \end{aligned} \tag{6}$$

where the second constraint enforces the law of total probability and the last constraint imposes an upper bound on the mutual information between the private parameter and the pseudoparameter. We refer to the last constraint as the privacy constraint since it limits the amount of information that can be inferred about Θ based on $\tilde{\Theta}$. In what follows, we refer to I_0 as the leakage level of private information.

The next theorem investigates the structural properties of the optimization problem (6).

Theorem 2: The objective function in (6) is linear in the randomization probabilities. Also, the privacy constraint is a convex constraint.

Proof: See Appendix C. ∎

Theorem 2 shows that the optimization problem (6) is a convex optimization problem. Hence, the optimal randomization probabilities can be computed efficiently.

Remark 2: In practice, the elements of $\tilde{\boldsymbol{\Theta}}$ can be selected using a nested optimization problem where the optimal elements of $\tilde{\boldsymbol{\Theta}}$ and the optimal randomization probabilities are computed recursively such that the total distortion is minimized while the privacy constraint is met. Although this results in a nonconvex optimization problem, one can obtain a locally optimal solution by alternating between the nested optimization problems.

Remark 3: It may not be always possible to find a closed-form expression for the total distortion due to the structure of the distortion function and the distribution of the sensor measurements. However, it is straightforward to approximate the total distortion accurately using Monte Carlo simulations.

C. Privacy Level of Θ

In this subsection, we study the privacy level of the parameter Θ under the proposed framework. To this end, let $\hat{\Theta}(\tilde{Y}_{1:T})$ denote an arbitrary estimator of Θ, based on the output of the privacy filter over the horizon $1,\ldots,T$, which is defined as a mapping from $\mathbb{R}^{T\times d}$ to $\boldsymbol{\Theta} = \{\theta_1,\ldots,\theta_m\}$. The next theorem establishes a lower bound on the error probability of the estimator $\hat{\Theta}(\tilde{Y}_{1:T})$.

Theorem 3: Let $\mathsf{Pr}(\Theta \neq \hat{\Theta}(\tilde{Y}_{1:T}))$ denote the error probability of the estimator $\hat{\Theta}(\tilde{Y}_{1:T})$. Then, we have

$$\mathsf{Pr}\left(\Theta \neq \hat{\Theta}\left(\tilde{Y}_{1:T}\right)\right) \geq \frac{\mathsf{H}\left[\Theta\right] - I_0 - 1}{\log|\boldsymbol{\Theta}|}$$

where $\mathsf{H}[\Theta]$ is the discrete entropy of Θ, I_0 is the leakage level of private information in (6), and $|\boldsymbol{\Theta}|$ denotes the cardinality of the set $\boldsymbol{\Theta}$.

Proof: See Appendix D. ∎

According to Theorem 3, the error probability of any estimator of the private parameter based on the output of the privacy filter is limited by the leakage level of private information I_0. The lower bound in Theorem (3) increases as the leakage level of private information becomes small. Thus, for a relatively small value of I_0, no estimator can reliably infer the private parameter Θ which implies that the proposed framework is capable of ensuring the privacy of Θ.

In Appendix D, we show that the mutual information between Θ and $\tilde{Y}_{1:T}$ can be upper bounded as

$$\mathsf{I}\left[\Theta;\tilde{Y}_{1:T}\right] \leq \mathsf{I}\left[\Theta;\tilde{\Theta}\right]. \tag{7}$$

Note that $\mathsf{I}[\Theta;\tilde{Y}_{1:T}]$ quantifies the amount of information that can be inferred about Θ by observing $\tilde{Y}_{1:T}$. Thus, the inequality above implies that the leakage of information about the private parameter via the output of the privacy filter is limited by the mutual information between the input and the output of the randomizer. Hence, the upper bound on the mutual information between Θ and $\tilde{\Theta}$ in (6) ensures the privacy of Θ.

Remark 4: To prove Theorem 3, we first show that the Markov chain $\Theta \longrightarrow \tilde{\Theta} \longrightarrow \tilde{Y}_{1:T}$ holds (see Appendix A for more details). This Markov chain along with the data processing inequality [26] allow us to establish the inequality in (7). Finally, the lower bound in Theorem 3 is obtained by combining Fano's inequality with inequality (7).

IV. Special Cases

In Section III, we studied the structure of the nonlinear transformation without imposing any restriction on the joint pdf of the measurements. In this section, we study the structure of the nonlinear transformation in the special cases of i.i.d., Markovian, and Gauss–Markov measurements.

A. Independent and Identically Distributed Measurements

In this subsection, we assume that $Y_1,\ldots,Y_T$ is a sequence of i.i.d. random variables with the common pdf $p_{\theta_i}(y^1,\ldots,y^d)$ when Θ is equal to θ_i. Given $\Theta = \theta_i$ and $\tilde{\Theta} = \tilde{\theta}_j$, the privacy filter at time-step k generates y_k^l according to

$$\begin{aligned} \tilde{y}_k^l &= F_{l,\tilde{\theta}_j}^{-1}\left(u_k^l\left|\tilde{y}_k^{1:l-1}\right.\right), \\ u_k^l &= F_{l,\theta_i}\left(y_k^l\left|y_k^{1:l-1}\right.\right) \end{aligned}$$

where

$$F_{l,\theta_i}\left(z\left|y_k^{1:l-1}\right.\right) = \int_{-\infty}^{z} p_{\theta_i}\left(y^l\left|y_k^{1:l-1}\right.\right)dy^l.$$

Note that $\tilde{Y}_k$ is generated only using Y_k when the measurements are i.i.d.

B. Markovian Measurements

When measurements are Markovian, the conditional joint pdf of $Y_1,\ldots,Y_T$ given $\Theta = \theta_i$ factorizes as

$$p_{\theta_i}\left(y_1,\ldots,y_T\right) = p_{1,\theta_i}\left(y_1\right)\prod_{k=1}^{T-1} p_{k+1,\theta_i}\left(y_{k+1}\left|y_k\right.\right).$$

In this case, the privacy filter at time-step k generates $\tilde{y}_k^l$ according to

$$\tilde{y}_k^l = F^{-1}_{l,k,\tilde{\theta}_j}\left(u_k^l \left| \tilde{y}_k^{1:l-1}, \tilde{y}_{k-1}\right.\right),$$
$$u_k^l = F_{l,k,\theta_i}\left(y_k^l \left| y_k^{1:l-1}, y_{k-1}\right.\right)$$

where

$$F_{l,k,\theta_i}\left(z \left| y_k^{1:l-1}, y_{k-1}\right.\right) = \int_{-\infty}^{z} p_{k,\theta_i}\left(y^l \left| y_k^{1:l-1}, y_{k-1}\right.\right) dy^l.$$

Note that, different from the general case, $\tilde{Y}_k$ is generated using Y_{k-1} and Y_k in the Markovian case.

C. Gauss–Markov Measurements

Given $\Theta = \theta_i$ in the Gauss–Markov case, the measurements are generated by the following model:

$$\begin{aligned} X_{k+1} &= A_i X_k + W_k \\ Y_k &= C_i X_k + V_k \end{aligned} \tag{8}$$

where $\{W_k\}_k$ and $\{V_k\}_k$ are sequences of i.i.d. zero mean Gaussian random vectors with the covariance matrices Q_w^i and Q_v^i, respectively. We assume that X_0 is a zero mean Gaussian random vector with the covariance matrix Q_0^i. We also assume that $\{W_k\}_k$ and $\{V_k\}_k$ are mutually independent and X_0 is independent of $\{W_k, V_k\}_k$. For all i, we assume that A_i is Schur stable, (A_i, C_i) is observable, and the matrices $Q_w^i, Q_v^i,$ and Q_0^i are positive definite.

In the Gauss–Markov case, the parameter $\theta_i = (A_i, C_i, Q_w^i, Q_v^i, Q_0^i)$ characterizes the joint pdf of the measurements. Given $\tilde{\Theta} = \tilde{\theta}_j = (A_j, C_j, Q_w^j, Q_v^j, Q_0^j)$, the objective of the nonlinear transformation is to sequentially generate $\tilde{Y}_1, \ldots, \tilde{Y}_T$ such that their joint pdf is the same as the joint pdf of a sequence generated by a Gauss–Markov model with the parameters $(A_j, C_j, Q_w^j, Q_v^j, Q_0^j)$.

Before proceeding with the structure of the nonlinear transformation in the Gauss–Markov case, we introduce the structure of the optimal output predictor of the model in (8). Let $\hat{x}_{k|k}$ denote the Kalman estimate of X_k based on $\{Y_1 = y_1, \ldots, Y_k = y_k, \Theta = \theta_i\}$ which is given by (9) with $\hat{x}_{1|0} = 0$ where $\Sigma^i_{k|k-1}$ is the one-step ahead prediction error covariance which satisfies the algebraic Riccati recursion in (10).

$$\begin{aligned} \hat{x}_{k|k} &= \hat{x}_{k|k-1} + \Sigma^i_{k|k-1} C_i^\top \left(C_i \Sigma^i_{k|k-1} C_i^\top + Q_v^i\right)^{-1} \\ &\quad \times \left(y_k - C_i \hat{x}_{k|k-1}\right). \\ \hat{x}_{k|k-1} &= A_i \hat{x}_{k-1|k-1}. \end{aligned} \tag{9}$$

$$\begin{aligned} \Sigma^i_{k+1|k} &= A_i \Sigma^i_{k|k-1} A_i^\top + Q_w^i - A_i \Sigma^i_{k|k-1} C_i^\top \\ &\quad \left(C_i \Sigma^i_{k|k-1} C_i^\top + Q_v^i\right)^{-1} C_i \Sigma^i_{k|k-1} A_i^\top. \end{aligned} \tag{10}$$

$$\begin{aligned} \hat{\tilde{y}}_{k|k-1} &= C_j \hat{\tilde{x}}_{k|k-1}, \\ \hat{\tilde{x}}_{k|k} &= \hat{\tilde{x}}_{k|k-1} + \Sigma^j_{k|k-1} C_j^\top \left(C_j \Sigma^j_{k|k-1} C_j^\top + Q_v^j\right)^{-1} \\ &\quad \times \left(\tilde{y}_k - C_j \hat{\tilde{x}}_{k|k-1}\right). \\ \hat{\tilde{x}}_{k|k-1} &= A_j \hat{\tilde{x}}_{k-1|k-1}. \end{aligned} \tag{11}$$

$$\begin{aligned} \Sigma^j_{k+1|k} &= A_j \Sigma^j_{k|k-1} A_j^\top + Q_w^j - A_j \Sigma^j_{k|k-1} C_j^\top \\ &\quad \times \left(C_j \Sigma^j_{k|k-1} C_j^\top + Q_v^j\right)^{-1} C_j \Sigma^j_{k|k-1} A_j^\top. \end{aligned} \tag{12}$$

Let $\hat{y}_{k|k-1}$ denote the one-step ahead Kalman predictor of Y_k based on $\{Y_1 = y_1, \ldots, Y_{k-1} = y_{k-1}, \Theta = \theta_i\}$ which is given by

$$\hat{y}_{k|k-1} = C_i \hat{x}_{k|k-1}$$

with the output prediction error covariance matrix $\Sigma^{o,i}_{k|k-1}$ defined as

$$\Sigma^{o,i}_{k|k-1} = C_i \Sigma^i_{k|k-1} C_i^\top + Q_v^i.$$

The next lemma studies the pdf of the lth component of Y_k given $\{Y_k^{1:l-1} = y_k^{1:l-1}, Y_{1:k-1} = y_{1:k-1}, \Theta = \theta_i\}$. This result will be used to study the structure of the nonlinear transformation in the Gauss–Markov case.

Lemma 1: The conditional pdf of Y_k^l given $\{Y_k^{1:l-1} = y_k^{1:l-1}, Y_{1:k-1} = y_{1:k-1}, \Theta = \theta_i\}$ is a Gaussian pdf with mean $\hat{\mu}_k^{li}$ and variance $\sigma_k^{l,i}$ given by

$$\begin{aligned} \hat{\mu}_k^{l,i} &= \hat{y}^l_{k|k-1} - \Delta_k^{l,i} \left[\Sigma^{o,i}_{k|k-1}\right]^{-1}_{l-1} \left(y_k^{1:l-1} - \hat{y}^{1:l-1}_{k|k-1}\right) \\ \sigma_k^{l,i} &= \Sigma^{l,o,i}_{k|k-1} - \Delta_k^{l,i} \left[\Sigma^{o,i}_{k|k-1}\right]^{-1}_{l-1} \left(\Delta_k^{l,i}\right)^\top \end{aligned} \tag{13}$$

where $\hat{y}^l_{k|k-1}$ is the lth element of $\hat{y}_{k|k-1}$, $\hat{y}^{1:l-1}_{k|k-1}$ is the vector of the first $l-1$ elements of $\hat{y}_{k|k-1}$, $[\Sigma^{o,i}_{k|k-1}]_{l-1}$ is a matrix formed by the elements in the first $l-1$ rows and the first $l-1$ columns of $\Sigma^{o,i}_{k|k-1}$ with $[\Sigma^{o,i}_{k|k-1}]^{-1}_0 = 0$, $\Sigma^{l,o,i}_{k|k-1}$ is the lth diagonal entry of $\Sigma^{o,i}_{k|k-1}$, and $\Delta_k^{l,i}$ is the vector of the first $l-1$ elements in the lth row of $\Sigma^{o,i}_{k|k-1}$.

Proof: See Appendix E. ■

Lemma 1 implies that $\hat{\mu}_k^{l,i}$ is the minimum mean square error (mmse) predictor of Y_k given $\{Y_k^{1:l-1} = y_k^{1:l-1}, Y_{1:k-1} = y_{1:k-1}, \Theta = \theta_i\}$ and $\sigma_k^{l,i}$ is the error variance associated with $\hat{\mu}_k^{l,i}$. The next theorem studies the structure of the nonlinear transformation in the Gauss–Markov case.

Theorem 4: Consider the Gauss–Markov model with $\Theta = \theta_i$ and $\tilde{\Theta} = \tilde{\theta}_j$. Then, for $1 \le l \le d$, the lth component of the $\tilde{Y}_k^l$ is generated according to

$$\tilde{y}_k^l = F^{-1}\left(u_k^l, \hat{\mu}_k^{l,j}, \sigma_k^{l,j}\right)$$
$$u_k^l = F\left(y_k^l, \hat{\mu}_k^{l,i}, \sigma_k^{l,i}\right)$$

where $F(\cdot, \hat{\mu}, \sigma)$ is the cdf of a Gaussian random variable with mean $\hat{\mu}$ and variance σ, $F^{-1}(\cdot, \hat{\mu}, \sigma)$ is the inverse function of $F(\cdot, \hat{\mu}, \sigma)$, $\hat{\tilde{y}}_{k|k-1}$ and $\Sigma^{o,j}_{k|k-1}$ are given by (11) and (12) with $\hat{\tilde{x}}_{1|0} = 0$, $\hat{\mu}_k^{l,i}$ and $\sigma_k^{l,i}$ are defined in (13), and $\sigma_k^{l,j}$ and $\hat{\mu}_k^{l,j}$ are defined in a similar way.

Proof: According to Lemma 1, the conditional pdf of Y_k^l given $\{Y_k^{1:l-1} = y_k^{1:l-1}, Y_{1:k-1} = y_{1:k-1}, \Theta = \theta_i\}$ is a Gaussian pdf with mean $\hat{\mu}_k^{li}$ and variance $\sigma_k^{l,i}$. Thus, we have

$$p_{\theta_i}\left(x \left| y_k^{1:l-1}, y_{1:k-1}\right.\right) = \frac{1}{\sqrt{2\pi\sigma_k^{l,i}}} \mathrm{e}^{-\frac{1}{2\sigma_k^{l,i}}\left(x - \hat{\mu}_k^{li}\right)^2}.$$

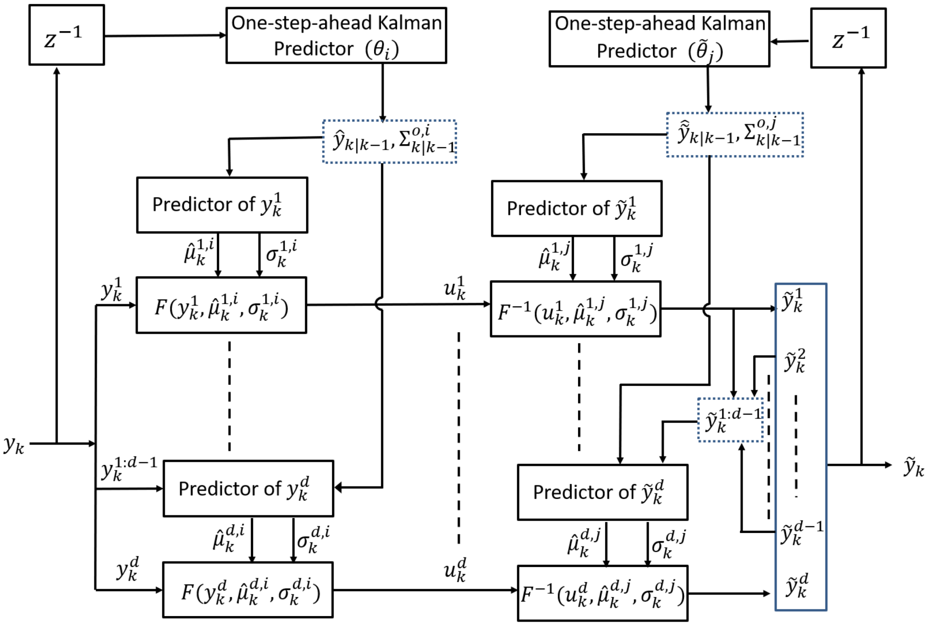
z^{-1}
One-step-ahead Kalman Predictor (θ_i)
$\hat{y}_{k|k-1}, \Sigma^{o,i}_{k|k-1}$
Predictor of y_k^1
$\hat{\mu}_k^{1,i}$
$\sigma_k^{1,i}$
y_k^1
$F(y_k^1, \hat{\mu}_k^{1,i}, \sigma_k^{1,i})$
u_k^1
y_k
$y_k^{1:d-1}$
Predictor of y_k^d
$\hat{\mu}_k^{d,i}$
$\sigma_k^{d,i}$
y_k^d
$F(y_k^d, \hat{\mu}_k^{d,i}, \sigma_k^{d,i})$
u_k^d
One-step-ahead Kalman Predictor $(\tilde{\theta}_j)$
z^{-1}
$\hat{\tilde{y}}_{k|k-1}, \Sigma^{o,j}_{k|k-1}$
Predictor of $\tilde{y}_k^1$
$\hat{\mu}_k^{1,j}$
$\sigma_k^{1,j}$
$F^{-1}(u_k^1, \hat{\mu}_k^{1,j}, \sigma_k^{1,j})$
$\tilde{y}_k^1$
$\tilde{y}_k^2$
$\tilde{y}_k^{1:d-1}$
Predictor of $\tilde{y}_k^d$
$\hat{\mu}_k^{d,j}$
$\sigma_k^{d,j}$
$F^{-1}(u_k^d, \hat{\mu}_k^{d,j}, \sigma_k^{d,j})$
$\tilde{y}_k^{d-1}$
$\tilde{y}_k^d$
$\tilde{y}_k$

Fig. 5. Structure of the nonlinear transformation in the Gauss–Markov case.

This implies that $F_{l,k,\theta_i}(z|y_k^{1:l-1}, y_{1:k-1}) = F(z, \hat{\mu}_k^{l,i}, \sigma_k^{l,i})$ where $F(\cdot, \hat{\mu}, \sigma)$ is the cdf of a Gaussian random variable with mean $\hat{\mu}$ and variance σ. Using a similar argument, it is straightforward to show $F^{-1}_{l,k,\tilde{\theta}_j}(z|\tilde{y}_k^{1:l-1}, \tilde{y}_{1:k-1}) = F^{-1}(z, \hat{\mu}_k^{l,j}, \sigma_k^{l,j})$ where $F^{-1}(\cdot, \hat{\mu}, \sigma)$ is the inverse function of $F(\cdot, \hat{\mu}, \sigma)$. Using these observations and Theorem 1, $\tilde{y}_k^l$, in the Gauss–Markov case, is generated according to

$$\tilde{y}_k^l = F^{-1}\left(u_k^l, \hat{\mu}_k^{l,j}, \sigma_k^{l,j}\right),$$
$$u_k^l = F\left(y_k^l, \hat{\mu}_k^{l,i}, \sigma_k^{l,i}\right)$$

which completes the proof. ■

According to Theorem 4, the structure of the nonlinear transformation in the Gauss–Markov case is characterized by two types of predictors: one-step ahead Kalman predictors and one-component ahead predictors. At time-step k, the Kalman predictors compute $\hat{y}_{k|k-1}$ and $\hat{\tilde{y}}_{k|k-1}$ which, respectively, are the optimal prediction of Y_k given $\{Y_{1:k-1} = y_{1:k-1}, \Theta = \theta_i\}$ and the optimal prediction of $\tilde{Y}_k$ given $\{\tilde{Y}_{1:k-1} = \tilde{y}_{1:k-1}, \tilde{\Theta} = \tilde{\theta}_j\}$. At time-step k, the optimal one-component ahead predictors use $\hat{y}_{k|k-1}$ and $\hat{\tilde{y}}_{k|k-1}$ to compute $\hat{\mu}^{l,i}$ and $\hat{\mu}^{l,j}$ which are the mmse prediction of Y_k^l given $\{Y_k^{1:l-1} = y_k^{1:l-1}, Y_{1:k-1} = y_{1:k-1}, \Theta = \theta_i\}$ and the mmse prediction of $\tilde{Y}_k^l$ given $\{\tilde{Y}_k^{1:l-1} = \tilde{y}_k^{1:l-1}, \tilde{Y}_{1:k-1} = \tilde{y}_{1:k-1}, \tilde{\Theta} = \tilde{\theta}_j\}$, respectively. The outputs of these predictors are used to compute the parameters of the nonlinear transformation in the Gauss–Markov case as shown in Fig. 5.

Remark 5: The Kalman predictors significantly reduce the computational complexity of generating the disguised measurements. In the Gauss–Markov case, the proposed filtering scheme requires the conditional pdfs $p_{\theta_i}(y_k|y_{1:k-1})$ and $p_{\theta_i}(\tilde{y}_k|\tilde{y}_{1:k-1})$ which can be computed (recursively) with low computational costs using the filtering equations in (9)–(12). Note that the direct computation of these pdfs becomes prohibitive when k is large.

V. Numerical Results

In this section, we numerically evaluate the performance of the proposed framework in ensuring the occupancy privacy of the CO_2 measurements in a building automation application. To this end, let Y_k denote the CO_2 measurement inside a room at time-step k which evolves according to

$$X_{k+1} = 0.95X_k + W_k + 10\Theta,$$
$$Y_k = X_k + V_k$$

where Θ denotes the number of occupants in the room, $\{W_k, V_k\}_k$ is a sequence of i.i.d. Gaussian random variables with zero mean and variance 10^{-1}, and X_0 is a Gaussian random variable, independent of $\{W_k, V_k\}_k$, with mean 100 and variance 1. The occupancy parameter Θ is assumed to take values in $\{0, 1\}$ with the probabilities $\mathsf{Pr}(\Theta = 1) = \mathsf{Pr}(\Theta = 0) = 0.5$. We also assume that the pseudo occupancy parameter $\tilde{\Theta}$ takes values in $\{0, 0.2, 0.4, 0.6, 0.8, 1\}$. The horizon length T was set to 50 in our simulations.

Fig. 6 shows the percentage of the relative distortion, due to the privacy filter, as a function of the leakage level of private information I_0. The relative distortion is defined as the ratio of the average distortion between the input and output of the filter over the average CO_2. According to Fig. 6, the relative distortion increases as the leakage of private information becomes small since the privacy constraint in (6) becomes tight in this case. The maximum distortion and maximum privacy level are achieved when leakage of private information is equal to zero. Note that Θ and $\tilde{\Theta}$ are independent when $I_0 = 0$ which results in a high level of relative distortion. The minimum distortion and minimum privacy levels are achieved for $I_0 = 0.69$. In this case, the privacy

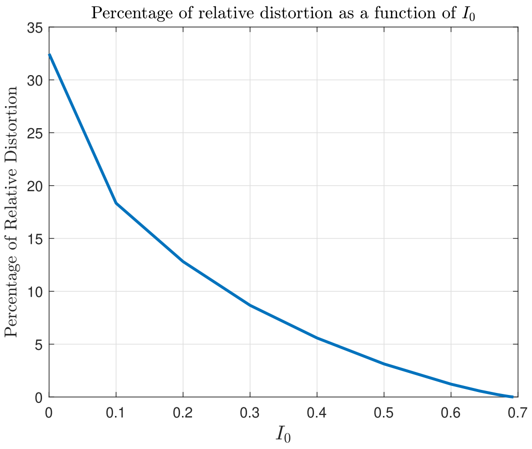

Fig. 6. Percentage of relative distortion versus the leakage level of private information I_0.

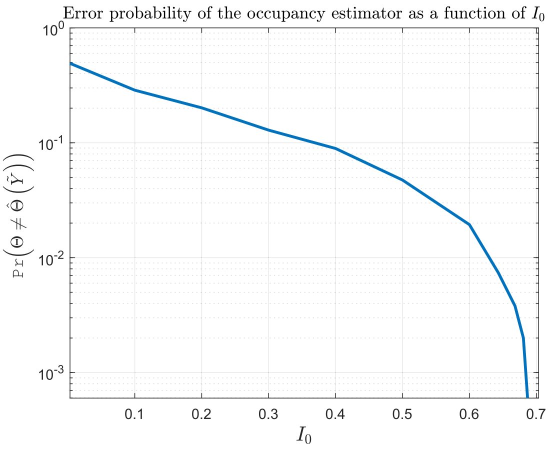

Fig. 7. Error probability of the occupancy estimator versus the leakage level of private information I_0.

constraint is relaxed, and, however, Θ can be correctly inferred from $\tilde{\Theta}$.

To study the performance of an adversary in estimating the occupancy using $\tilde{Y}_k$, we consider the following estimator of the occupancy:

$$\hat{\Theta}\left(\tilde{Y}\right) = \arg\min_{\theta\in\{0,1\}} \left|\frac{\bar{Y} - \bar{\bar{Y}}}{0.95} - \theta\right|$$

where $\bar{Y}$ and $\bar{\bar{Y}}$ are given by

$$\bar{Y} = \frac{1}{10}\sum_{i=41}^{50} \tilde{Y}_i,$$

$$\bar{\bar{Y}} = \frac{1}{10}\sum_{i=40}^{49} \tilde{Y}_i.$$

Fig. 7 shows the error probability of the proposed occupancy estimator as a function of the leakage level of private information I_0. Based on this figure, the proposed estimator can reliably estimate the occupancy information when the leakage of private information is high. However, as I_0 decreases from 0.68 to 0.3, the performance of the occupancy estimator degrades by more than two orders of magnitude while the distortion due to the privacy filter at $I_0 = 0.3$ is approximately 0.8% . This is due to the fact that the output of the filter ceases to be a reliable source of information for estimating the occupancy as the leakage of private information decreases.

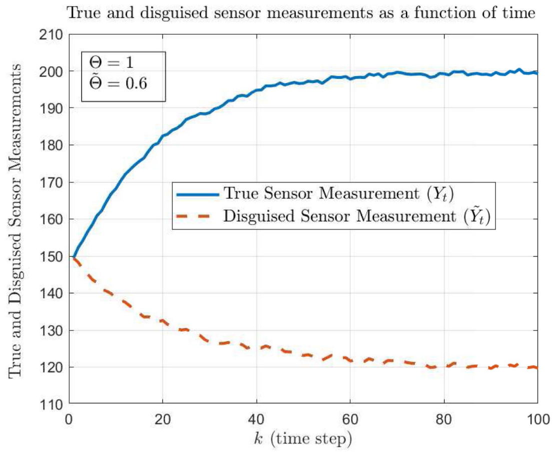

Fig. 8. Realizations of the true and disguised sensor measurements as a function of time for $\Theta = 1$ and $\tilde{\Theta} = 0.6$.

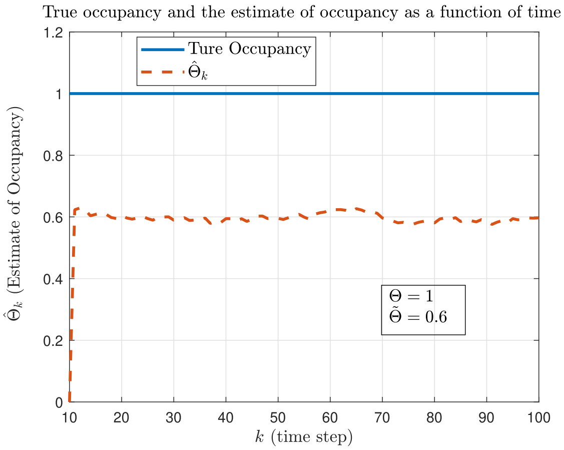

Fig. 9. Output of the occupancy estimator, based on the disguised measurements, as a function of time.

Fig. 8 shows realizations of the true and disguised sensor measurements as a function of time for $\Theta = 1$ and $\tilde{\Theta} = 0.6$. According to this figure, the privacy filter results in a certain level of distortion between the true and disguised measurements. Fig. 9 shows the output of the occupancy estimator in (2) when the output of the nonlinear transformation is used as the input to the estimator. According to this figure, the occupancy estimator cannot accurately infer the occupancy based on the disguised measurements.

VI. Conclusion

In this article, we proposed a privacy filter design framework to ensure the statistical parameter privacy of a sequence of sensor measurements. Under the proposed framework, the privacy filter has two components: a randomizer and a nonlinear transformation. The optimal design of randomizer was studied under a privacy constraint and it was shown than the optimal randomizer is the solution of a convex optimization problem. The privacy level of the proposed framework was examined using information-theoretic inequalities. We also studied the structure of the nonlinear transformation in special cases.

An important direction for our future research is to investigate the optimal design of the randomization probabilities and nonlinear transformation for the privacy-aware closed-loop control problem. Another important avenue for our future work is the optimal design of the randomizer when the private parameter is time-varying and the randomizer has the causal knowledge of the private parameter, i.e., at each time-step, the randomizer only has the knowledge of the current and past values of the private parameter.

Appendix A
Proof of Theorem 1

To prove this result, we first show that the Markov chain $\Theta \longrightarrow \tilde{\Theta} \longrightarrow \tilde{Y}_{1:T}$ holds. To this end, we derive an expression for the conditional pdf of $\tilde{Y}_k^l$ given the event $\{\tilde{Y}_k^{1:l-1} = \tilde{y}_k^{1:l-1}, \tilde{Y}_{1:k-1} = \tilde{y}_{1:k-1}, \Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j\}$ in the next lemma.

Lemma 2: Let $p_{\tilde{Y}_k^l}(\tilde{y}_k^l|\tilde{y}_k^{1:l-1}, \tilde{y}_{1:k-1}, \Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j)$ denote the conditional pdf of $\tilde{Y}_k^l$ given the event $\{\tilde{Y}_k^{1:l-1} = \tilde{y}_k^{1:l-1}, \tilde{Y}_{1:k-1} = \tilde{y}_{1:k-1}, \Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j\}$. Then, we have

$$p_{\tilde{Y}_k^l}\left(\tilde{y}_k^l \left| \tilde{y}_k^{1:l-1}, \tilde{y}_{1:k-1}, \Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j\right.\right)$$
$$= p_{\tilde{\theta}_j}\left(\tilde{y}_k^l \left| \tilde{y}_k^{1:l-1}, \tilde{y}_{1:k-1}\right.\right)$$

where $p_{\tilde{\theta}_j}(\tilde{y}_k^l|\tilde{y}_k^{1:l-1}, \tilde{y}_{1:k-1})$ is obtained from $p_{\tilde{\theta}_j}(\tilde{y}_1, \ldots, \tilde{y}_T)$ using the Bayes' rule and marginalization.

Proof: See Appendix B ■

Let $p_{\tilde{Y}_{1:T}}(\tilde{y}_{1:T}|\Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j)$ denote the conditional pdf of $\tilde{Y}_{1:T}$ given $\{\Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j\}$ which can be written as follows:

$$p_{\tilde{Y}_{1:T}}\left(\tilde{y}_{1:T} \left| \Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j\right.\right)$$
$$\overset{(a)}{=} \prod_{k=1}^{T} \prod_{l=1}^{d} p_{\tilde{Y}_k^l}\left(\tilde{y}_k^l \left| \tilde{y}_k^{1:l-1}, \tilde{y}_{1:k-1}, \Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j\right.\right)$$
$$\overset{(b)}{=} \prod_{k=1}^{T} \prod_{l=1}^{d} p_{\tilde{\theta}_j}\left(\tilde{y}^l \left| \tilde{y}_k^{1:l-1}, \tilde{y}_{1:k-1}\right.\right)$$
$$\overset{(c)}{=} p_{\tilde{\theta}_j}\left(\tilde{y}_1, \ldots, \tilde{y}_T\right) \tag{14}$$

where (a) and (c) follow from the Bayes' rule and (b) follows from Lemma 2. This implies that the conditional pdf of $\tilde{Y}_1, \ldots, \tilde{Y}_T$ given Θ and $\tilde{\Theta}$ only depends on $\tilde{\Theta}$. Thus, the following Markov chain holds $\Theta \longrightarrow \tilde{\Theta} \longrightarrow \tilde{Y}_{1:T}$ and, for all i, j, we have

$$p_{\tilde{Y}_{1:T}}\left(\tilde{y}_{1:T} \left| \tilde{\Theta} = \tilde{\theta}_j\right.\right) = p_{\tilde{Y}_{1:T}}\left(\tilde{y}_{1:T} \left| \Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j\right.\right)$$
$$= p_{\tilde{\theta}_j}\left(\tilde{y}_1, \ldots, \tilde{y}_T\right)$$

which completes the proof.

Appendix B
Proof of Lemma 2

Note that the conditional cdf of $\tilde{Y}_k^l$ given the event $\{\tilde{Y}_k^{1:l-1} = \tilde{y}_k^{1:l-1}, \tilde{Y}_{1:k-1} = \tilde{y}_{1:k-1}, \Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j\}$ can be written as (15), where (a) follows from the definition of $\tilde{Y}_k^l$

$$\tilde{Y}_k^l = F^{-1}_{l,k,\tilde{\theta}_j}\left(U_k^l \left| \tilde{Y}_k^{1:l-1}, \tilde{Y}_{1:k-1}\right.\right)$$
$$U_k^l = F_{l,k,\theta_i}\left(Y_k^l \left| Y_k^{1:l-1}, Y_{1:k-1}\right.\right)$$

and (b) follows from the fact that $F_{l,k,\tilde{\theta}_j}(\cdot|\cdot)$ is invertible with respect to its first argument. We next show that U_k^l given $\{\tilde{Y}_k^{1:l-1} = \tilde{y}_k^{1:l-1}, \tilde{Y}_{1:k-1} = \tilde{y}_{1:k-1}, \Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j\}$ is a uniformly distributed random variable. To this end, let z denote a real number from the interval $[0, 1]$. Then, we have (16) where (a) follows from the fact that $y_k^{1:l-1}, y_{1:k}$ can be uniquely obtained from $\tilde{y}_k^{1:l-1}, \tilde{y}_{1:k}$ since each nonlinear mapping $F_{l,k,\theta}(\cdot|\cdot)$ is invertible with respect to its first argument and (b) follows from the Markov chain $\tilde{\Theta} \longrightarrow (\Theta, Y_k^{1:l-1}, Y_{1:k}) \longrightarrow Y_k^l$.

Combining (15) and (16), we have

$$\mathsf{Pr}\left(\tilde{Y}_k^l \leq \tilde{y}_k^l \left| \tilde{y}_k^{1:l-1}, \tilde{y}_{1:k-1}, \Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j\right.\right)$$
$$= F_{l,k,\tilde{\theta}_j}\left(\tilde{y}_k^l \left| \tilde{y}_k^{1:l-1}, \tilde{y}_{1:k-1}\right.\right)$$

which implies that the conditional pdf of $\tilde{Y}_k^l$ given the event $\{\tilde{Y}_k^{1:l-1} = \tilde{y}_k^{1:l-1}, \tilde{Y}_{1:k-1} = \tilde{y}_{1:k-1}, \Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j\}$ is given by $p_{\tilde{\theta}_j}(\tilde{y}_k^l|\tilde{y}_k^{1:l-1}, \tilde{y}_{1:k})$.

$$\mathsf{Pr}\left(\tilde{Y}_k^l \leq \tilde{y}_k^l \left| \tilde{y}_k^{1:l-1}, \tilde{y}_{1:k-1}, \Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j\right.\right)$$
$$\overset{(a)}{=} \mathsf{Pr}\left(F^{-1}_{l,k,\tilde{\theta}_j}\left(U_k^l \left| \tilde{y}_k^{1:l-1}, \tilde{y}_{1:k-1}\right.\right) \leq \tilde{y}_k^l \right| \tilde{y}_k^{1:l-1},$$
$$\left. \tilde{y}_{1:k-1}, \Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j\right)$$
$$\overset{(b)}{=} \mathsf{Pr}\left(U_k^l \leq F_{l,k,\tilde{\theta}_j}\left(\tilde{y}_k^l \left| \tilde{y}_k^{1:l-1}, \tilde{y}_{1:k-1}\right.\right) \right| \tilde{y}_k^{1:l-1},$$
$$\left. \tilde{y}_{1:k-1}, \Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j\right) \tag{15}$$

$$\mathsf{Pr}\left(U_k^l \leq z \left| \tilde{y}_k^{1:l-1}, \tilde{y}_{1:k-1}, \Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j\right.\right)$$
$$\overset{(a)}{=} \mathsf{Pr}\left(Y_k^l \leq F^{-1}_{\theta_i}\left(z \left| y_k^{1:l-1}, y_{1:k-1}\right.\right) \right| y_k^{1:l-1},$$
$$\left. y_{1:k-1}, \Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j\right)$$
$$\overset{(b)}{=} F_{\theta_i}\left(F^{-1}_{\theta_i}\left(z \left| y_k^{1:l-1}, y_{1:k-1}\right.\right) \left| y_k^{1:l-1}, y_{1:k-1}\right.\right).$$
$$= z \tag{16}$$

Appendix C
Proof of Theorem 2

The objective function in (6) can be written as

$$\frac{1}{T}\sum_{k=1}^{T} \mathsf{E}\left[d\left(Y_k, \tilde{Y}_k\right)\right]$$
$$= \frac{1}{T}\sum_{k=1}^{T}\sum_{i,j} \mathsf{E}\left[d\left(Y_k, \tilde{Y}_k\right) \middle| \Theta = \theta_i, \tilde{\Theta} = \tilde{\theta}_j\right] P_{ji} \mathsf{Pr}\left(\Theta = \theta_i\right)$$

$$
\begin{aligned}
&= \frac{1}{T}\sum_{k=1}^{T}\sum_{i,j}\mathsf{E}\left[d\left(Y_k, \Phi_k\left(Y_{1:k}, \theta_i, \tilde{\theta}_j\right)\right)\right] P_{ji}p_i \\
&= \frac{1}{T}\sum_{k=1}^{T}\sum_{i,j} L_k\left(\theta_i, \tilde{\theta}_j\right) P_{ji}p_i \qquad (17)
\end{aligned}
$$

where $\Phi_k(\cdot, \theta, \tilde{\theta})$ is the vector-valued transformation that generates $\tilde{Y}_k$ using $(Y_{1:k}, \Theta, \tilde{\Theta})$ and $L_k(\theta_i, \tilde{\theta}_j) = \mathsf{E}[d(Y_k, \Phi_k(Y_{1:k}, \theta_i, \tilde{\theta}_j))]$. Thus, the objective function is linear in the randomization probabilities. Also, it can be shown that the privacy constraint is convex in the optimization variables [26]. Thus, the optimization problem in (6) is convex.

Appendix D
Proof of Theorem 3

Using Fano's inequality [26], we can lower bound the error probability of any estimator of Θ, based on $\tilde{Y}_{1:T}$, as

$$
\begin{aligned}
\mathsf{Pr}\left(\Theta \neq \hat{\Theta}\left(\tilde{Y}_{1:T}\right)\right) &\geq \frac{\mathsf{H}\left[\Theta \middle| \tilde{Y}_{1:T}\right] - 1}{\log|\boldsymbol{\Theta}|}, \\
&\overset{(a)}{=} \frac{\mathsf{H}[\Theta] - \mathsf{I}\left[\Theta; \tilde{Y}_{1:T}\right] - 1}{\log|\boldsymbol{\Theta}|} \qquad (18)
\end{aligned}
$$

where $\mathsf{H}[\Theta|\tilde{Y}_{1:T}]$ denotes the conditional entropy of Θ given $\tilde{Y}_{1:T}$ and (a) follows from the definition of the mutual information. In Appendix A, we show that the following Markov chain holds

$$
\Theta \longrightarrow \tilde{\Theta} \longrightarrow \tilde{Y}_{1:T}.
$$

Hence, using the data processing inequality [26], we have

$$
\mathsf{I}\left[\Theta; \tilde{Y}_{1:T}\right] \leq \mathsf{I}\left[\Theta; \tilde{\Theta}\right]. \qquad (19)
$$

Combining (19) and (18), we have

$$
\begin{aligned}
\mathsf{Pr}\left(\Theta \neq \hat{\Theta}\left(\tilde{Y}_{1:T}\right)\right) &\geq \frac{\mathsf{H}[\Theta] - \mathsf{I}\left[\Theta; \tilde{\Theta}\right] - 1}{\log|\boldsymbol{\Theta}|} \\
&\overset{(a)}{\geq} \frac{\mathsf{H}[\Theta] - I_0 - 1}{\log|\boldsymbol{\Theta}|}
\end{aligned}
$$

where (a) follows from the fact that the mutual information between the private parameter and the pseudoparameter is upper bounded by the leakage level of private information in (6).

Appendix E
Proof of Lemma 1

Note that the distribution of Y_k given $\{Y_{1:k-1} = y_{1:k-1}, \Theta = \theta_i\}$ is Gaussian with mean $\hat{y}_{k|k-1}$ and covariance $\Sigma^{o,i}_{k|k-1}$ where

$$
\begin{aligned}
\hat{y}_{k|k-1} &= \mathsf{E}\left[Y_k | \, y_{1:k-1}\right] \\
&= C_i \hat{x}_{k|k-1}
\end{aligned}
$$

and

$$
\begin{aligned}
\Sigma^{o,i}_{k|k-1} &= \mathsf{E}\left[\left(Y_k - \hat{y}_{k|k-1}\right)\left(Y_k - \hat{y}_{k|k-1}\right)^\top\right] \\
&= C_i \Sigma^i_{k|k-1} C_i^\top + Q^i_v
\end{aligned}
$$

where $\hat{x}_{k|k-1}$ is the optimal predictor of X_k based on $\{Y_{1:k-1} = y_{1:k-1}, \Theta = \theta_i\}$ and $\Sigma^i_{k|k-1}$ is the error covariance matrix associated with $\hat{x}_{k|k-1}$. Note that the joint pdf of $Y_k^{1:l}$ given $\{Y_{1:k-1} = y_{1:k-1}, \Theta = \theta_i\}$ is also a Gaussian distribution with mean $\hat{y}^{1:l}_{k|k-1}$ and covariance $[\Sigma^{o,i}_{k|k-1}]_l$ where $[\Sigma^{o,i}_{k|k-1}]_l$ is a matrix formed by the elements in the first l rows and the first l columns of $\Sigma^{o,i}_{k|k-1}$. Thus, using the conditional distribution formula for Gaussian random variables, the conditional pdf of Y_k^l given $\{Y_k^{1:l-1} = y_k^{1:l-1}, Y_{1:k-1} = y_{1:k-1}, \Theta = \theta_i\}$ is a Gaussian distribution with mean $\hat{\mu}^{l,i}_{lk}$ and variance $\sigma^{i,l}_{lk}$ where

$$
\begin{aligned}
\hat{\mu}^{l,i}_k &= \hat{y}^l_{k|k-1} - \Delta^{l,i}_k \left[\Sigma^{o,i}_{k|k-1}\right]^{-1}_{l-1}\left(y_k^{1:l-1} - \hat{y}^{1:l-1}_{k|k-1}\right) \\
\sigma^{l,i}_k &= \Sigma^{l,o,i}_{k|k-1} - \Delta^{l,i}_k \left[\Sigma^{o,i}_{k|k-1}\right]^{-1}_{l-1}\left(\Delta^{l,i}_k\right)^\top.
\end{aligned}
$$

$\Sigma^{l,o,i}_{k|k-1}$ is the lth diagonal entry of $\Sigma^{o,i}_{k|k-1}$ and $\Delta^{l,i}_k$ is the vector of the first $l-1$ elements in the lth row of $\Sigma^{o,i}_{k|k-1}$.

References

[1] A. Ebadat, G. Bottegal, D. Varagnolo, B. Wahlberg, and K. H. Johansson, "Regularized deconvolution-based approaches for estimating room occupancies," *IEEE Trans. Automat. Sci. Eng.*, vol. 12, no. 4, pp. 1157–1168, Oct. 2015.

[2] X. He, W. P. Tay, and M. Sun, "Privacy-aware decentralized detection using linear precoding," in *Proc. IEEE Sensor Array Multichannel Signal Process. Workshop*, 2016, pp. 1–5.

[3] M. Sun and W. P. Tay, "Privacy-preserving nonparametric decentralized detection," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2016, pp. 6270–6274.

[4] X. He and W. P. Tay, "Multilayer sensor network for information privacy," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2017, pp. 6005–6009.

[5] J. Liao, L. Sankar, V. Y. F. Tan, and F. P. Calmon, "Hypothesis testing in the high privacy limit," in *Proc. 54th Annu. Allerton Conf. Commun., Control, Comput.*, 2016, pp. 649–656.

[6] E. Nekouei, M. Pirani, H. Sandberg, and K. H. Johansson, "A randomized filtering strategy against inference attacks on active steering control systems," *IEEE Trans. Inf. Forensics Secur.*, vol. 17, pp. 16–27, 2022, doi: 10.1109/TIFS.2021.3130439.

[7] Z. Li and T. J. Oechtering, "Privacy-aware distributed Bayesian detection," *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 7, pp. 1345–1357, Oct. 2015.

[8] Z. Li and T. J. Oechtering, "Privacy-constrained parallel distributed Neyman-Pearson test," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 3, no. 1, pp. 77–90, Mar. 2017.

[9] R. Mochaourab and T. J. Oechtering, "Private filtering for hidden Markov models," *IEEE Signal Process. Lett.*, vol. 25, no. 6, pp. 888–892, Jun. 2018.

[10] T. Tanaka, M. Skoglund, H. Sandberg, and K. H. Johansson, "Directed information as privacy measure in cloud-based control," KTH Roy. Inst. Technol., Sweden, Tech. Rep., 2017. [Online]. Available: https://arxiv.org/abs/1705.02802

[11] P. Venkitasubramaniam, "Privacy in stochastic control: A Markov decision process perspective," in *Proc. 51st Annu. Allerton Conf. Commun., Control, Comput.*, 2013, pp. 381–388.

[12] K. Kalantari, L. Sankar, and O. Kosut, "On information-theoretic privacy with general distortion cost functions," in *Proc. IEEE Int. Symp. Inf. Theory*, 2017, pp. 2865–2869.

[13] Y. O. Basciftci, Y. Wang, and P. Ishwar, "On privacy-utility tradeoffs for constrained data release mechanisms," in *Proc. Inf. Theory Appl. Workshop*, 2016, pp. 1–6.

[14] F. du Pin Calmon and N. Fawaz, "Privacy against statistical inference," in *Proc. 50th Annu. Allerton Conf. Commun., Control, Comput.*, 2012, pp. 1401–1408.

[15] B. Moraffah and L. Sankar, "Information-theoretic private interactive mechanism," in *Proc. 53rd Annu. Allerton Conf. Commun., Control, Comput.*, 2015, pp. 911–918.

[16] E. Nekouei, T. Tanaka, M. Skoglund, and K. H. Johansson, "Information-theoretic approaches to privacy in estimation and control," *Annu. Rev. Control*, vol. 47, pp. 412–422, 2019.

[17] J. L. Ny and G. J. Pappas, "Differentially private filtering," *IEEE Trans. Autom. Control*, vol. 59, no. 2, pp. 341–354, Feb. 2014.

[18] H. Sandberg, G. Dán, and R. Thobaben, "Differentially private state estimation in distribution networks with smart meters," in *Proc. 54th IEEE Conf. Decis. Control*, 2015, pp. 4492–4498.

[19] E. Nozari, P. Tallapragada, and J. Cortés, "Differentially private average consensus: Obstructions, trade-offs, and optimal algorithm design," *Automatica*, vol. 81, pp. 221–231, 2017.

[20] Y. Mo and R. M. Murray, "Privacy preserving average consensus," *IEEE Trans. Autom. Control*, vol. 62, no. 2, pp. 753–765, Feb. 2017.

[21] Y. Wang, Z. Huang, S. Mitra, and G. E. Dullerud, "Differential privacy in linear distributed control systems: Entropy minimizing mechanisms and performance tradeoffs," *IEEE Control Netw. Syst.*, vol. 4, no. 1, pp. 118–130, Mar. 2017.

[22] G. Bassi, M. Skoglund, and P. Piantanida, "Lossy communication subject to statistical parameter privacy," in *Proc. IEEE Int. Symp. Inf. Theory*, 2018, pp. 1031–1035.

[23] I. Ziemann and H. Sandberg, "Parameter privacy versus control performance: Fisher information regularized control," in *Proc. Amer. Control Conf.*, 2020, pp. 1259–1265.

[24] G. Bassi, E. Nekouei, M. Skoglund, and K. H. Johansson, "Statistical parameter privacy," in *Privacy in Dynamical Systems*. Singapore: Springer, 2020, pp. 65–82.

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

[26] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Hoboken, NJ, USA: Wiley-Interscience, 2006.

Ehsan Nekouei (Member, IEEE) received the B.S. degree from the Shahid Bahonar University of Kerman, Kerman, Iran, the M.S. degree from Tarbiat Modares University, Tehran, Iran, and the Ph.D. degree in electrical engineering from the University of Melbourne, Melbourne, VIC, Australia, in 2003, 2006, and 2013, respectively.

He is currently an Assistant Professor with the Department of Electrical Engineering, City University of Hong Kong, Hong Kong. From 2014 to 2019, he held postdoctoral positions at the KTH Royal Institute of Technology, Stockholm, Sweden, and the University of Melbourne, Melbourne, Australia.

His research interests include privacy in networked control systems and integrated processing of human decision-making data.

Henrik Sandberg (Senior Member, IEEE) received the M.Sc. degree in engineering physics and the Ph.D. degree in automatic control from Lund University, Lund, Sweden, in 1999 and 2004, respectively.

He is currently a Professor with the Division of Decision and Control Systems, KTH Royal Institute of Technology, Stockholm, Sweden. From 2005 to 2007, he was a Postdoctoral Scholar with the California Institute of Technology, Pasadena, CA, USA. In 2013, he was a Visiting Scholar with the Laboratory for Information and Decision Systems (LIDS), MIT, Cambridge, MA, USA. He has also held visiting appointments at The Australian National University, Canberra, Australia, and the University of Melbourne, Melbourne, Australia. His research interests include security of cyber–physical systems, power systems, model reduction, and fundamental limitations in control.

Dr. Sandberg was a recipient of the Best Student Paper Award from the IEEE Conference on Decision and Control in 2004, an Ingvar Carlsson Award from the Swedish Foundation for Strategic Research in 2007, and a Consolidator Grant from the Swedish Research Council in 2016. He has served on the editorial boards of IEEE TRANSACTIONS ON AUTOMATIC CONTROL and the *IFAC Journal Automatica*.

Mikael Skoglund (Fellow, IEEE) received the Ph.D. degree in communication systems from the Chalmers University of Technology, Sweden, in 1997.

In 1997, he joined the Royal Institute of Technology (KTH), Stockholm, Sweden, where he was appointed to the Chair in Communication Theory in 2003. At KTH, he heads the Division of Information Science and Engineering, and the Department of Intelligent Systems.

Dr. Skoglund has worked on problems in source-channel coding, coding and transmission for wireless communications, Shannon theory, information and control, and statistical signal processing. He has authored and coauthored more than 185 journal and 400 conference papers.

From 2003 to 2008, he was an Associate Editor for IEEE TRANSACTIONS ON COMMUNICATIONS. From 2008 to 2012, he was on the editorial board for IEEE TRANSACTIONS ON INFORMATION THEORY and starting in the Fall of 2021 he joined it once more. He has served on numerous technical program committees for IEEE sponsored conferences, he was general co-chair for IEEE ITW 2019, and in 2022 he is serving as TPC co-chair for IEEE ISIT.

Karl Henrik Johansson (Fellow, IEEE) received M.Sc. degree in electrical engineering and Ph.D in automatic control from Lund University.

He is currently a Professor with the School of Electrical Engineering and Computer Science at KTH Royal Institute of Technology in Sweden and Director of Digital Futures. He has held visiting positions at UC Berkeley, Caltech, NTU, HKUST Institute of Advanced Studies, and NTNU.

His research interests are in networked control systems and cyber-physical systems with applications in transportation, energy, and automation networks. He is President of the European Control Association and member of the IFAC Council, and has served on the IEEE Control Systems Society Board of Governors and the Swedish Scientific Council for Natural Sciences and Engineering Sciences. He has received several best paper awards and other distinctions from IEEE, IFAC, and ACM. He has been awarded Swedish Research Council Distinguished Professor, Wallenberg Scholar with the Knut and Alice Wallenberg Foundation, Future Research Leader Award from the Swedish Foundation for Strategic Research, the triennial IFAC Young Author Prize, and IEEE Control Systems Society Distinguished Lecturer. He is Fellow of the IEEE and the Royal Swedish Academy of Engineering Sciences.