

Directed Information and Privacy Loss in Cloud-based Control

Takashi Tanaka¹

Mikael Skoglund²

Henrik Sandberg³

Karl Henrik Johansson⁴

Abstract—We consider a cloud-based control framework in which individual clients own their local plants that must be controlled by a public authority. Individual clients desire to keep the local state information as private as possible, as long as the cloud-based controller can provide a given level of quality of service. Based on an axiomatic argument, we show that Kramer’s notion of causally conditioned directed information from the state random variable to a random variable disclosed to the public authority is an appropriate measure of privacy loss. For a special case with the Linear-Quadratic-Gaussian (LQG) setting, we provide a procedure to construct a “privacy filter” that attains the optimal trade-off between privacy loss and control quality.

I. INTRODUCTION

Leveraged by modern cloud computing technologies, the concept of cloud-based control has attracted much attention in several industrial contexts. While it enhances conventional control technologies in various ways, it also introduces new concerns related to privacy. The purpose of this paper is to discuss an appropriate privacy notion for cloud-based control and a framework for privacy protection.

Cloud-based control is a powerful solution in the following industrial situations.

- (i) The clients do not have sufficient computational resources to perform control tasks.
- (ii) Control performance is improved by utilizing global information/sharing operational data with other clients.
- (iii) New kind of services become available by mining large-scale operational data.

Situation (i) occurs when clients are in charge of operating highly complex plants that requires solving large-scale optimization problems in real-time [1]. Another example of this category is a virtual server (cloud) providing multiple web services (clients) with computational resources [2], [3]. In this example, the assigned computational resource to each client can be viewed as a control input, by which the state (e.g., throughput) of the client’s web service is maintained. An example of cloud-based control in category (ii) is traffic navigation and management [4]. Since multiple vehicles (clients) share a common infrastructure, the overall control performance is drastically improved by “centralizing” the decision-making mechanisms. Category (iii) contains various examples, including the concept of *predictive manufacturing* [5]. This is an idea of collecting and analyzing large amount

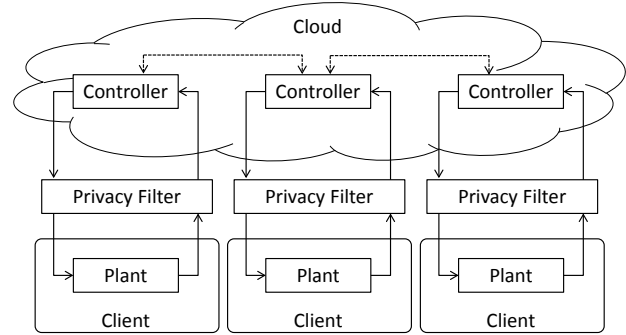


Fig. 1. Cloud-based control.

of operational data from machines in production lines, targeting at improving productivity and safety by predicting failures before they occur.

A. Privacy concerns in cloud-based control

Since cloud-based control mandates clients to share local plant information with the cloud operator, the issue of privacy naturally arises. While careless treatment of privacy leads to undesirable social and economic consequences, controlling privacy is a subtle and challenging engineering task. The difficulty comes from the following reasons.

First, establishing the adequacy of the existing privacy notions (e.g., differential privacy [6], k -anonymity [7], information theoretic privacy [8] among others) in the context of particular applications (in our case, cloud-based control) often requires subtle discussions. In fact, many of the available privacy notions and their validities are sensitive, explicitly or implicitly, to the problem settings where those notions were originally introduced and premises therein. For instance, [9] shows by simple counterexamples that k -anonymity is fragile against side-information. Likewise, [10] and [8] demonstrate that, under several distinct settings, differential privacy does not provide any guarantee in terms of the information theoretic privacy.

Second, privacy notions targeting at single-stage data disclosure mechanisms are in general not sufficient to accommodate privacy issues in cloud-based control. A good privacy notion must also respect the fact that privacy leakage occurs over multiple time steps, and the data from the past stored in the cloud can potentially be used to threaten privacy at the present time. Also, the existence of information feedback must be carefully taken into account. Namely, the cloud has certain influences (through control input) on the future private information (state of the client’s plant), and hence appropriate statistical conditioning is needed to distinguish private information from public information.

School of Electrical Engineering, KTH Royal Institute of Technology. {¹ttanaka, ²skoglund, ³hsan, ⁴kallej}@kth.se. This work is supported by the Knut and Alice Wallenberg Foundation, the Swedish Foundation for Strategic Research, the Swedish Research Council, and the EU CHIST-ERA project COPEs.

Finally, we note that various techniques of data encryption could enhance the security of cloud-based control, but are not sufficient to resolve privacy issues completely. In fact, some applications (e.g., traffic navigation) requires the cloud to have an access to decrypted data (e.g., client's location) to provide an appropriate service. Hence privacy is not protected unless the cloud operator is completely trustworthy and the decrypted data is never transferred to the third party. This is why randomization (e.g., private data is intentionally corrupted by noise before disclosed) is a widely used scheme to trade-off privacy and service utility. However, we also note that there is an emerging control technology by which the cloud is able to perform necessary computation using encrypted data only [11]–[13].

For these reasons, protecting privacy requires sophisticated technologies and hence, it is not practical nor safe to leave the decision on which data to be disclosed/kept as confidential up to each client's discretion. A possible solution is to introduce an additional layer (*privacy filter*) bridging the cloud and clients, which has a dedicated role to control the leakage of private data (Fig. 1). By an appropriate joint design of privacy filter and control algorithms, the overall system is able to balance utility of cloud-based control and privacy losses.

In order to discuss privacy issues quantitatively, this paper introduces an axiomatic characterization of privacy. More specifically, we propose a set of postulates, which is a set of natural properties to be satisfied by a reasonable notion of privacy, and show that a particular function, namely Kramer's notion of *causally conditioned directed information*, arises as a unique candidate to quantify privacy loss in a compatible manner. An axiomatic characterization also provides a convenient interface between theory and practice of privacy discussions. As discussed above, it is often difficult to judge whether a *given* notion of privacy is appropriate for individual applications. In contrast, axioms are often easier to discuss in practical contexts. Axioms also provide a solid mathematical basis on which rigorous theory of privacy can be developed.

B. Related work

Privacy has been extensively studied in the database literature in recent years. While ad hoc approaches for privacy (sub-sampling, aggregation, and suppression) have a long history, one of the first formal definition of privacy is given by k -anonymity [7]. Extensions of this notion include t -closeness and l -diversity [9]. Differential privacy [6] has been particularly popular since its introduction, partly because of its convenient property that no prior on the database content is needed nor used. Information-theoretic privacy in database is also considered in [8] and [14].

Privacy has only relatively recently become a topic of concern in the control-engineering literature. Some of the first works in the area treated consensus algorithms, and how participating agents can maintain some level of privacy despite sharing information with neighbors, see [15], [16], [17]. Differential privacy, which was originally developed

for database privacy, can quite generally be adapted to a control-theoretic context as shown in [18], [19], and also in particular filtering and control applications, see [15], [20], [17], [21], for example. Based on game theory, alternative rigorous notions of privacy in a control and filtering context have been obtained in [22], [23].

A general introduction to information-theoretic security, secrecy and privacy can be found in [24]. Information theory has been used to analyze various aspects of privacy in several different problems settings. The problem of private information-retrieval [25] was considered for example in [26] (and references therein). Recent work reported in [27] introduced the notion of capacity of private information-retrieval, and characterized corresponding fundamental bounds. Information-theoretic tools have also been utilized in the context of differential privacy [28], [29]. Very recent work summarized in [30] studies the relation between differential privacy and privacy quantified in terms of mutual information. This paper also relates these two notions to the concept of identifiability. The work reported in [14] introduced a general framework for establishing a relation between privacy and utility based on rate-distortion arguments. Similarly, [31] developed analytic tools to support the characterization of leakage of privacy in biometric systems.

C. Contribution of this paper

Contributions of this paper are summarized as follows.

- (a) We exemplify a set of postulates (Postulates 1-4) characterizing basic properties of a privacy measure in cloud-based control, and elucidate how Kramer's causally conditioned directed information arises as a unique candidate satisfying it.
- (b) We formulate an optimization problem characterizing optimal joint control and privacy filter policies, and derive its explicit solution in the LQG case.

Although we show that the causally conditioned directed information is the only candidate satisfying the considered set of postulates, we do *not* claim that the considered postulates are the only possible characterization of privacy. In fact, it is our important future work to examine carefully, possibly using real-world incidents of privacy attacks, whether the considered privacy postulates are appropriate or not. At the same time it is worth studying how a different set of postulates leads to a different notion of privacy. We also note that axiomatic consistency is not the only criteria that determines usefulness of privacy notions. For instance, to design a privacy filter according to our privacy notion we need to have a precise knowledge about the system model (e.g., distributions of process noises). This is a weakness compared to the mechanisms based on differential privacy, which does not require prior knowledge of the system.

D. Notation

Random variables are indicated by upper case symbols such as X . We denote by P_X , $P_{X,Y}$ and $P_{X|Y}$ the probability distribution of X , the joint probability distribution of X and Y , and the stochastic kernel of X given Y , respectively.

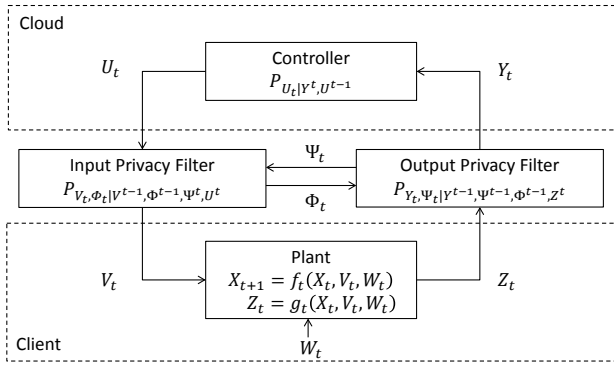


Fig. 2. Privacy filter: General model.

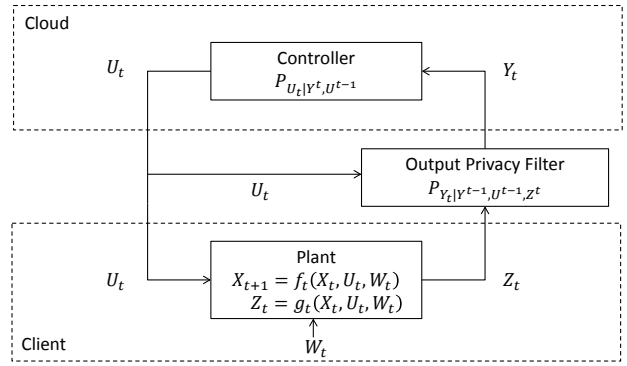


Fig. 3. Privacy filter: Output filter only.

We use notation $P_{X|y}$ to emphasize that it is the conditional probability distribution of X given $Y = y$. We write $H(X|z)$ and $I(X; Y|z)$ to denote the entropy and mutual information evaluated under $P_{X, Y|z}$, and define conditional entropy and conditional mutual information by $H(X|Z) := \mathbb{E}_{P_Z} H(X|z)$ and $I(X; Y|Z) := \mathbb{E}_{P_Z} I(X; Y|z)$. If f is a function of a random variable X , denote by $\mathbb{E}_{P_X} f(X)$ or $\mathbb{E}_{P_X} f(x)$ the mathematical expectation.

II. PROBLEM SETTING

In this paper, a cloud-based control system is modeled by a discrete-time nonlinear stochastic control system. We say that a random variable is *public at time t* if its realization is known to the cloud operator at time t . In contrast, by *private random variable at time t* , we refer to random variables that the client wishes to keep confidential (in an appropriate sense discussed below) at time t .¹ This classification reflects our premise that the cloud operator is “honest but curious” in that it will perform designated control actions faithfully, but will try to learn private data and even to transfer it to the third party. In this paper, we treat the state sequence $X^t \triangleq (X_1, \dots, X_t)$ of the local plant up to time t as the private random variable at time t . We wish to introduce an appropriate measure of privacy loss that occurs during the operation of cloud-based control over a period $1 \leq t \leq T$.

Fig. 2 shows a general structure of privacy filters. An output filter prevents raw sensor data to be disclosed to the cloud. An input privacy filter replaces the control input U_t with a different value V_t to enhance privacy. In general, the input and output filters can communicate with each other via messages Ψ_t and Φ_t . Privacy filters and controller algorithms are in general randomized policies and have memories of the past observations. Thus, we model them as stochastic kernels of the forms specified in Fig. 2. Fig. 3 shows a simpler form of a privacy filter in which the control input commanded by the cloud is directly applied to the plant. Since there is no input filter, this architecture is easier to implement. For the rest of the paper, we focus on this simple architecture in Fig. 3, and discuss privacy notions and privacy filter design problems exclusively for this architecture. In Section III,

¹According to this definition, note that random variables are public or private, or neither.

we characterize our privacy notion axiomatically, and then formulate a joint controller and output privacy filter design problem in Section IV. We derive an optimal form of joint controller and output privacy filter in the LQG regime in Section V.

III. AXIOMATIC CHARACTERIZATION OF PRIVACY

A meaningful notion of privacy must satisfy some basic properties. In this section, we first consider a single-stage data disclosure mechanism and show that the only candidate function that satisfies the natural set of postulates (axioms) is Shannon’s mutual information between private and published random variables. Our argument is aligned with the logic developed in [32], where the mutual information arises as a unique function that characterizes the value of side information in inference problems. The set of axioms used there is simple, and thus we argue that it can be naturally used as a set of axioms for privacy. Then, we apply this observation to multi-stage feedback control systems and show the unique candidate characterizing privacy loss in cloud-based control in a satisfactory manner is the *causally conditioned directed information*.

A. Single-stage case

Suppose X and Y are \mathcal{X} - and \mathcal{Y} -valued random variables with joint distribution $P_{X, Y}$. We temporarily assume that \mathcal{X} and \mathcal{Y} are countably finite sets, and denote by $\mathcal{P}_{\mathcal{X}}$ the space of probability distributions on \mathcal{X} . Assuming that X is a private random variable, we wish to quantify the privacy loss due to the disclosure of a random variable Y .

First, we quantify the “hardness” of inferring X using the notion of *loss function*. Generally speaking, a random variable X is hard to infer if the expected posterior value of observation (i.e., the degree of “surprise” that occurs when observing a realization $x \in \mathcal{X}$) cannot be made small. The posterior value of the observation is a function of the observed realization $x \in \mathcal{X}$ and a prior distribution $Q_X \in \mathcal{P}_{\mathcal{X}}$ assumed by the observer. We refer to such a function $\ell : \mathcal{X} \times \mathcal{P}_{\mathcal{X}} \rightarrow \mathbb{R}$ as a *loss function*. In the literature, it is also called the scoring rule [33] or self-information [34]. Note that for a given choice of ℓ , the task of inference is to minimize ℓ by properly assuming Q_X . Among many

options, the *logarithmic loss function* $\ell(x, Q_X) = \log \frac{1}{Q_X(x)}$ is frequently used in the literature.

Let $P_X \in \mathcal{P}_{\mathcal{X}}$ be the true probability distribution of X , and $Q_X \in \mathcal{P}_{\mathcal{X}}$ be the assumed distribution. In general $P_X \neq Q_X$. If $S(P_X, Q_X) \triangleq \mathbb{E}_{P_X} \ell(X, Q_X)$, the quantity $\inf_{Q_X} S(P_X, Q_X)$ is referred to as the *Bayes envelope*. A loss function is said to be proper if $\inf_{Q_X} S(P_X, Q_X) = S(P_X, P_X)$. It can easily be shown that the logarithmic loss function is proper, and the associated Bayes envelope coincides with the entropy $H(X)$ of X :

$$\inf_{Q_X} S(P_X, Q_X) = S(P_X, P_X) = \mathbb{E}_{P_X} \log \frac{1}{P_X(x)} = H(X).$$

Now we introduce the first postulate characterizing our privacy notion. It states that privacy loss due to disclosing Y is measured by the expected difference in Bayes envelope evaluated before and after observing Y . Given a loss function ℓ and the joint distribution $P_{X,Y}$, we refer to the privacy loss evaluated this way as the *privacy leakage function*, and denote it by $L(\ell, P_{X,Y})$.

Postulate 1: The privacy leakage function is in the form

$$L(\ell, P_{X,Y}) = \inf_{Q_X} \mathbb{E}_{P_X} [\ell(X, Q_X)] - \mathbb{E}_{P_Y} \inf_{Q_{X|Y}} \mathbb{E}_{P_{X|Y}} [\ell(X, Q_{X|Y})].$$

The first term on the right hand side is the Bayes envelope evaluated without side information Y , while in the second term, the assumed distribution $Q_{X|Y}$ is allowed to depend on Y . Hence, $L(\ell, P_{X,Y})$ is understood to be the improvement in the estimation quality due to the side information Y . If the loss function ℓ is logarithmic, the privacy leakage function $L(\ell, P_{X,Y})$ defined above coincides with the mutual information between X and Y , i.e.,

$$L(\ell, P_{X,Y}) = H(X) - H(X|Y) = I(X; Y).$$

Up to now, the logarithmic loss function is just an example among many other possible choices of loss functions. It turns out that it is the only option that satisfies the following natural postulate.

Postulate 2: (Data-processing axiom [32]) For any distribution $P_{X,Y}$ on $\mathcal{X} \times \mathcal{Y}$, the information leakage function $L(\ell, P_{X,Y})$ satisfies

$$L(\ell, P_{T(X),Y}) \leq L(\ell, P_{X,Y}) \quad (1)$$

for every $T : \mathcal{X} \rightarrow \mathcal{X}$ such that $T(X)-X-Y$ and $X-T(X)-Y$ form Markov chains. In (1), the joint distribution $P_{T(X),Y}$ on $\mathcal{X} \times \mathcal{Y}$ is defined by

$$P_{T(X),Y}(T(\mathcal{B}_X) \times \mathcal{B}_Y) = P_{X,Y}(\mathcal{B}_X \times \mathcal{B}_Y)$$

for all subsets \mathcal{B}_X and \mathcal{B}_Y of \mathcal{X} and \mathcal{Y} .

Theorem 1: (Justification of mutual information [32]) Let \mathcal{X} be a finite set with $|\mathcal{X}| \geq 3$. Under Postulate 2, the privacy leakage function is uniquely determined by the mutual information

$$L(\ell, P_{X,Y}) = I(X; Y)$$

up to a positive multiplicative factor.

Proof: See [32]. ■

Remark 1: The result of Theorem 1 can be extended to the case with continuous random variables X and Y using a formula [35, Ch. 2.5], [36, Ch. 3.5], [37, Ch. 7.1]:

$$I(X; Y) = \sup I([X], [Y]). \quad (2)$$

The right-hand-side of (2) denotes the supremum of mutual information between discrete random variables $[X]$ and $[Y]$ over all finite quantizations. If we consider a supremum achieving sequence of quantizers, and require the data-processing axiom to be satisfied by each element of the sequence, we obtain $I(X; Y)$ as the unique privacy leakage function for continuous random variables X and Y .

B. Multi-stage case

Based on the discussion in the previous section, in this section we propose a privacy measure suitable for cloud-based control (Fig. 3). To proceed, we introduce the following additional postulates.

Postulate 3: The private random variable at time t is X^t , while Y^{t-1} and U^{t-1} are public at time t .

We first characterize the instantaneous privacy loss at time step t due to the disclosure of Y_t . By Postulate 3, we need to characterize the privacy leakage function for X^t due to disclosing Y_t under the joint distribution $P_{X^t, Y_t | Y^{t-1}, U^{t-1}}$. Notice that, by Postulate 3, y^{t-1} and u^{t-1} are public knowledge.

Let ℓ be a loss function as in the preceding subsection. For every realization (y^{t-1}, u^{t-1}) , Postulate 2 requires the privacy leakage function $L(\ell, P_{X^t, Y_t | y^{t-1}, u^{t-1}})$ to satisfy

$$L(\ell, P_{T(X^t), Y_t | y^{t-1}, u^{t-1}}) \leq L(\ell, P_{X^t, Y_t | y^{t-1}, u^{t-1}})$$

whenever $T(X^t) \in \mathcal{X}^t$ is a sufficient statistic of X^t for Y_t given $Y^{t-1} = y^{t-1}$ and $U^{t-1} = u^{t-1}$, i.e., the following Markov chains hold under $P_{X^t, Y_t | y^{t-1}, u^{t-1}}$:

$$T(X^t) - X^t - Y_t, \quad X^t - T(X^t) - Y_t.$$

From Theorem 1, we conclude that the only loss function (up to positive multiplicative factors) that satisfies the above inequality is the logarithmic one, and with necessity we have

$$L(\ell, P_{X^t, Y_t | y^{t-1}, u^{t-1}}) = I(X^t; Y_t | y^{t-1}, u^{t-1}).$$

Thus, if Y^{t-1} and U^{t-1} have a joint distribution $P_{Y^{t-1}, U^{t-1}}$, the expected privacy loss at time step t is

$$\mathbb{E}_{P_{Y^{t-1}, U^{t-1}}} L(\ell, P_{X^t, Y_t | y^{t-1}, u^{t-1}}) = I(X^t; Y_t | Y^{t-1}, U^{t-1}).$$

Finally, we assume that our privacy notion satisfies the following natural property.

Postulate 4: The expected total privacy loss over the horizon $t = 1, 2, \dots, T$ has a stage-additive form over the expected instantaneous privacy losses.

Under Postulate 4, the expected total privacy loss is

$$\sum_{t=1}^T I(X^t; Y_t | Y^{t-1}, U^{t-1}) =: I(X^T \rightarrow Y^T || U^{T-1}). \quad (3)$$

The notation on the right hand side of (3) is introduced in [38] to denote the quantity appearing on the left hand side of (3). We refer to this quantity as (Kramer's) causally conditioned directed information. Thus, we obtain:

Proposition 1: Under Postulates 1-4, causally conditioned directed information $I(X^T \rightarrow Y^T \| U^{T-1})$ is the only function (up to positive multiplicative factors) quantifying the expected privacy loss in the cloud-based control in Fig. 3.

IV. PRIVACY-PRESERVING CLOUD-BASED CONTROL DESIGN

Suppose that the performance of the cloud-based control system is measured by a stage-wise additive cost function $\sum_{t=1}^T \mathbb{E}c(X_{t+1}, U_t)$. Then, privacy loss in the cloud-based control with a given control performance requirement δ is minimized by solving

$$\min I(X^T \rightarrow Y^T \| U^{T-1}) \quad (4a)$$

$$\text{s.t.} \quad \sum_{t=1}^T \mathbb{E}c(X_{t+1}, U_t) \leq \delta. \quad (4b)$$

Likewise, the best achievable control performance under the privacy constraint is characterized by flipping the constraint and objective functions in (4). In both cases, the optimization domain is the space of the sequence of Borel measurable stochastic kernels

$$\mathcal{D} = \{P_{U_t|Y^t, U^{t-1}}, P_{Y_t|Y^{t-1}, U^{t-1}, Z_t}\}_{t=1}^T \quad (5)$$

characterizing joint controller and output privacy filter policies.² Since (4) is an infinite dimensional optimization problem, it is in general difficult to obtain an explicit form of an optimal solution. In Section V, we consider a special case in which it is possible.

A. Implication

So far we have provided a justification of $I(X^T \rightarrow Y^T \| U^{T-1})$ as a measure of privacy loss. In this subsection, we study how this quantity imposes a fundamental limitation in estimating private random variables.

Consider an optimal joint controller and output privacy filter policy solving (4), and let γ be the optimal value. By Postulate 4, the total privacy loss γ can be written as $\gamma = \sum_{t=1}^T \gamma_t$, where

$$\gamma_t = I(X^t; Y_t | Y^{t-1}, U^{t-1}) (\geq 0) \quad (6)$$

is the privacy loss at time t . To see how (6) guarantees privacy against inferring X^t at time t even after disclosing Y_t , consider an estimate of X^t of the form $\hat{X}^t : \mathcal{Y}^t \times \mathcal{U}^{t-1} \rightarrow \mathcal{X}^t$. Since realizations of Y^{t-1} and U^{t-1} are prior knowledge at time t , \hat{X}^t can be viewed as a function of Y_t alone, and thus $X^t - Y_t - \hat{X}^t$ forms a Markov chain given (Y^{t-1}, U^{t-1}) . By the data-processing inequality,

$$I(X^t; \hat{X}^t | Y^{t-1}, U^{t-1}) \leq I(X^t; Y_t | Y^{t-1}, U^{t-1}) = \gamma_t.$$

²We consider $P_{U_1|Y^1, U^0} = P_{U_1|Y^1}$ and $P_{Y_1|Y^0, U^0, Z^1} = P_{Y_1|Z^1}$.

In other words, the expected mutual information between X^t and \hat{X}^t is bounded by γ_t :

$$\mathbb{E}_{P_{Y^{t-1}, U^{t-1}}} I(X^t; \hat{X}^t | y^{t-1}, u^{t-1}) \leq \gamma_t. \quad (7)$$

This inequality imposes a fundamental limitation of estimation in the following sense. Let $\rho_t : \mathcal{X}^t \times \mathcal{X}^t \rightarrow [0, \infty)$ be an arbitrary distortion function. For a given source distribution $P_{X^t|y^{t-1}, u^{t-1}}$, let $D_t : [0, \infty) \rightarrow [0, \infty)$ be the distortion-rate function [39]. By definition of the distortion-rate function, for any joint distribution $P_{X^t, \hat{X}^t|y^{t-1}, u^{t-1}}$, we have

$$\mathbb{E}_{P_{X^t, \hat{X}^t|y^{t-1}, u^{t-1}}} \rho_t(X^t, \hat{X}^t) \geq D_t(I(X^t; \hat{X}^t | y^{t-1}, u^{t-1})).$$

Taking expectation with respect to $P_{Y^{t-1}, U^{t-1}}$, we have

$$\mathbb{E} \rho_t(X^t, \hat{X}^t) \geq \mathbb{E}_{P_{Y^{t-1}, U^{t-1}}} D_t(I(X^t; \hat{X}^t | y^{t-1}, u^{t-1})) \quad (8a)$$

$$\geq D_t(\mathbb{E}_{P_{Y^{t-1}, U^{t-1}}} I(X^t; \hat{X}^t | y^{t-1}, u^{t-1})) \quad (8b)$$

$$\geq D_t(\gamma_t). \quad (8c)$$

Recall that distortion-rate functions are in general convex and non-increasing [39, Lemma 10.4.1]. Thus, (8b) follows from Jensen's inequality, and (8c) follows from (7).

Hence, under our privacy notion, (6) ensures that the estimation error corresponding to any estimator \hat{X}^t based on all information available in the cloud at time t cannot be smaller than the distortion-rate function $D_t(\gamma_t)$.

V. LQG CASE

In this section, we consider a special case in which (4) becomes a tractable optimization problem. Suppose the plant in Fig. 3 is a fully observable linear dynamics system

$$X_{t+1} = A_t X_t + B_t U_t + W_t, \quad Z_t = X_t$$

where $W_t \sim \mathcal{N}(0, \Sigma_t^W)$ is a sequence of independent Gaussian random variables. Assume also that $c(\cdot, \cdot)$ is a convex quadratic function, and that the problem (4) can be written as

$$\min I(X^T \rightarrow Y^T \| U^{T-1}) \quad (9a)$$

$$\text{s.t.} \quad \sum_{t=1}^T \mathbb{E}(\|X_{t+1}\|_{Q_t}^2 + \|U_t\|_{R_t}^2) \leq \delta. \quad (9b)$$

The domain of optimization is (5). Problem (9) is identical to the problem considered in [41], except that in [41], an optimal solution is provided under the restriction that the stochastic kernels in (5) are Linear-Gaussian. In what follows, we provide a solution to (9) without such an assumption.

To this end, we consider a related optimization problem

$$\min I(X^T \rightarrow U^T) \quad (10a)$$

$$\text{s.t.} \quad \sum_{t=1}^T \mathbb{E}(\|X_{t+1}\|_{Q_t}^2 + \|U_t\|_{R_t}^2) \leq \delta \quad (10b)$$

where the domain of optimization is \mathcal{D} given by (5). Problem (10) is studied in [40], where it is shown that an optimal solution is a Linear-Gaussian randomized policy shown in

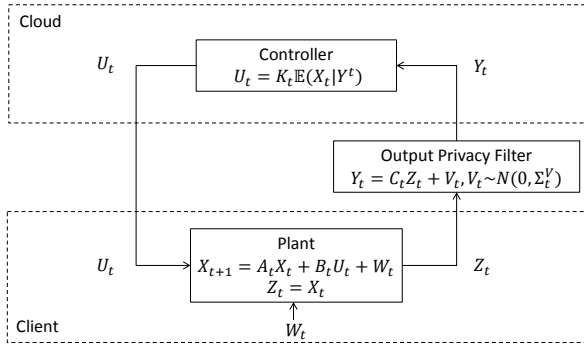


Fig. 4. Structure of optimal policy for problem (10). Matrix parameters $C_t, \Sigma_t^V, L_t, K_t$ for $t = 1, \dots, T$ are determined by an algorithm based on semidefinite programming. See [40] for the details. This structure can also be viewed as an optimal joint controller policy and output privacy filter for the LQG case (9). Although the output privacy filter is allowed to utilize public random variable U^{t-1} (as shown in Fig. 3), it turns out that this information need not be used.

Fig. 4, which can be synthesized by means of semidefinite programming.³ Now we claim the following.

Proposition 2: The policy shown in Fig. 4, which is an optimal solution to (10), is also an optimal solution to (9).

Proof: (Outline only) This result is obtained in the course of proving the main result of [40]. By [40, equation (22)], it is shown that for an optimal policy P^* shown in Fig. 4, an equality

$$I(X^T \rightarrow U^T) = \sum_{t=1}^T I(X_t; Y_t | Y^{t-1})$$

holds. However, for P^* , it is also shown in [40, Section VI-E] that

$$I(X^T \rightarrow Y^T \| U^{T-1}) = \sum_{t=1}^T I(X_t; Y_t | Y^{t-1}).$$

A complete proof is rather lengthy, and must be deferred to [40, Section VI]. ■

Proposition 2 suggests an explicit form of the joint control and output privacy filter policy solving (9).

Proposition 3: An optimal joint controller and output privacy filter characterized by an optimal solution to (9) is in the form shown in Fig. 4. An optimal choice of matrices C_t, Σ_t^V, L_t (Kalman gains) and K_t (feedback control gains) are obtained by semidefinite programming.

Notice that the privacy filter shown in Fig. 4 is similar to privacy protecting mechanisms considered in various other contexts (e.g., [6]) in that it is adding noise V_t before disclosing data. Proposition 3 shows that the optimal noise distribution is Gaussian in the LQG case (9).

VI. NUMERICAL EXAMPLE

In this section, we consider a simple scalar system

$$X_{t+1} = X_t + U_t + W_t, \quad t = 1, \dots, T$$

³In [40], problem (10) is studied with more general optimization domain $\mathcal{D}' \triangleq \{P_{U_t|X^t, U^{t-1}}\}_{t=1}^T$. Note that $\mathcal{D} \subset \mathcal{D}'$ in that every element in \mathcal{D} can be, by compositions of stochastic kernels, mapped to an element of \mathcal{D}' . Since it is shown that a policy P^* shown in Fig. 4 is optimal over \mathcal{D}' , and since P^* is in fact an element of \mathcal{D} , it can be concluded that P^* is an optimal solution to (10).




Fig. 5. SNR of the optimal privacy filter and the privacy loss.

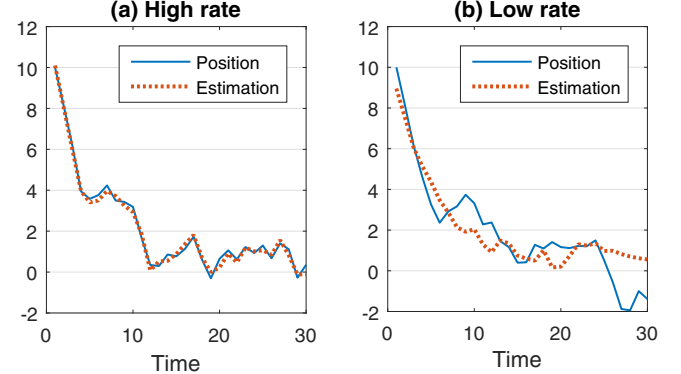


Fig. 6. Simulated trajectory of X_t and its estimate $\mathbb{E}(X_t|Y^t)$.

with a process noise $W_t \sim \mathcal{N}(0, 0.3)$. This example is motivated by a cloud-based navigation service, where the state variable X_t is interpreted as the position of the client at time t , whereas U_t is the navigation signal provided by the cloud. Assuming the initial position is $X_1 = 10$, the cloud-based controller navigates the client to the origin withing $T = 30$ steps only using the output of a privacy filter

$$Y_t = C_t X_t + V_t, \quad V_t \sim \mathcal{N}(0, \Sigma_t^V).$$

The optimal privacy filter is different depending on the choice of δ in (9). We consider two scenarios in which control requirements are stringent ($\delta = 24.9$) and mild ($\delta = 31.9$). In both cases, we use the same control cost function with $Q_t = 1, R_t = 20$ for all t . The former case requires higher data rate (measured in directed information). In each scenario, we compute the optimal sequence $\{C_t, \Sigma_t^V\}_{t=1, \dots, 30}$ by solving (9) using semidefinite programming. The sequence of signal-to-noise ratios $\text{SNR}_t = C_t^2 / \Sigma_t^V$ is plotted in Fig. 5 (a). The total loss of privacy in the high rate scenario is 45.6 [bits] (the area of blue region in Fig. 5 (b)) while it is 8.1 [bits] in the low rate case (the area of red region in Fig. 5 (b)). Fig. 6 shows the closed-loop performance in each scenario. In the high rate case, the cloud estimates the position of the client accurately, and consequently the control performance is preferable. In the low rate case, the estimate is not accurate (privacy is better protected) and the control performance is poor.

VII. SUMMARY AND FUTURE WORK

In this paper, we showed that Kramer's causally conditioned directed information arises as the unique candidate

(up to positive multiplicative factor) for a measure of privacy loss in cloud-based control if Postulates 1-4 (including the *data-processing axiom* [32]) are to be satisfied. This result is the first step towards axiomatic privacy theory, which is a prospective approach providing a convenient interface between theory and practice in privacy discussions. There are numerous further opportunities in the same line of research. Notice that the set of postulates we have selected in this paper is not the only possible characterization of privacy. In fact, in [8], the “maximum” type of privacy leakage function

$$L(\ell, P_{X,Y}) = \inf_{Q_X} \mathbb{E}_{P_X}[\ell(X, Q_X)] - \min_{y \in \mathcal{Y}} \inf_{Q_{X|Y}} \mathbb{E}_{P_{X|Y}}[\ell(X, Q_{X|Y})].$$

is considered in parallel with the “average” type of privacy leakage function we assumed in Postulate 1. It is of great interest whether there exists a valid notion of privacy satisfying the corresponding new set of postulates.

REFERENCES

- [1] T. Hegazy and M. Hefeeda, “Industrial automation as a cloud service,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 26, no. 10, pp. 2750–2763, 2015.
- [2] M. A. Kjær, M. Kihl, and A. Robertsson, “Resource allocation and disturbance rejection in web servers using slas and virtualized servers,” *IEEE Transactions on Network and Service Management*, vol. 6, no. 4, pp. 226–239, 2009.
- [3] H. C. Lim, S. Babu, J. S. Chase, and S. S. Parekh, “Automated control in cloud computing: challenges and opportunities,” *ACM Proceedings of the 1st workshop on Automated control for datacenters and clouds*, 2009.
- [4] B. Besselink, V. Turri, A. A. S.H. van de Hoef, K.-Y. Liang, J. Martensson, and K. H. Johansson, “Cyber-physical control of road freight transport,” *Proceedings of IEEE*, vol. 104, no. 5, pp. 1128–1141, 2016.
- [5] J. Lee, E. Lapira, B. Bagheri, and H.-a. Kao, “Recent advances and trends in predictive manufacturing systems in big data environment,” *Manufacturing Letters*, vol. 1, no. 1, pp. 38–41, 2013.
- [6] C. Dwork, “Differential privacy: A survey of results,” *International Conference on Theory and Applications of Models of Computation*, 2008.
- [7] 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.
- [8] F. du Pin Calmon and N. Fawaz, “Privacy against statistical inference,” *The 50th IEEE Annual Allerton Conference on Communication, Control, and Computing*, 2012.
- [9] A. Machanavajjhala, D. Kifer, J. Gehrke, and M. Venkatasubramanian, “l-diversity: Privacy beyond k-anonymity,” *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 1, no. 1, p. 3, 2007.
- [10] A. De, “Lower bounds in differential privacy,” *Theory of Cryptography Conference*, 2012.
- [11] K. Kogiso and T. Fujita, “Cyber-security enhancement of networked control systems using homomorphic encryption,” *The 54th IEEE Conference on Decision and Control (CDC)*, pp. 6836–6843, 2015.
- [12] F. Farokhi, I. Shames, and N. Batterham, “Secure and private cloud-based control using semi-homomorphic encryption,” *The 6th IFAC Workshop on Distributed Estimation and Control in Networked Systems (NecSys)*, 2016.
- [13] Y. Shoukry, K. Gatsis, A. Alanwar, G. J. Pappas, S. A. Seshia, M. Srivastava, and P. Tabuada, “Privacy-aware quadratic optimization using partially homomorphic encryption,” *The 55th IEEE Conference on Decision and Control (CDC)*, pp. 5053–5058, 2016.
- [14] L. Sankar, S. R. Rajagopalan, and H. V. Poor, “Utility-privacy tradeoffs in databases: An information-theoretic approach,” *IEEE Transactions on Information Forensics and Security*, vol. 8, no. 6, pp. 838–852, June 2013.
- [15] Z. Huang, S. Mitra, and G. Dullerud, “Differentially private iterative synchronous consensus,” *Proceedings of the 2012 ACM Workshop on Privacy in the Electronic Society*, 2012.
- [16] N. E. Manitaras and C. N. Hadjicostis, “Privacy-preserving asymptotic average consensus,” *The 2013 European Control Conference (ECC)*, 2013.
- [17] Y. Mo and R. M. Murray, “Privacy preserving average consensus,” *IEEE Transactions on Automatic Control*, 2016, to appear.
- [18] J. L. Ny and G. J. Pappas, “Differentially private filtering,” *IEEE Transactions on Automatic Control*, vol. 59, no. 2, pp. 341–354, Feb 2014.
- [19] Y. Wang, Z. Huang, S. Mitra, and G. E. Dullerud, “Entropy-minimizing mechanism for differential privacy of discrete-time linear feedback systems,” *The 53rd IEEE Conference on Decision and Control (CDC)*, 2014.
- [20] H. Sandberg, G. Dán, and R. Thobaben, “Differentially private state estimation in distribution networks with smart meters,” *The 54th IEEE Conference on Decision and Control (CDC)*, 2015.
- [21] E. Nozari, P. Tallapragada, and J. Cortés, “Differentially private distributed convex optimization via objective perturbation,” *2016 American Control Conference (ACC)*, 2016.
- [22] E. Akyol, C. Langbort, and T. Basar, “Privacy constrained information processing,” *The 54th IEEE Conference on Decision and Control (CDC)*, 2015.
- [23] F. Farokhi, H. Sandberg, I. Shames, and M. Cantoni, “Quadratic Gaussian privacy games,” *The 54th IEEE Conference on Decision and Control (CDC)*, 2015.
- [24] M. Bloch and J. Barros, *Physical-layer Security*. Cambridge Univ. Press, 2011.
- [25] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan, “Private information retrieval,” in *Proc. IEEE Symposium on Foundations of Computer Science*, 1995.
- [26] C. Cachin, S. Micali, and M. Stadler, “Computationally private information retrieval with polylogarithmic communication,” *International Conference on the Theory and Applications of Cryptographic Techniques*, 1999.
- [27] H. Sun and S. A. Jafar, “The capacity of private information retrieval,” *arXiv:1602.09134 [cs.IT]*.
- [28] M. S. Alvim and M. E. Andrés, “On the relation between differential privacy and quantitative information flow,” *International Colloquium on Automata, Languages, and Programming*, 2011.
- [29] G. Barthe and B. Kopf, “Information-theoretic bounds for differentially private mechanisms,” *The IEEE 24th Computer Security Foundations Symposium*, 2011.
- [30] W. Wang, L. Ying, and J. Zhang, “On the relation between identifiability, differential privacy, and mutual-information privacy,” *IEEE Transactions on Information Theory*, vol. 62, no. 9, pp. 5018–5029, Sept 2016.
- [31] T. Ignatenko and F. M. J. Willems, “Biometric systems: Privacy and secrecy aspects,” *IEEE Transactions on Information Forensics and Security*, vol. 4, no. 4, pp. 956–973, Dec 2009.
- [32] J. Jiao, T. A. Courtade, K. Venkat, and T. Weissman, “Justification of logarithmic loss via the benefit of side information,” *IEEE Transactions on Information Theory*, vol. 61, no. 10, pp. 5357–5365, 2015.
- [33] T. Gneiting and A. E. Raftery, “Strictly proper scoring rules, prediction, and estimation,” *Journal of the American Statistical Association*, vol. 102, no. 477, pp. 359–378, 2007.
- [34] N. Merhav and M. Feder, “Universal prediction,” *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2124–2147, 1998.
- [35] R. G. Gallager, *Information theory and reliable communication*. Springer, 1968, vol. 2.
- [36] M. Pinsker, *Information and information stability of random variables and processes*. Holden-Day, 1964.
- [37] R. Gray, *Entropy and information theory*. Springer, 1990.
- [38] G. Kramer, “Capacity results for the discrete memoryless network,” *IEEE Transactions on Information Theory*, vol. 49, no. 1, pp. 4–21, 2003.
- [39] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York, NY, USA: Wiley-Interscience, 1991.
- [40] T. Tanaka, P. M. Esfahani, and S. K. Mitter, “LQG control with minimum directed information: Semidefinite programming approach,” *arXiv preprint arXiv:1510.04214*, 2015.
- [41] T. Tanaka and H. Sandberg, “SDP-based joint sensor and controller design for information-regularized optimal LQG control,” *The 54th IEEE Conference on Decision and Control (CDC)*, 2015.