Optimal Sensor-Controller Codesign for Privacy in Dynamical Systems

Ruoxi Jia, Roy Dong, Shankar Sastry and Costas Spanos

Abstract— **We study the problem of jointly designing the sensor and controller for a dynamical system driven by a privacy-sensitive input process. This problem is motivated by the modern thermostat control example where home's occupancy is continually monitored and leveraged to tailor thermostat behaviors for better energy savings and comfort, which, however, arouses users' concern over privacy. We start by quantifying the instantaneous privacy loss in a control system under standard inference attacks. We present the closed form of privacy loss for linear Gaussian systems and propose a sampling-based method to approximate privacy loss for general dynamical systems. The optimal control and sensor query strategy for a private-input-driven system is then characterized, and we further prove the validity of separation principle for a linear system with Gaussian disturbance and quadratic cost under the privacy loss proposed in this paper. We close the paper by demonstrating the flexibility of the joint sensor-controller policy in the occupancy-based thermostat control example and providing some insights on the tradeoff among energy, comfort, and privacy.**

I. Introduction

Occupancy sensing forms a core aspect of the fabric of modern thermostats. Nest [1], as a popular example, exploits occupancy information derived from the built-in motion sensor or tracking users' smartphone location to automatically switch thermostat behaviors for increased energy saving and better thermal comfort. The collection of users' activity data naturally causes the concern about privacy. Currently, the thermostat companies rely on the encryption of data transmission to guard against malicious intercept of private data, and the institution of privacy policies to provide users with "notice and choice" [2]. However, neither of these measures will prevent users' private data from being eavesdropped by an insider who has legitimate access to the data.

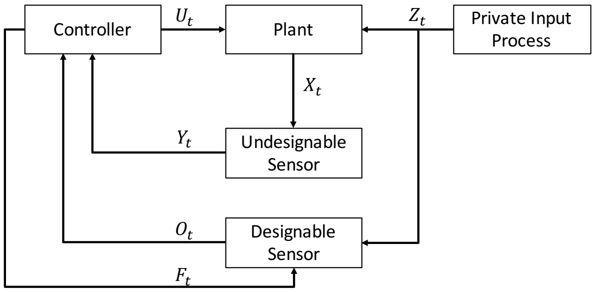

Fig. 1. Private-input-driven system diagram.

In recent years, privacy-differentiated goods have surfaced as a way to address the tension between privacy requirements and data utility expectation of users. For instance, AT&T has announced a price model that allows users to pay for an opt-out from the default setting where their web browsing activities are wiretapped and utilized for targeted advertisement; Telematics offers a more cost-effective car insurance plan to users who are willing to share their driving data. Inspired by these applications, we prototype a privacy-differentiated occupancy sensing module that adjusts the precision of occupancy data revealed to the controller in response to users' priority of energy saving, thermal comfort, and privacy.

Building upon the example of occupancy-based thermostat control, we consider a more general formulation which is referred to as private-input-driven system shown in Fig. 1. The system contains a local plant whose dynamics are influenced by a private input process. In the thermostat control example, the private input process may represent users' presence. The control of the plant is accomplished by a controller held by an honest but curious third party. The controller acquires the knowledge about the plant state and issues the control demands based on the noisy observations from one sensor that measures the plant state and another that measures the private input process. To accommodate the privacy needs, the controller supplies another control signal to the sensor that directly monitors the private input process, and tailors the sensor configuration according to users' privacy preference.

A quantification scheme of privacy leakage is essential to modeling and analysis of privacy-preserving systems. Two popular privacy metrics have been explored in the context of dynamical systems. One is differential privacy [3], which characterizes the change in the probability distribution of observables of a system by adding any single user's data; the other is information theoretic privacy [4], [5], which captures the reduction of adversary's uncertainty about private data after observing publicly available measurements.

Differential privacy relies on hiding an individual's private parameters in an aggregate of several different users, often known as "hiding in the crowd" [6]. It also has a one-size-fits-all privacy model, where the overall system provides the

This research is funded by the Republic of Singapore's National Research Foundation through a grant to the Berkeley Education Alliance for Research in Singapore (BEARS) for the Singapore-Berkeley Building Efficiency and Sustainability in the Tropics (SinBerBEST) Program. BEARS has been established by the University of California, Berkeley as a center for intellectual excellence in research and education in Singapore. This research is also supported in part by FORCES (Foundations Of Resilient CybEr-Physical Systems), which receives support from the National Science Foundation (NSF award numbers CNS-1238959, CNS-1238962, CNS-1239054, CNS-1239166).

R. Jia, R. Dong, S. Sastry and C. Spanos are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley ruoxijia@berkeley.edu, roydong@eecs.berkeley.edu, sastry@eecs.berkeley.edu,spanos@berkeley.edu

same level of privacy to all users. As we are interested in customizing the privacy loss for individual users instead of database-level privacy, information theoretic privacy is more suitable to our context. However, current information theoretic measures of time series data [7], [8] are not amendable to the optimal control framework where the instantaneous privacy leakage from sequentially generated data needs to be properly characterized. Mutual information or directed information between private time series and public measurements are used in [7], [8] as privacy metrics; however, these metrics essentially compute the expected privacy leakage over publicly measured random variables, and do not take into account the fact that at time t the public measurements before t have already been realized rather than remaining random.

The contributions of this paper are three-fold. Firstly, we propose a measure for instantaneous privacy loss that captures the privacy leakage of time series data on-the-fly. An analytic form of the privacy loss is derived for a linear Gaussian (LG) systems. We also develop an approximation framework for computing the privacy loss for more general dynamical systems. Secondly, we formulate the problem of finding optimal plant and measurement control policies for the private-input-driven systems using the framework of Partially Observed Markov Decision Processes (POMDPs), and give the optimal policies via dynamical programming. We further prove the separation of measurement control, plant control, and state estimation for LG systems with quadratic plant control cost. Thirdly, we leverage the techniques developed to investigate the trade-off between privacy, comfort and energy performance for occupancy-based thermostat control.

The problem studied in this paper has a close connection to some research work conducted in the control and information theory communities. A popular topic in control studies the case where the information used for controller's decision making consumes resources that need to be explicitly dealt with, such as bandwidth, power, or delay incurred by potential network traffic. A line of research [9] focuses on control under channel capacity constrained channels, which also results in an information theoretic constraint similar to the privacy loss proposed herein. However, it is often assumed that the sensor model and/or channel model is given a priori. This is different from our objective which is to jointly design the controller and sensor/channel. Another line of research, so-called "optimal sensing", coincides with our objective. The seminal work in [10] considers an adaptive measurement system where control is available for both the plant and measurement subsystem, and proves that the optimization of plant control can be carried out independently of the measurement control optimization for linear Gaussian system with quadratic costs. More recent work has focused on optimal measurement control [11] in which there are a set of sensors with different levels of precision and operation costs and the controller can access one sensor at a time to receive observation. The main results provided in [10], [11] assume the cost is only a function of the state and control actions. However, in our case, the privacy cost depends also on the sensor observations. [12] adopts a similar information constraint to the proposed privacy loss in this paper, and studies the problem of joint sensor and controller design. Our work differs from [12] in that we only consider a partial state to be private which is modeled as private input process. In addition, previous work has largely focused on LGQ control problems as it leads to separation of control and estimation and thereby an analytic solution. Our paper tries to tackle the general dynamical systems and provides an approximation framework to compute the optimal control policy. Our paper is also partially inspired by [13] which considers the state of a system as being private and studies the privacy-aware control in a POMDP setting. However, the privacy is protected by randomizing the control action instead of adding random disturbance to sensory data as in this paper.

II. General Problem Formulation

Notation. Throughout the paper, we will use capital letters to indicate random variables and lower case ones to refer to the realizations of the corresponding random variables. $X_{0:t}$ is used as a shorthand for $\{X_0, \cdots, X_t\}$.

Consider the discrete-time control problem depicted in Fig. 1. The dynamics of the plant are described by

$$X_{t+1} = f_t^X(X_t, Z_t, U_t, W_t^X) \tag{1}$$

where f_t^X is some function of possibly nonlinear form, $\{X_t\}$ is a plant state process, $\{U_t\}$ is a plant control process, $\{W_t^X\}$ is an independent and identically distributed (i.i.d.) disturbance. Additionally, the dynamics of the plant are steered by an exogenous input process $\{Z_t\}$, which represents personal attributes or behaviors and is thus *privacy-sensitive*. $\{Z_t\}$ is assumed to evolve according to the known dynamics given by

$$Z_{t+1} = f_t^Z(Z_t, W_t^Z) \tag{2}$$

where W_t^Z are i.i.d. random disturbance of private input process $\{Z_t\}$. The state of the plant is observed via an *undesignable* noisy sensor

$$Y_t = h_t^X(X_t, V_t^X) \tag{3}$$

where h_t^X represents the measurement model for the plant, and $\{V_t^X\}$ is a i.i.d measurement noise process. To cater for the needs of privacy, the measurement of the private input process $\{Z_t\}$ is regulated by a *designable* sensor. The control over the designable sensor is denoted by $\{F_t\}$, and the noisy observation for private input is thus given by

$$O_t = h_t^Z(Z_t, F_{t-1}, V_t^Z) \tag{4}$$

where h_t^Z represents the measurement model for the private input, and $\{V_t^Z\}$ is the i.i.d. measurement noise.

The controller can get access to noisy measurements for the plant and private input and its past control actions. We denote the information available to the controller at time t by $I_t = (Y_{0:t}, O_{0:t}, U_{0:t-1}, F_{0:t-1})$. The controls are assumed to be a deterministic function of the available information. Let the control policies for the plant and designable measurement system be μ and α, respectively. Then we have $U_k = \mu(I_k)$ and $F_k = \alpha(I_k)$.

Our aim is to find the optimal balance between the privacy loss due to measurements and the savings in plant operation costs made possible by the measurements. Suppose the control horizon is T, the plant control cost associated with policy $\eta = (\mu, \alpha)$ is

$$C^\eta = \sum_{t=0}^{T-1} c_t(X_t^\eta, Z_t^\eta, U_t^\eta) + c_T(X_T^\eta, Z_T^\eta) \tag{5}$$

where the superscript η indicates that the corresponding random processes are induced by policy η. c_t represents the instantaneous cost at time t. Let the total privacy loss during horizon T be denoted by G^η, our objective is to solve

$$\inf_\eta \mathbb{E}[(1-\gamma)C^\eta + \gamma G^\eta] \tag{6}$$

for a desired weight γ. As γ is increased from 0 to 1, the controller gradually changes from being completely utility driven to one prioritizing privacy.

We note the closed-loop system under control law $\eta = (\mu, \alpha)$ determines the information available I_t^η as well as the dynamics and costs. In the following sections, when we are analyzing the performance of a fixed controller η, or considering the effects of particular control actions (u_t, f_t), we will suppress the η or (u_t, f_t) dependence for notational cleanliness when the context is clear.

III. Inferential privacy

In this section, we present a statistical inference view to capture the privacy threat incurred by releasing streaming noisy measurements to achieve certain control objectives.

A. Threat model

We start by defining the threat model. An adversary observes I_T by constantly eavesdropping the communication link between the controller and the local infrastructure including the plant and sensors. The adversary is interested in inferring the private input process $Z_{0:T}$ from I_T. To formalize the privacy threat model, we assume the standard statistical inference threat model in [14].

Definition 1: (Inference attack). An inference attack on $Z_{0:T}$ given the observation I_T takes as input the joint distribution $p(Z_{0:T}, I_T)$ and the observation $I_T = i_T$, and outputs a probability distribution q^* defined over the domain of $Z_{0:T}$, as the solution to the minimization

$$q^* = \arg\min_q \sum_{z_{0:T}} p(z_{0:T}|i_T)\Psi(z_{0:T}, q) \tag{7}$$

for some cost function $\Psi(z_{0:T}, q)$.

Remark 1: $\Psi(z_{0:T}, q)$ is a generic notation for inference loss function. E.g., $\Psi(z_{0:T}, q) = -\log q(z_{0:T}|i_T)$ if the logarithm loss is used. The inference attack generates a belief distribution q over the private process $Z_{0:T}$ given the observation I_T by minimizing the expected inference loss.

We proceed to define the privacy leakage as follows: Without observing i_T, the adversary's belief about the private process can be captured by the solution of the minimization

$$\Psi_0^* = \min_q \sum_{z_{0:T}} p(z_{0:T})\Psi(z_{0:T}, q) \tag{8}$$

After observing i_T, the adversary would update the belief q such that it minimizes

$$\Psi_{i_T}^* = \min_q \sum_{z_{0:T}} p(z_{0:T}|i_T)\Psi(z_{0:T}, q) \tag{9}$$

Note that $\Psi_{i_T}^*$ is determined by the realization i_T, whereas we wish to consider the average privacy loss across the distribution of I_T. In order to quantify how much an adversary gains in terms of inference of the private process $Z_{0:T}$ by virtue of observing I_T, we consider the average cost gain

$$\Delta\Psi = \Psi_0^* - \sum_{i_T} p(i_T)\Psi_{i_T}^* \tag{10}$$

Definition 2: (Expected total privacy loss). The expected total privacy loss of $Z_{0:T}$ from the observation I_T is given by $\Delta\Psi$.

The log-loss is used as the loss function in recent work, e.g., [4], to compute the privacy loss, as it results in a natural measure of relevance - mutual information.

Proposition 1: If an adversary uses the log-loss $\Psi(Z_{0:T}, q) = -\log q(Z_{0:T})$, the total expected privacy loss of $Z_{0:T}$ from I_T is equivalent to the mutual information between $Z_{0:T}$ and I_T, i.e.,

$$\Delta\Psi = \mathbb{I}(Z_{0:T}; I_T) \tag{11}$$

and $\mathbb{I}(Z_{0:T}; I_T) \triangleq \mathbb{H}(Z_{0:T}) - \mathbb{H}(Z_{0:T}|I_T)$ represents the mutual information between the two sequences $Z_{0:T}$ and I_T.

In this paper, we will also use mutual information as a measure of privacy loss. Mutual information is shown to be the *only* measure that satisfies the data processing axiom which is needed to properly define the benefit of side information in an inference problem [15]. In addition, mutual information is closely related to the probability of success of an inference algorithm used by the adversary via Fano's Inequality. Let $\hat{Z}_{0:T}$ be the adversary's inference based on the observation I_T. Assuming Z_t has a finite alphabet size denoted by $|Z|$, Fano's Inequality states a lower bound on the probability of inference error

$$P(Z_{0:T} \neq \hat{Z}_{0:T}) \geq \frac{\mathbb{H}(Z_{0:T}) - \mathbb{I}(Z_{0:T}; I_T) - 1}{T \log |Z|} \tag{12}$$

It is clear from (12) that the bound on the probability of error is maximized when $\mathbb{I}(Z_{0:T}; I_T) = 0$, i.e., the two sequences are independent.

B. Instantaneous privacy loss

Solutions to classical POMDPs reduce the full horizon policy optimization to the Bellman equation that finds the optimal control at each single time step iteratively. A key element in the reduction is the additive form of the classical cost function definitions. To benefit from the Bellman-type reduction, the following lemma presents a stepwise decomposition of total expected privacy loss.

Lemma 1: For the dynamical model in (1) and (2) and the measurement model in (3) and (4), assuming η is deterministic, $\mathbb{I}(Z_{0:T}; I_T)$ can be decomposed into the sum of

$$\mathbb{I}(Z_{0:T}; I_T) = \sum_{t=0}^{T} \mathbb{I}(Z_{0:t}; Y_t, O_t | I_{t-1}) \tag{13}$$

Proof: Note that $I_T = (I_{T-1}, Y_T, O_T, U_{T-1}, F_{T-1})$, hence

$$\begin{aligned}
\mathbb{I}(Z_{0:T}; I_T) &= \mathbb{I}(Z_{0:T}; I_{T-1}, Y_T, O_T, U_{T-1}, F_{T-1}) && (14)\\
&= \mathbb{I}(Z_{0:T-1}; I_{T-1}) + \mathbb{I}(Z_T; I_{T-1}|Z_{0:T-1})\\
&\quad + \mathbb{I}(Z_{0:T}; Y_T, O_T, U_{T-1}, F_{T-1}|I_{T-1}) && (15)\\
&= \mathbb{I}(Z_{0:T-1}; I_{T-1})\\
&\quad + \mathbb{I}(Z_{0:T}; Y_T, O_T, U_{T-1}, F_{T-1}|I_{T-1}) && (16)\\
&= \mathbb{I}(Z_{0:T-1}; I_{T-1}) + \mathbb{I}(Z_{0:T}; Y_T, O_T|I_{T-1}) && (17)
\end{aligned}$$

where (15) is by applying the chain rule for mutual information, (16) follows from the fact that Z_T is conditionally independent of I_{T-1} given $Z_{0:T-1}$, and (17) follows from the presumption that U_{T-1} and F_{T-1} are a deterministic function of I_{T-1}. We can recursively break down $\mathbb{I}(Z_{0:T-1}; I_{T-1})$ in a similar fashion to obtain (13). ■

Definition 3: (Instantaneous privacy loss). The instantaneous privacy loss at time t due to the observation $I_t = i_t$ is given by

$$g_t(i_t, f_t, u_t) = \mathbb{I}(Z_{0:t+1}; Y_{t+1}, O_{t+1}|I_t = i_t) \tag{18}$$

Remark 2: It is obvious that the RHS of the equation above is a function of i_t. The reason the RHS also depends on f_t and u_t is that the distribution of Y_{t+1} and O_{t+1} hinges on the control actions exerted from time 0 to t, while the dependence on $\{u_i, f_i\}_{i=0}^{t-1}$ are implicitly encapsulated into i_t. G^η in (6) is thus defined as $\sum_{t=0}^{T-1} g_t(i_t, f_t, u_t)$.

We can represent the instantaneous privacy loss in the form of entropy difference:

$$\mathbb{H}(Z_{0:t+1}|I_t = i_t) - \mathbb{H}(Z_{0:t+1}|Y_{t+1}, O_{t+1}, I_t = i_t) \tag{19}$$

from which we can obtain an intuitive interpretation of the instantaneous privacy loss at time t - it indicates the anticipated reduction in adversary's uncertainty about the private process up to time $t+1$ due to the upcoming measurements given all the information received up to time t.

C. Special case: linear Gaussian system

For a general dynamical system stated in Section II, the instantaneous privacy loss is a function of the information available to the controller as well as the plant and measurement control policy. However, if the plant and the measurement systems are linear, and the disturbance and measurement noise are Gaussian, then the instantaneous privacy loss at each time step is independent of plant control policy.

Suppose $X_t \in \mathbb{R}^{n_X}$, $Z_t \in \mathbb{R}^{n_Z}$, $U_t \in \mathbb{R}^{n_U}$, $Y_t \in \mathbb{R}^{n_Y}$, $O_t \in \mathbb{R}^{n_O}$, and the LG system is described by the following equations:

$$\begin{aligned}
X_{t+1} &= A_t^X X_t + A_t^{XZ} Z_t + B_t^X U_t + W_t^X && (20)\\
Z_{t+1} &= A_t^Z Z_t + W_t^Z && (21)\\
Y_t &= C_t^X X_t + V_t^X && (22)\\
O_t &= C_t^Z Z_t + V_t^Z && (23)
\end{aligned}$$

where $W_t^X \in \mathbb{R}^{n_X}$, $W_t^Z \in \mathbb{R}^{n_Z}$, $V_t^X \in \mathbb{R}^{n_Y}$ and $V_t^Z \in \mathbb{R}^{n_O}$ are zero-mean Gaussian random variables with covariance M_t^X, M_t^Z, N_t^X and N_t^Z, respectively. N_t^Z and C_t^Z are determined by $F_{t-1} \in \mathbb{R}^{n_F}$.

We introduce the following notations. Let

$$\begin{aligned}
&A_t = \begin{bmatrix} A_t^X & A_t^{XZ} \\ \mathbf{0}_{n_Z \times n_X} & A_t^Z \end{bmatrix} B_t = \begin{bmatrix} B_t^X \\ \mathbf{0}_{n_Z \times n_U} \end{bmatrix} && (24)\\
&C_t = \begin{bmatrix} C_t^X & C_t^Z \end{bmatrix} && (25)\\
&M_t = \begin{bmatrix} M_t^X & \mathbf{0}_{n_X \times n_X} \\ \mathbf{0}_{n_Z \times n_Z} & M_t^Z \end{bmatrix} N_t = \begin{bmatrix} N_t^X & \mathbf{0}_{n_Y \times n_Y} \\ \mathbf{0}_{n_O \times n_O} & N_t^Z \end{bmatrix} && (26)
\end{aligned}$$

where $\mathbf{0}_{\cdot\times\cdot}$ stands for a zero matrix and the associated subscripts represent its dimension.

Proposition 2: The instantaneous privacy loss induced by measurement control policy α at time t of the LQ system presented in (20)-(23) is given by

$$g_{t,lin} = \frac{1}{2} \log \frac{|\Sigma_{t+1|t,sub}|}{|\Sigma_{t+1|t+1,sub}|} \tag{27}$$

where $\Sigma_{t+1|t,sub}, \Sigma_{t+1|t+1,sub} \in \mathbb{R}^{m \times m}$ are the lower-right submatrices of $\Sigma_{t+1|t}$ and $\Sigma_{t+1|t+1}$, respectively, which are the error covariance matrices in Kalman filter and can be iteratively computed from

$$\begin{aligned}
&\Sigma_{t+1|t} = A_t \Sigma_{t|t} A_t' + M_t && (28)\\
&\Sigma_{t+1|t+1} = \Sigma_{t+1|t}\\
&- \Sigma_{t+1|t} C_{t+1}' (C_{t+1} \Sigma_{t+1|t} C_{t+1}' + N_{t+1})^{-1} C_{t+1} \Sigma_{t+1|t} && (29)
\end{aligned}$$

Remark 3: The proof of the proposition is provided in Appendix I. From (27) we can see that $g_{t,lin}$ only depends on the measurement control signal $(f_0, \cdots, f_t)$. This is because the recursive structure of $\Sigma_{t+1|t+1}$ makes it a function of $N_0, \cdots, N_{t+1}$ and $C_0, \cdots, C_{t+1}$, which, in turn, hinge on $(f_0, \cdots, f_t)$. An important consequence of this result, as we will prove in the later section, is that the measurement control policy can be designed separately from the plant control policy.

D. Privacy loss approximation for general control systems

Unlike LG systems, a general dynamical system does not enjoy an analytic form of privacy loss. Direct evaluation of the exact privacy loss is computationally intractable as it involves calculating the joint probabilities of a sequence of random variables whose size grows exponentially with the control horizon. Rather than computing the exact privacy loss, we propose to approximate it using Gibbs sampling and “plug-in” estimators, and thereby avoid operating on exponentially many state patterns at some cost in accuracy.

More specifically, given the information i_t at time t, we first estimate the following probabilities distributions: $P(Z_{0:t+1}|i_t)$, $P(Z_{0:t+1}|Y_{t+1}, O_{t+1}, i_t)$, and $P(Y_{t+1}, O_{t+1}|i_t)$ by sampling, and then plug the estimates into (19) to obtain a “plug-in” estimator of $g_t(i_t, u_t, a_t)$. To acquire the samples from the aforementioned probabilities, we consider two probabilities $P(X_{0:t}, Z_{0:t+1}|i_t)$ and $P(X_{0:t+1}, Z_{0:t+1}, Y_{t+1}, O_{t+1}|i_t)$, which can be efficiently sampled via Gibbs sampling due to the Markovian structure of the problem. For detailed implementation procedure of Gibbs sampling, we refer the readers to [16].

Note that in contrast to the second probability, the first one does not depend on the control action (u_t, f_t). So the samples from the first one can be used to approximate $P(Z_{0:t+1}|i_t)$, which is also not dependent on the current control action, by simply considering the samples of $Z_{0:t+1}$ and ignoring the rest. Similarly, $P(Z_{0:t+1}|y_{t+1}, o_{t+1}, i_t)$, and $P(y_{t+1}, o_{t+1}|i_t)$ can be estimated by the samples from $P(X_{0:t+1}, Z_{0:t+1}, Y_{t+1}, O_{t+1}|i_t)$.

IV. Characterization of optimal policy

The optimal policy for the general problem formulation stated in Section II with the proposed instantaneous privacy loss is presented as follows.

Proposition 3: If (18) is used as the privacy loss at time t, then the optimal plant and measurement control policy η in (6) for the system described by Equation (1)-(4) is obtained by minimizing the right-hand side of the following Bellman equations (if exists):

$$J_t(i_t) = \min_{u_t, f_t} g_t(i_t, u_t, f_t) + \mathbb{E}\bigg[c_t(X_t, Z_t, U_t) \\ + J_{t+1}(i_t, Y_{t+1}, O_{t+1}, u_t, f_t)|i_t, u_t, f_t\bigg] \quad (30)$$

for $t = 1, \cdots, T-1$, and

$$J_T(i_T) = \mathbb{E}\bigg[c_T(X_T, Z_T)|i_T\bigg] \quad (31)$$

Proof: This can be shown by treating (u_t, f_t) as a combined control action of the dynamical system with the state (X_t, Z_t) and then applying dynamic programming. ■

Proposition 3 indicates that the solutions of optimal plant control and measurement control are tightly coupled for a general dynamical system as the privacy loss is determined by both plant and measurement controls. The coupling of solutions also arises from the interaction between control and estimation. The measurement control signal affects knowledge of states through noisy observations, which, in turn, affect control actions exerted to the plant.

By combining the result that the instantaneous privacy loss is only a function of measurement control signals shown in Proposition 2 and the famous control-estimation separation theorem for LG systems with quadratic costs, we conjecture that the plant and measurement control policies can be solved separately in LGQ systems. The following proposition states and proves the conjecture above in a rigorous manner. In addition, it shows that the optimal measurement control signal can be determined a priori, regardless of the measurement received on-the-fly. The proof of the following proposition is provided in Appendix II.

TABLE I

Parameter	Meaning	Value & Units
R	Average thermal resistance	$2°C/kW$
C	Average thermal capacitance	$10kWh/°C$
P	Average energy transfer rate	$14kW$
C_p	Average thermal load per person	$0.01°C/h$
X_a	Ambient temperature	$32°C$
η	Load efficiency	2.5
σ_{model}	Noise std of temperature model	$10^{-5}°Cs^{-0.5}$
σ_{meas}	Noise std of temperature measurement	$0.5°C$
X_d	Desired temperature	$25°C$

Proposition 4: Consider the LQ system described by (20)-(23). Defining $S_t = [X_t; Z_t]$, the optimal plant and measurement controls u_t^* and f_t^* that minimize the expected value of the sum of the quadratic operation cost and the privacy loss

$$\mathbb{E}[S_N' Q_N S_N + \sum_{t=0}^{T-1} (S_t' Q_t S_t + U_t' R_t U_t + \gamma g_{t,lin})] \quad (32)$$

are given by

$$u_t^* = L_t \mathbb{E}[S_t|i_t], t = 0, \cdots, T-1 \quad (33)$$

where L_t is defined by

$$L_t = -(R_t + B_t' K_{t+1} B_t)^{-1} B_t' K_{t+1} A_t \quad (34)$$

with the matrices K_t given recersively by the Riccati equation

$$K_T = Q_T \quad (35)$$

$$K_t = Q_t + A_t' K_{t+1} A_t - P_t \quad (36)$$

$$P_t = A_t' K_{t+1} B_t (R_t + B_t' K_{t+1} B_t)^{-1} B_t' K_{t+1} A_t \quad (37)$$

and f_t^* ($t = 0, \cdots, T-1$) that solve the following deterministic optimization problem

$$\min_{\substack{f_t \\ t=1,\cdots,T-1}} \sum_{t=1}^{T-2} Tr\big(P_{t+1}\Sigma_{t+1|t+1}\big) + \sum_{t=1}^{T-1} \gamma g_{t,lin} \quad (38)$$

V. Case study: Occupancy-based Thermostat Control

In this section, we design a privacy-preserving control law for a thermostat that utilizes occupancy information to achieve energy savings and comfort improvement.

A. Thermostat model

The thermostat model presented herein follows [17] closely. With reference to the notations in Table I, the temperature dynamics are modeled by a linear equation

$$X_{t+1} = aX_t + (1-a)(X_a - U_t RP) + C_p h Z_t + W_t^X \quad (39)$$

where $a = \exp(-h/CR)$ governs the thermal characteristics of the thermal mass and h is the time elapsed per discrete time step. $U_t \in \{0, 1\}$ indicates the ON/OFF actions. W_t^X is a zero-mean Gaussian noise process with variance $h\sigma^2_{model}$, accounting for all heat gain and loss not modeled explicitly. The term $C_p h Z_t$ captures the heat contributed by human presence, where Z_t indicates the number of people. For simplicity, we consider Z_t to be binary, which evolves as a Markov chain with transition probability

$$P(Z_{t+1}|Z_t) = \begin{cases} 1-q & \text{if } Z_{t+1} \neq Z_t \\ q & \text{if } Z_{t+1} = Z_t \end{cases} \tag{40}$$

The temperature is measured via a noisy sensor given by

$$Y_t = X_t + V_t^X \tag{41}$$

where $V_t^X \sim \mathcal{N}(0, \sigma^2_{meas})$. To preserve privacy, the occupancy data Z_t will be obfuscated randomly in situ before being sent to the controller. To be specific, the occupancy sensor measurement O_t is modeled by

$$P(O_t|Z_t) = \begin{cases} 1-F_{t-1} & \text{if } O_t \neq Z_t \\ F_{t-1} & \text{if } O_t = Z_t \end{cases} \tag{42}$$

where $F_t \in \{0.5, 1\}$ is the control action on the occupancy sensor. U_t and F_t are designed by minimizing the expectation of the following objective

$$\sum_{t=0}^{T-1} \left[\underbrace{\frac{1}{\eta} P h U_t}_{\text{energy cost}} + \gamma_p \underbrace{g_t(I_t, U_t, F_t)}_{\text{privacy loss}} + \gamma_c \underbrace{Z_t (X_t - X_d)^2}_{\text{comfort loss}} \right] \tag{43}$$

where γ_p and γ_c stand for the amount of energy people are willing to pay in exchange of one bit of private information leakage, and one unit of uncomfortableness measured in the squared deviation of current temperature from the desired temperature X_d, respectively. Z_t is included in the comfort loss term to accommodate the fact that people only sustain comfort loss when they are present.

The optimal policy of the thermostat system presented above is intractable. To obtain a simple suboptimal controller, we resort to open-loop feedback control [18] and a heuristic that calibrates the estimate of future control cost by taking into account the effect of occupancy noise on the estimation of comfort loss.

B. Results

Since the thermostat model is not linear Gaussian, we use the sampling-based method presented in Section III-D to approximate the instantaneous privacy loss at each time step, where sample size is a free parameter that manifests the tension between computational efficiency and approximation accuracy. Fig. 2 provides a reference for the selection of the sample size. For each sample size and horizon length, we conduct 100 Monte Carlo simulations of privacy loss for each one of the 100 randomly generated action traces. The vertical axis of Fig. 2 is obtained by first calculating the standard deviation of privacy loss over all Monte Carlo simulations for every single action trace, and then computing the average of different action traces' privacy loss variation. It can be seen that a larger sample size can reduce the variation of the sampling-based privacy loss approximation. Moreover, for a given variation tolerance, more samples are needed for a longer horizon. In the following experiments, we choose the sample size to be 5000.

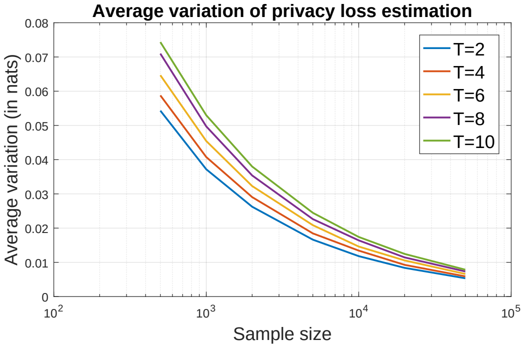

Fig. 2. The variation of privacy loss approximation for different sample sizes and horizon lengths.

We then demonstrate the flexibility of jointly optimized control over plant and occupancy sensing fidelity, by comparing it with two other policies: (1) a minimum-privacy policy that reports the ground truth occupancy at all times (i.e., $F_t = 1, \forall t$) and optimizes the plant control action at each time step; (2) a maximum-privacy policy that always flips the true occupancy with probability 0.5 (i.e., $F_t = 0.5, \forall t$) and optimizes the plant control. The costs incurred by the three policies are simulated for users with different privacy preferences. We consider three different types of users, namely, the privacy fundamentalist who is very concerned about privacy, the pragmatist that values the privacy to a moderate degree, and the unconcerned who does not care about privacy loss. We fix the comfort weight $\gamma_c = 1$ and varies the privacy weight $\gamma_p = 5, 1, 0$ for the three type of users. We cluster simulations into two scenarios, namely mostly occupied or mostly empty, according to the ground truth occupancy of the space, and examine how the operation costs and privacy loss of different policies change for different types of users.

Figure 3a and 3b illustrate the superiority of "optimally" flipped occupancy reports in terms of the total cost. Reporting the true occupancy tends to be the best strategy for users who are unconcerned about privacy loss while reporting randomly flipped occupancy is better if the user is a privacy fundamentalist. Figure 3c and 3d show an interesting occupancy bias introduced by random occupancy flipping. It can be seen that when the space is mostly occupied the controller with randomly flipped occupancy ($F_t = 0.5$) incurs lower energy cost compared with the one that receives true occupancy reports; to the contrary, when the space is mostly empty the energy cost with randomly flipping occupancy is higher. This is because random flipping introduces artifacts of low (or high) occupancy when the space is actually occupied (or empty). Similar occupancy bias also manifests itself in Fig. 3e and 3f by driving the comfort loss up when the space is occupied and bringing

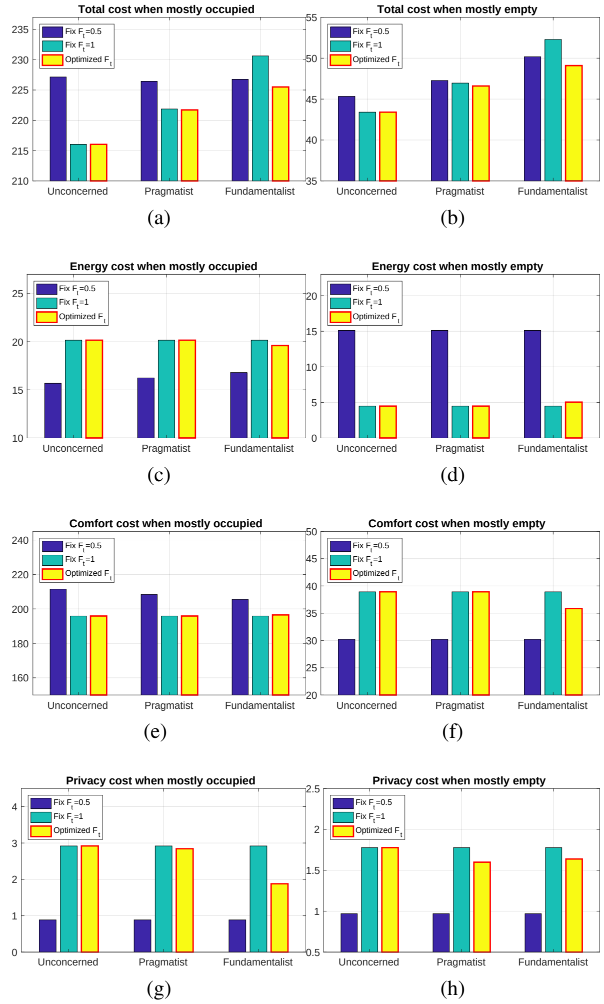

Fig. 3. The total cost, energy cost, comfort cost, and privacy cost incurred by three different policies: (1) flipping the occupancy with probability 0.5 and optimizing the plant controls (maximum privacy), (2) always reporting the true occupancy and optimizing the plant controls (minimum privacy), and (3) jointly optimizing the occupancy flipping probability and plant controls for users with different privacy preferences (optimized privacy), under different space occupation states.

it down when the space is empty. The optimized flipping policy gradually reduces the privacy loss as the user puts more weight on the privacy loss, as shown in Fig. 3g and 3h while the other two policies lack the agility to tune the privacy loss in response to users' preference.

VI. Conclusion and Future Work

This paper studies the joint sensor-controller design problem in private-input-driven dynamical systems, which are motivated by occupancy-based thermostat control applications. We propose an instantaneous privacy loss measure that characterizes the privacy leakage of sequential data streams under standard inference attack, and provide a sampling-based method to approximate privacy loss of general systems with a possibly nonlinear form and arbitrary disturbance. We also characterize the optimal joint design policy and prove the separation of sensor design, plant control and system state estimation in LGQ systems. The numeric example on thermostat control demonstrates the tradeoff between plant operation cost and privacy leakage. The privacy-aware thermostat controller presented in the paper is shown to realize better privacy protection while maintaining energy efficiency and thermal comfort.

For future work, it will be meaningful to develop a suboptimal controller with lower computational overhead, especially when the control actions are continuous. We also intend to implement the privacy-aware thermostat control strategy in a real-world testbed.

Appendix I
Proof of Proposition 2

Proof: Since

$$Cov(X_{t+1}, Z_{t+1}|I_t = i_t) = \Sigma_{t+1|t} \tag{44}$$
$$Cov(X_{t+1}, Z_{t+1}|I_{t+1} = i_{t+1}) = \Sigma_{t+1|t+1} \tag{45}$$

and the fact that $p(X_{t+1}, Z_{t+1}|I_t = i_t)$ and $p(X_{t+1}, Z_{t+1}|I_{t+1} = i_{t+1})$ are Gaussian distributions, we have

$$Cov(Z_{t+1}|I_t = i_t) = \Sigma_{t+1|t,sub} \tag{46}$$
$$Cov(Z_{t+1}|I_{t+1} = i_{t+1}) = \Sigma_{t+1|t+1,sub} \tag{47}$$

It follows that

$$H(Z_{t+1}|I_t = i_t) = \frac{1}{2}\log(2\pi e)^{nz}|\Sigma_{t+1|t,sub}| \tag{48}$$
$$\begin{aligned} &H(Z_{t+1}|Y_{t+1}, O_{t+1}, I_t = i_t) \\ &= \sum_{y_{t+1}, o_{t+1}} p(y_{t+1}, o_{t+1}|i_t) H(Z_{t+1}|i_{t+1}) \\ &= H(Z_{t+1}|i_{t+1}) \\ &= \frac{1}{2}\log(2\pi e)^{nz}|\Sigma_{t+1|t+1,sub}| \end{aligned} \tag{49}$$

where the second equality in (49) is because $\Sigma_{t+1|t+1,sub}$ and thereby $H(Z_{t+1}|i_{t+1})$ are not functions of the observations received, i.e., y_{t+1} and o_{t+1}. Then by (19) and simple algebraic manipulation, we complete the proof. ■

Appendix II
Proof of Proposition 4

Proof: Note that the optimization problem (38) is equivalent to the following deterministic dynamic programming:

$$G_{T-1}(\theta_{T-1}) = \min_{f_{T-1}} \gamma g_{T-1,lin} \tag{50}$$
$$G_t(\theta_t) = \min_{f_t} Tr(P_{t+1}\Sigma_{t+1|t+1}) + \gamma g_{t,lin} + G_{t+1}(\theta_t, f_t) \tag{51}$$

where $\theta_t = (\theta_{t-1}, f_{t-1})$. Both $g_{t,lin}$ and $\Sigma_{t+1|t+1}$ are a function of (θ_t, f_t). For the optimal measurement policy proof, we will prove the equivalent characterization (50) and (51).

We will prove the result by induction. Let $W_t = [W_t^X, W_t^Z]$. By (30) and (31) in Proposition 3, we have

$$\begin{aligned}
&J_{T-1}(i_{T-1}) = \mathbb{E}[S'_{T-1}(Q_{T-1} + A'_{T-1}Q_TA_{T-1})S_{T-1}|i_{T-1}] \\
&+ \mathbb{E}[W'_{T-1}Q_TW_{T-1}] + \min_{f_{T-1}} \gamma g_{T-1,lin} \\
&+ \min_{u_{T-1}} \{u'_{T-1}(R_{T-1} + B'_{T-1}Q_TB_{T-1})u_{T-1} \\
&+ 2\mathbb{E}[S_{T-1}|i_{T-1}]A'_{T-1}Q_TB_{T-1}u_{T-1}\} \qquad (52)
\end{aligned}$$

Hence, the optimal measurement control f^*_{T-1} is the solution of (50). By taking the derivative of the last two lines in (52) and setting to zero, we get $u^*_{T-1} = L_{T-1}\mathbb{E}[S_{T-1}|i_{T-1}]$.

Now, assume that the cost-to-go function at time $t+1$ takes the form

$$\begin{aligned}
&J_{t+1}(i_{t+1}) = \mathbb{E}[S'_{t+1}K_{t+1}S_{t+1}|i_{t+1}] \\
&+ \mathbb{E}[(S_{t+1} - \mathbb{E}[S_{t+1}|i_{t+1}])'P_{t+1}(S_{t+1} - \mathbb{E}[S_{t+1}|i_{t+1}])|i_{t+1}] \\
&+ \sum_{\tau=t+1}^{T-1} \mathbb{E}[W'_\tau K_{\tau+1}W_\tau] + G_{t+1}(\theta_{t+1}) \qquad (53)
\end{aligned}$$

By Proposition 3, the cost-to-go at time t is

$$\begin{aligned}
&J_t(i_t) = \mathbb{E}[S'_tQ_tS_t|i_t] + \mathbb{E}[S'_tA'_tK_tA_tS_t|i_t] \\
&+ \min_{u_t}\{u'_t(R_t + B'_tK_{t+1}B_t)u_t + \mathbb{E}[u'_tB'_tK_{t+1}A_tS_t|i_t, u_t]\} \\
&+ \min_{f_t}\{\gamma g_{t,lin} + G_{t+1}(\theta_t, f_t)\} \\
&+ \min_{u_t,f_t} \mathbb{E}[(S_{t+1} - \mathbb{E}[S_{t+1}|I_{t+1}])'P_{t+1}\cdot \\
&\qquad (S_{t+1} - \mathbb{E}[S_{t+1}|I_{t+1}])|i_t, u_t, f_t] \\
&+ \sum_{\tau=t}^{T-1} \mathbb{E}[W'_\tau K_{\tau+1}W_\tau] \qquad (54)
\end{aligned}$$

Since $\mathbb{E}[(S_{t+1} - \mathbb{E}[S_{t+1}|I_{t+1}])'P_{t+1}(S_{t+1} - \mathbb{E}[S_{t+1}|I_{t+1}])]$ is weighted error covariance produced by the Kalman filter, which does not depend the measurement received and plant control actions but rather on the plant and measurement parameters, it follows that

$$\begin{aligned}
&\min_{u_t,f_t} \mathbb{E}[(S_{t+1} - \mathbb{E}[S_{t+1}|I_{t+1}])'P_{t+1}\cdot \\
&\qquad (S_{t+1} - \mathbb{E}[S_{t+1}|I_{t+1}])|i_t, u_t, f_t] \\
&= \min_{f_t} Tr(P_{t+1}\Sigma_{t|t}) \qquad (55)
\end{aligned}$$

Therefore, f^*_t is given by the solution of the following optimization problem

$$\min_{f_t} Tr(P_{t+1}\Sigma_{t|t}) + \gamma g_{t,lin} + G_{t+1}(\theta_t, f_t) \qquad (56)$$

By setting the derivative of the second line in (54) to zero, we get the optimal plant control

$$u^*_t = -L_t\mathbb{E}[S_t|i_t] \qquad (57)$$

Plug f^*_t and u^*_t back to $J_t(i_t)$, we have

$$\begin{aligned}
&J_t(i_t) = \mathbb{E}[S'_tK_tS_t|i_t] + \mathbb{E}[(S_t - \mathbb{E}[S_t|i_t])'P_t(S_t - \mathbb{E}[S_t|i_t])|i_t] \\
&\quad + \sum_{i=t}^{T-1} \mathbb{E}[W'_iK_{i+1}W_i] + G_t(\theta_t) \qquad (58)
\end{aligned}$$

Herein, we have proved the correctness of the proposed form of the cost-to-go at each time step and obtained the expression of optimal plant and measurement control. ■

References

[1] https://nest.com/support/article/Learn-more-about-Home-Away-Assist.

[2] L. F. Cranor, "Necessary but not sufficient: Standardized mechanisms for privacy notice and choice," *J. on Telecomm. & High Tech. L.*, vol. 10, p. 273, 2012.

[3] J. Le Ny and G. J. Pappas, "Differentially private filtering," *IEEE Transactions on Automatic Control*, vol. 59, no. 2, pp. 341–354, 2014.

[4] F. du Pin Calmon and N. Fawaz, "Privacy against statistical inference," in *Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference on*. IEEE, 2012, pp. 1401–1408.

[5] R. Jia, R. Dong, S. S. Sastry, and C. J. Spanos, "Privacy-enhanced architecture for occupancy-based hvac control," *arXiv preprint arXiv:1607.03140*, 2016.

[6] C. Dwork, "Differential privacy," in *Proc. of the Int. Colloq. on Automata, Languages and Programming*. Springer, 2006, pp. 1–12.

[7] M. A. Erdogdu and N. Fawaz, "Privacy-utility trade-off under continual observation," in *Information Theory (ISIT), 2015 IEEE International Symposium on*. IEEE, 2015, pp. 1801–1805.

[8] L. Sankar, S. R. Rajagopalan, and S. Mohajer, "Smart meter privacy: A theoretical framework," *IEEE Transactions on Smart Grid*, vol. 4, no. 2, pp. 837–846, 2013.

[9] S. Tatikonda, A. Sahai, and S. Mitter, "Stochastic linear control over a communication channel," *IEEE transactions on Automatic Control*, vol. 49, no. 9, pp. 1549–1561, 2004.

[10] L. Meier, J. Peschon, and R. Dressler, "Optimal control of measurement subsystems," *IEEE Transactions on Automatic Control*, vol. 12, no. 5, pp. 528–536, 1967.

[11] W. Wu and A. Arapostathis, "Optimal sensor querying: General markovian and lqg models with controlled observations," *IEEE Transactions on Automatic Control*, vol. 53, no. 6, pp. 1392–1405, 2008.

[12] T. Tanaka and H. Sandberg, "Sdp-based joint sensor and controller design for information-regularized optimal lqg control," in *Decision and Control (CDC), 2015 IEEE 54th Annual Conference on*. IEEE, 2015, pp. 4486–4491.

[13] P. Venkitasubramaniam, J. Yao, and P. Pradhan, "Information-theoretic security in stochastic control systems," *Proceedings of the IEEE*, vol. 103, no. 10, pp. 1914–1931, 2015.

[14] S. Salamatian, A. Zhang, F. du Pin Calmon, S. Bhamidipati, N. Fawaz, B. Kveton, P. Oliveira, and N. Taft, "Managing your private and public data: Bringing down inference attacks against your privacy," *IEEE Journal of Selected Topics in Signal Processing*, vol. 9, no. 7, pp. 1240–1255, 2015.

[15] 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.

[16] C. M. Bishop, "Pattern recognition," *Machine Learning*, vol. 128, pp. 1–58, 2006.

[17] D. S. Callaway, "Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy," *Energy Conversion and Management*, vol. 50, no. 5, pp. 1389–1400, 2009.

[18] D. P. Bertsekas, D. P. Bertsekas, D. P. Bertsekas, and D. P. Bertsekas, *Dynamic programming and optimal control*. Athena Scientific Belmont, MA, 1995, vol. 1, no. 2.