Leveraging Unique CPS Properties to Design Better Privacy-Enhancing Algorithms

Jairo Giraldo
University of Texas at Dallas
800 W. Campbell Rd.
Richardson, Texas 75080
jairo.giraldo@utdallas.edu

Alvaro Cardenas
University of Texas at Dallas
800 W. Campbell Rd.
Richardson, Texas 75080
alvaro.cardenas@utdallas.edu

Murat Kantarcioglu
University of Texas at Dallas
800 W. Campbell Rd.
Richardson, Texas 75080
muratk@utdallas.edu

ABSTRACT

Cyber-Physical Systems (CPS) have unique properties that can be exploited to design new privacy-enhancing technologies that minimize the negative impact to the utility of CPS. In this paper we show two examples of these properties.

The first example looks at how differential privacy degrades CPS performance due to the large noise addition, and we then show how the inherent noise of CPS can be leveraged to reduce the additional noise added by differential privacy algorithms, and therefore, minimize the negative impact on the system utility and safety. In the second example we look at the ability to sample at sensor readings on demand, and how this flexibility can be used to design adaptive sensor sampling algorithms that hide sensitive information without the need to add noise.

CCS CONCEPTS

• Security and privacy → Privacy-preserving protocols; Intrusion detection systems;

ACM Reference format:

Jairo Giraldo, Alvaro Cardenas, and Murat Kantarcioglu. 2017. Leveraging Unique CPS Properties to Design Better Privacy-Enhancing Algorithms. In *Proceedings of HoTSoS, Hanover, MD, USA, April 04 - 05, 2017*, 12 pages. DOI: <http://dx.doi.org/10.1145/3055305.3055313>

1 INTRODUCTION

Cyber-Physical Systems (CPS) are modernizing our infrastructures by optimizing the use of resources in power grids, transportation systems, buildings, and other critical infrastructures [16]. To achieve their goal, CPS uses sensors to collect large-scale, fine-grained data from a variety of systems including human activities, posing serious privacy concerns. To protect the privacy of people whose activities are captured by CPS, we can apply tools developed for general information technology problems, like differential privacy. However, their direct application without leveraging the unique CPS setting will cause unnecessary negative impacts to the utility of CPS, and in some cases, it can lead to potentially unsafe operating conditions

(e.g. a smart grid system whose frequency deviates significantly from 60Hz).

To address these concerns, in this paper we show how privacy in cyber-physical systems present different challenges and opportunities compared to traditional database privacy. Due to the physical system dynamics, feedback control, and inherent noise, different tools from control theory can be extended to design novel privacy-enhancing technologies.

In particular, we focus on stochastic control systems and on how the inherent uncertainties of the system (e.g., sensor noise, and disturbances) can be used to ensure certain levels of privacy. In particular, we introduce a methodology to inject the minimum amount of differential-privacy noise to ensure our desired levels of privacy by leveraging the noise already available in the system. In the second part of the paper we show that for some systems, the addition of external noise can prevent them from operating safely, and therefore we propose a new definition of privacy focused on obscuring the presence or absence of events and show that we can meet that definition without adding noise. In this case, we show how a discretionary sampling, consensus-based control strategy enables us to hide information about specific events and ensure privacy with minimal costs to stability and convergence rates.

Privacy from a control-theory perspective has started to receive some attention in the last couple of years [2, 12]. Differential privacy has been a particularly popular tool proposed for estimation and control [4]. These previous papers have not considered how the inherent noise of stochastic control systems can reduce the amount of noise that needs to be added in differential privacy.

Our previous work considered the idea of minimizing data collection through the use of discretionary sampling [8]. In this paper we extend our previous work by showing proofs of convergence and stability. In addition we show how this discretionary sampling mechanism compares to differential privacy, and also show how adding noise can affect the safety of the system.

1.1 Preliminaries

We introduce some general definitions for Differential Privacy (DP) that are used in the rest of the paper.

Definition 1.1. Adjacency of two datasets [12]: A pair of numerical datasets $D_1, D_2 \subset \mathcal{D}^n$, for $\mathcal{D}^n$ the domain of the databases, and where $D_i = \{d_{i,1}, d_{i,2}, \dots\}$ is γ -adjacent $\gamma\text{-Adj}(D_1, D_2)$ if there exists a l such that $|d_{1,l} - d_{2,l}| \leq \gamma$ and $d_{1,k} = d_{2,k}$ for all $k \neq l$. In other words, they differ by at most γ in a single element. For streams of data, γ -adjacency is defined at each time step.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HoTSoS, Hanover, MD, USA

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM. 978-1-4503-5274-1/17/04...\$15.00

DOI: <http://dx.doi.org/10.1145/3055305.3055313>

Definition 1.2. Sensitivity [6]: Consider a multidimensional query function $q : \mathcal{D}^n \rightarrow \mathbb{R}^d$, where $\mathcal{D}^n$ is the domain of the databases and d is the dimension of the query output. For a given database $D \in \mathcal{D}^n$, the query output is

$$q(D) = (q_1(D), q_2(D), \dots, q_d(D))$$

for $q_i(D) \in \mathbb{R}$ for all i . The global sensitivity of q is

$$\Delta_{q,p} = \max_{Adj\{D_1, D_2\}} \|q(D_1) - q(D_2)\|_p$$

where p indicates the type of norm used.

Definition 1.3. Differential privacy [6]: Let $\eta \in \Omega$ be a random number drawn from the set Ω (it can be a finite or infinite set). A randomized mechanism $M : \mathcal{D}^n \times \Omega \rightarrow \mathcal{M}$ preserves (ϵ, δ) -differential privacy if for all adjacent datasets D_1, D_2 and for all subsets of possible answers $S \in \mathcal{M}$,

$$P(M(D_1) \in S) \leq e^\epsilon P(M(D_2) \in S) + \delta$$

where $\delta > 0$. ϵ can be considered as an upper bound on the amount of influence an individual's record has on the information released and δ is an approximation parameter that relaxes the privacy definition.

LEMMA 1.4. Gaussian Mechanism [6]: For a dataset D and a query q , a Gaussian mechanism $M = q(D) + \eta$ preserves (ϵ, δ) -differential privacy if η is drawn from a zero-mean Gaussian distribution with $\sigma \geq \sqrt{2 \ln(1.25/\delta)} \Delta_{q,2}/\epsilon$.

Graph theory: Let $\mathcal{G} = \{\mathcal{V}, \mathcal{E}, \mathcal{A}_{\mathcal{G}}\}$ represent an undirected graph, where $\mathcal{V} = \{1, \dots, N\}$ is the set of nodes or vertices, and $\mathcal{E} = \{(i, j) | i, j \in \mathcal{V}\}$ is the set of edges or adjacent nodes. The adjacency matrix $\mathcal{A}_{\mathcal{G}} = [a_{ij}]$ is the symmetric matrix $N \times N$, where $a_{ij} = 1$ if (i, j) are adjacent, $a_{ij} = 0$ otherwise, and $a_{ii} = 0$ for all $i \in \mathcal{V}$. For the i^{th} node, the degree of a vertex d_i is the number of neighbors that are adjacent to i , i.e., $d_i = \sum_{j=1}^N a_{ij}$. A sequence of edges $(i_1, i_2), (i_2, i_3), \dots, (i_{r-1}, i_r)$ is called a path from node i_1 to node i_r . The graph $\mathcal{G}$ is said to be connected if for any $i, j \in \mathcal{V}$ there is a path from i to j . The degree matrix is $\mathcal{D} = \text{diag}(d_1, d_2, \dots, d_N)$, and the Laplacian of $\mathcal{G}$ is defined as $\mathcal{L} = \mathcal{D} - \mathcal{A}_{\mathcal{G}}$. Disconnected graphs can be divided into c connected subgraphs $\mathcal{G} = \mathcal{G}_1 \cup \dots \cup \mathcal{G}_c$. A right stochastic matrix is a real square matrix, with each row summing to 1.

2 DYNAMICAL SYSTEMS WITH DIFFERENTIAL PRIVATE OUTPUTS

CPS monitor physical processes and take control decisions based on these measurements in order to drive the process to a specific state, or to behave in a specific manner. A general way to describe CPS is as a feedback control system, where the process evolves over time generating streams of data, such that at each time instant k , $\mathbf{x}(k) = [x_1(k), x_2(k), \dots, x_n(k)]$ describes the system states; for instance, hourly power consumption of a set of electricity users or monthly weights of a group of individuals. The information that is transmitted to a controller is $\mathbf{y}(k) = [y_1(k), \dots, y_m(k)]$, which is the output of the system (e.g., sensor measurements). Each $y_i(k)$ is equivalent to a query that takes the entire dataset (states) $\mathbf{x}(k)$, and discloses statistical information about the dataset. For example, the

Figure 1: Feedback control system with differential privacy mechanism.

average of the vector $\mathbf{x}(k)$. The measurements $\mathbf{y}(k)$ are then sent to a controller who computes a control signal $\mathbf{u}(k)$ to drive the states of the system to a desired state (e.g., a control strategy modifies the price of electricity in order to shape the power consumption).

We consider $\mathbf{x}(k)$ that contain sensitive information, and therefore their values have to be protected against an adversary trying to make inferences about individual states (e.g., each user consumption) with the possible help of side information. To deal with this problem, differential privacy has emerged as a framework to ensure privacy by perturbing the exact output before it is released. More specifically, DP adds random noise $\eta_i(k) \in \mathbb{R}$ to each output $y_i(k)$, to produce a new private output $\tilde{y}_i(k) = y_i(k) + \eta_i(k)$, as illustrated in Figure 1.

On the other hand, one of the unique properties of CPS is that a dynamic system possesses inherent sources of uncertainties and noise, such as the noise of the sensor readings or statistical perturbations to the state of the system. These inherent uncertainties are modeled in classical control theory as being random stochastic process, and they can reduce or eliminate the noise required by differential privacy, which we call inherent differential privacy. In other words, the output $\mathbf{y}(k)$ is already random, even without the addition of differential privacy noise $\boldsymbol{\eta}(k)$, and guarantees a certain level of differential privacy.

Notice also another difficulty of adding external differential privacy noise to CPS: in many practical applications, adding external noise to sensor measurements may not be easy. For instance, modifying thousands of smart meters, or replacing thousands of loop detectors in a highway such that they start including new random DP noise may incur in large costs. For this reason, taking advantage of inherent noise already in the system may help to preserve privacy without adding or changing the sensors.

In the following section we will define inherent differential privacy for linear-time invariant control systems.

2.1 Linear Time-Invariant System with Inherent Noise

Linear Time Invariant (LTI) systems are a classical model for control systems. They can be used to model a large variety of physical processes and how they respond to a control input. The advantages of LTI systems is that they can be characterized by using linear algebra tools. The typical control-theoretic model of an LTI system

is the following,

$$\begin{aligned} \mathbf{x}(k+1) &= A\mathbf{x}(k) + B\mathbf{u}(k) + \boldsymbol{\omega}(k) \\ \mathbf{y}(k) &= C\mathbf{x}(k) + \mathbf{v}(k) \end{aligned} \quad (1)$$

where $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $C \in \mathbb{R}^{p \times n}$ are matrices that help to describe the system evolution over time, for a given control input $\mathbf{u}(k) \in \mathbb{R}^m$.

This system is stochastic, because $\boldsymbol{\omega}(k) \in \mathbb{R}^n$ and $\mathbf{v}(k) \in \mathbb{R}^p$ are zero-mean Gaussian noises with variances $\sigma_{\omega}^2, \sigma_v^2$. The randomness in the system is typically used to model system uncertainties, external perturbations, and sensor noise. Since $\mathbf{v}(k)$ is an i.i.d. vector of zero-mean Gaussian noise, we can define the *intensity matrix* (covariance matrix) R_v as a diagonal matrix with elements $\sigma_{v,i}^2$. Similarly we can define the intensity matrix of $\boldsymbol{\omega}(k)$, R_{ω} . Let C_{ij} be the elements of matrix C such that the i^{th} output is $y_i(k) = \sum_{j=1}^n C_{ij}x_j(k) + v_i(k)$. For instance, for the aggregation query, $C = [1, 1, \dots, 1]$, $y(k) = \sum_{j=1}^n x_j(k)$.

In order to analyze how the random variables $\boldsymbol{\omega}, \mathbf{v}$ will affect the system, it is necessary to define a control strategy. *Depending on the controller, the random noise may be amplified or attenuated, therefore the inherent privacy of the system is strongly connected to the control action.* In this work, we focus on output-feedback control, but our analysis can be easily extended to other types of controllers.

2.1.1 Output Feedback Control. We assume that the controller consists of an output feedback control of the form $\mathbf{u}(k) = K\mathbf{y}(k)$. Since $\mathbf{u}(k) = K\mathbf{y}(k) = K(C\mathbf{x}(k) + \mathbf{v}(k))$, we can simplify Eq. (1) by defining $\bar{A} = A + BKC$:

$$\mathbf{x}(k+1) = \bar{A}\mathbf{x}(k) + \underbrace{BK\mathbf{v}(k) + \boldsymbol{\omega}(k)}_{\varphi(k)}, \quad (2)$$

where φ is a linear combination of Gaussian random variables. The covariance of $\varphi(k)$ is described by

$$\begin{aligned} R_{\varphi} &= E[\varphi(k)\varphi(k)^T] = E[(BK\mathbf{v}(k) + \boldsymbol{\omega}(k))(BK\mathbf{v}(k) + \boldsymbol{\omega}(k))^T] \\ &= E[BK\mathbf{v}(k)\mathbf{v}(k)^T K^T B^T] + E[BK\mathbf{v}(k)\boldsymbol{\omega}(k)^T] + \\ &\quad E[\boldsymbol{\omega}(k)\mathbf{v}(k)^T K^T B^T] + E[\boldsymbol{\omega}(k)\boldsymbol{\omega}(k)^T] \\ &= BKE[\mathbf{v}(k)\mathbf{v}(k)^T]K^T B^T + E[\boldsymbol{\omega}(k)\boldsymbol{\omega}(k)^T] \\ &= BKR_v K^T B^T + R_{\omega} \end{aligned} \quad (3)$$

One particular property of feedback stochastic systems, is that the variance of the system states is not constant and it evolves over time. Therefore, it is possible to characterize the variance of $\mathbf{y}(k)$ to define the inherent level of privacy of the feedback system *without adding differential privacy* noise. Noise-driven systems have been widely studied in the control theory literature [3] and it is possible to find an expression of the variance evolution and how the inherent noise is amplified or attenuated by the feedback system.

Let $E[\mathbf{x}(k)] = \mathbf{m}(k)$ be the expected value of $\mathbf{x}(k)$. Since $E[\mathbf{v}(k)] = \mathbf{0}$ and $E[\boldsymbol{\omega}(k)] = \mathbf{0}$, we have that $\mathbf{m}(k+1) = \bar{A}\mathbf{m}(k)$ for $\mathbf{m}_0 = E[\mathbf{x}(0)]$ the initial state.

Since the state is a vector of size n , the variance of $\mathbf{x}(k)$ is defined by a matrix

$$\text{var}(\mathbf{x}(k)) = E[(\mathbf{x}(k) - \mathbf{m}(k))(\mathbf{x}(k) - \mathbf{m}(k))^T] = Q(k)$$

which leads to

$$Q(k) = E[\mathbf{x}(k)\mathbf{x}(k)^T - \mathbf{x}(k)\mathbf{m}(k)^T - \mathbf{m}(k)\mathbf{x}(k)^T + \mathbf{m}(k)\mathbf{m}(k)^T]. \quad (4)$$

Combining Eq. (4) for $k+1$ with Eq. (2) yields the dynamics of the variance matrix :

$$\begin{aligned} Q(k+1) &= E[\bar{A}\mathbf{x}(k)\mathbf{x}(k)^T \bar{A}^T + \bar{A}\mathbf{x}(k)\varphi(k)^T + \varphi(k)\mathbf{x}(k)^T \bar{A}^T \\ &\quad - \bar{A}\mathbf{x}(k)\mathbf{m}(k)^T \bar{A}^T - \varphi(k)\mathbf{m}(k)^T \bar{A}^T - \bar{A}\mathbf{m}(k)\mathbf{x}(k)^T \bar{A}^T \\ &\quad + \bar{A}\mathbf{m}(k)\mathbf{m}(k)^T \bar{A}^T - \bar{A}\mathbf{m}(k)\varphi(k)^T - \varphi(k)\varphi(k)^T]. \end{aligned} \quad (5)$$

Since the noise φ is independent of the states, the operator $E[\cdot]$ over all cross terms become zero. Combining with Eq. (4) we have that

$$Q(k+1) = \bar{A}Q(k)\bar{A}^T + R_{\varphi}. \quad (6)$$

If $\bar{A}$ is stable, $Q(k)$ will converge to Q [3], which is the solution of the Riccati equation of the form $\bar{A}Q\bar{A}^T + R_{\varphi} - Q = 0$.

These results help us for computing the variance of $\mathbf{y}(k)$, $\text{var}(\mathbf{y}(k)) = Q_y(k) = CQ(k)C^T + R_v$. Clearly, the output $\mathbf{y}(k)$ can be described by a Gaussian distribution with a time-varying standard deviation $\sigma_{y,i}(k)$ that corresponds to the square root of the diagonal elements of $Q_y(k)$.

2.2 Inherent Differential Privacy

We have obtained an expression for the time-varying standard deviation of the output $\sigma_{y,i}(k)$ that depends on $\bar{A}, C, R_{\varphi}$, and R_v . In order to relate $\sigma_{y,i}(k)$ with the differential privacy framework, we can define the sensitivity of the output as follows.

Let $\mathbf{x}, \mathbf{x}'$ be two γ -adjacent data sets according to Definition 1.1. Let $\mathbf{y} = [\mathbf{y}(0), \mathbf{y}(1), \dots, \mathbf{y}(T)]$ be the trace output vector of $\mathbf{x}$ and $\mathbf{y}'$ the output from $\mathbf{x}'$ for T iterations. We can define the sensitivity of the noiseless output for two γ -adjacent state traces $\mathbf{x}, \mathbf{x}'$ by

$$\Delta_{o,2} = \max_{\mathbf{x}, \mathbf{x}'} \|\mathbf{y} - \mathbf{y}'\|.$$

Recall from Definition 1.1 that at each time instant k , all the elements of $\mathbf{x}(k), \mathbf{x}'(k)$ are equal except for the l^{th} elements, such that $|x_l(k) - x'_l(k)| \leq \gamma$. Since each $y_i(k)$ is a linear combination of the elements of the dataset, we have that

$$|y_i(k) - y'_i(k)| = \left| \sum_{j=1}^n C_{ij}x_j(k) - C_{ij}x'_j(k) \right| \leq |C_{il}\gamma|.$$

The sensitivity is then bounded by

$$\Delta_{o,2} \leq \sqrt{\max_l \sum_{k=0}^T \sum_{i=1}^p |C_{il}\gamma|^2}.$$

Recall Lemma 1.4, where given δ and σ , we have the following ϵ level of privacy,

$$\epsilon \geq \frac{\sqrt{2 \ln(1.25/\delta)} \Delta_{o,2}}{\sigma}.$$

In our formulation, because the standard deviation of inherent noise of the output evolves over time, the total level of privacy of the entire time-series also evolves over time, and it can be defined by

$$\epsilon_y(k) \geq \frac{\sqrt{2 \ln(1.25/\delta)} \Delta_{o,2}}{\min_i \sigma_{y,i}(k)}. \quad (7)$$

Therefore, without the addition of differential privacy noise, we have obtained (ϵ_y, δ) -differential privacy that depends only on the inherent noise of the system.

Notice that the term $\min_i \sigma_{y,i}(k)$ indicates that privacy is dictated by the less noisy output. Also, it is worth to highlight that, according to Eq. (6), the variance of the states depends on the control parameter K (since $\bar{A}$ depends on K), such that tuning K may lead to a larger level of privacy.

2.3 Injecting minimum noise

We have shown how the typical uncertainties assumed in most control systems can help to maintain an inherent-level of differential privacy that evolves over time depending on the system dynamics and the control algorithm; however, some applications may require a desired level of privacy ϵ, δ . As a consequence, if $\epsilon_y < \epsilon$, it is necessary to inject additional noise to the system.

Recall that if two random variables are Gaussian, then their addition is also a Gaussian random variable. Because the noise typically inherent in control systems is Gaussian, by adding a Gaussian DP mechanism $\eta(k) \sim N(0, \sigma_\eta^2)$ to the system outputs, we are able to add the remaining Gaussian noise needed to achieve our desired ϵ, δ . In practice, noisy measurements may cause a degradation in the system that can wear down actuators or lead the states to undesired values. For this reason, it is important to inject the minimum amount of noise necessary to preserve a desired level of privacy. Since the variance of the output is time-varying, the variance of the added noise should also evolve over time, according to the following Theorem.

THEOREM 2.1. *Let ϵ, δ be the parameters of the desired level of privacy. Let σ^2 be the desired variance of the form*

$$\sigma \geq \sqrt{2 \ln(1.25/\delta) \Delta_{\alpha,2} / \epsilon}.$$

If at each time instant k , the differential privacy mechanism adds noise η with an intensity matrix that satisfies

$$R_\eta(k) = \sigma^2 I_N - CQ(k)C^\top - R_v, \quad (8)$$

the mechanism ensures (ϵ, δ) -Differential privacy.

Proof:

We can define $\tilde{\varphi}(k) = BK(v(k) + \eta(k)) + \omega(k)$ with a time-varying intensity matrix $\tilde{R}_\varphi(k) = BK(R_v + R_\eta(k))K^\top B^\top + R_\omega$, such that $Q(k+1) = \tilde{A}Q(k)\tilde{A}^\top + \tilde{R}_\varphi(k)$. Note that $Q_y(k) = CQ(k)C^\top + R_v + R_\eta(k)$, such that by using Eq. (8) we obtain $Q_y(k) = \sigma^2 I_N$ for all k , which satisfies the desired privacy level. ■

3 CASE STUDY: REAL-TIME PRICING IN SMART GRIDS

3.1 Demand Response Model

To show the generality of our approach, we include an example from smart grid problems. The goal in demand response systems is to incentivize consumers (industry or residential) to modify their electricity consumption $x(k)$ based on an electricity cost signal $\lambda(k)$ provided by the utility or by a demand-response company like EnerNOC.

In this section we follow the real-time pricing model from Tan et al. [21]. This model considers a market with N consumers of

electricity, a set of suppliers of electricity, and a third party entity—an Independent System Operator (ISO)—with the goal of matching supply and demand by setting the market price for electricity. The general assumption is that the ISO determines, at each time instant $k \in \mathbb{N}_+$, a clearing price $\lambda(k)$ valid for the period of time $[k \cdot \tau, (k+1) \cdot \tau]$ (this is called an *ex-ante* market) every τ hours (e.g., $\tau=0.5h$) and announces it to the suppliers and consumers.

The electricity demand of each user is characterized by two components: a baseline electricity consumption $b_i(k)$ that captures the electricity consumption that is independent of the pricing mechanism (i.e., the necessary power to satisfy the main consumer needs at each instant k such as refrigerator, cooking devices, light bulbs), and a price-responsive demand $r_i(\lambda(k))$, which captures the amount of electricity consumption that can be controlled by the pricing signal $\lambda(k)$. For instance, doing laundry when the price is low, or turning off the lights of rooms that are not being used.

The Constant Elasticity of Own-price (CEO) has been commonly adopted to characterize the total price-responsive demand [7, 13]. The CEO model is defined by

$$r_i(\lambda(k)) = D_i \lambda(k)^{\psi_i} \quad (9)$$

where $D_i > 0$ is a constant that properly scales $r_i(k)$ and $\psi_i \in (-1, 0)$ is the *price elasticity demand* that captures how the demand is affected by a specific price $\lambda(k)$ [13].

The electricity used by each consumer is given by the following model,

$$x_i^c(k+1) = \alpha x_i^c(k) + \beta_i(\psi_i) \lambda(k) + b_i(k) + \omega_i(k) \quad (10)$$

where $r_i(\lambda(k)) = \beta_i(\psi) \lambda(k)$ such that $\beta_i(\psi)$ linearly relates the CEO model, α indicates the effect of the past consumption in the current demand and we assume it is the same for all users, and $\omega_i(k)$ is a random noise that models the uncertainty in how consumers will react to market prices. The total power consumption is given by

$$\begin{aligned} x_T^c(k+1) &= \sum_{i=1}^N x_i^c(k+1) \\ &= \alpha x_T^c(k) + \beta_T \lambda(k) + b_T(k) + \omega_T(k) \end{aligned} \quad (11)$$

for $\beta_T = \sum_{i=1}^N \beta_i$, $b_T(k) = \sum_{i=1}^N b_i(k)$, and $\omega_T(k) = \sum_{i=1}^N \omega_i(k)$.

Similarly, for the supply of electricity, Tan et al. [21] propose a linear regression between supply and cost, a model they validated with the Australian Energy Market Operator and the electricity market in California. Under these assumptions the total supply of electricity can be modeled by $x_T^s(k+1) = p\lambda(k) + q$, where p and q are parameters estimated by historical market data from the area of study.

3.2 Control Objective

The control objective of the ISO is to send price signals $\lambda(k)$ to the users to keep the error between supply and demand of electric power $\mathcal{E}(k) = x_T^s(\lambda(k)) - x_T^c(\lambda(k))$ close to zero for every time instant k . This can be seen as a control problem in which the system to be controlled is the outcome of a market, the control variable is the price signal $\lambda(k)$ and measured variable is the error $y_{\mathcal{E}}(k) = \mathcal{E}(k) + v_T(k)$ for $v_T(k)$ being the total measured noise.

Figure 2 shows the feedback interaction between the consumers or smart meters, suppliers, and the ISO. The consumption of users is monitored by the smart meters and each smart meter sends x_i^c to

Figure 2: General feedback scheme for real time pricing.

an Energy Data Center (EDC) or directly to the ISO to obtain the aggregated value $x_T^c(k)$ [20] and calculate the measurement $y_E(k)$. The information released by the EDC can be used by a third party like the ISO to calculate new prices.

The price signal $\lambda(k)$ must be carefully designed in order to avoid oscillations or even instability [17, 21]. The proposed integral control strategy in [21] is described by

$$\lambda(k+1) = \lambda(k) - Ky_E(k), \quad (12)$$

which has been proven to keep the system stable with an appropriate selection of K . The compact representation of the real-time pricing model is given by

$$\underbrace{\begin{bmatrix} x_T^s(k+1) \\ x_T^c(k+1) \\ \lambda(k+1) \end{bmatrix}}_{x(k+1)} = \underbrace{\begin{bmatrix} 0 & 0 & p \\ 0 & \alpha & \beta_T \\ -K & K & 1 \end{bmatrix}}_{\bar{A}} \underbrace{\begin{bmatrix} x^s(k) \\ x^c(k) \\ \lambda(k) \end{bmatrix}}_{x(k)} + \underbrace{\begin{bmatrix} q \\ \omega_T(k) + b_T(k) \\ Kv_T(k) \end{bmatrix}}_{\varphi}$$

with output

$$y_E(k) = \underbrace{\begin{bmatrix} 1 & -1 & 0 \end{bmatrix}}_C x(k) + v_T(k).$$

Note that our previous analysis of variance and inherent privacy can be applied to this model for

$$R_\varphi = \begin{bmatrix} 0 & 0 & 0 \\ 0 & R_w & 0 \\ 0 & 0 & K^2 R_v \end{bmatrix}.$$

3.3 Differential Privacy in RTP

Smart meters allow the utility provider to monitor consumption in order to update prices or adjust generation. However, due to the accuracy and granularity of the data, it could be possible for a curious entity to estimate behavior profiles of each user, and identify their consumption patterns. For example, with the field of non-intrusive load monitoring, it is possible to extract information about the type of appliance that is being used [14]. Since the EDC derives aggregated statistics from consumer information, it is necessary to ensure that it is not possible to learn anything about the activities of individual households. For example, a third party that has access

to x_T^c cannot tell whether or when a user was doing laundry. The system uncertainties ω, v are able to maintain certain levels of privacy $\epsilon_y(k)$ for data aggregation, as it was defined in Eq. (7). In the case where additional noise is required, it is necessary to add noise $\eta(k)$ to the aggregated information $x_T^c(k)$ at each instant k , such that the new DP output is $\tilde{y}_E(k) = y_E(k) - \eta(k)$. This mechanism for smart meters is inspired by previous privacy mechanisms, such as the distributed smart grid differential privacy problem introduced in [1], where each smart meter adds noise from Gamma distributions, such that the aggregation results in a Laplace noise. However, in this work we focus on Gaussian mechanism, which can be easily deployed in a distributed fashion.

3.4 Experiments

For our experiments, we use a distribution feeder specification [18] which covers a moderately populated urban area composed by 1405 households. We model the consumption behavior of each user using the linear model introduced above. Based on [21], we draw the parameters from a normal distribution $D_i \sim N(7, 3.5^2)kW$, $\psi \sim N(-0.8, 0.1^2)$ where $\beta_i = D\psi\lambda_0^{\psi-1}$. $\alpha_i = 0.1kW$ and $K = 1$. Additionally, we use a half-hourly baseline consumption $b_i(k)$ based on historical data provided by the NY ISO such that $b_i(k) \in [0.28, 0.76] kW$. We assume that the maximum consumption is $x_{max}^c = 1.6 kW$ and $T = 24 h$, such that the sensitivity is $\Delta_{0,2} = 24x_{max}^c$. For the supply model, we use the parameters $p = 43.6e-4$, $q = 1.28$, which are scaled versions from experimental data from the NSW region in Australia (see [21] for further details on the selection of these parameters), $\omega_i(k) \sim N(0, 0.007^2)$ and $v_i(k) \sim N(0, 0.006^2)$.

Figure 3 depicts the supply-demand mismatch $\mathcal{E}(k)$ for a 24 h period and for the control parameter $K = 1$. The control is able to take the output close to zero for $\lambda_0 = 21 \$/MWh$. Note that due to the system and sensor noise, $\mathcal{E}(k)$ is a random variable from a Normal distribution. Since we only have one output, $\sigma_y(k) = \sqrt{Q_y}(k)$. This parameter evolves over time until it reaches its steady state, therefore according to Eq. (7), the inherent privacy level $\epsilon_y(k)$ also evolves in time.

The control parameter K plays a fundamental role not only in the stability of the system but also in its noise-attenuation or amplification properties also called *sensitivity* (which should not be confused with the sensitivity of differential privacy). Sensitivity is a tool from control theory usually implemented to analyze how an external input (e.g., disturbances) affect the system output. As it was studied in [9], the proposed RTP strategy tends to amplify high frequency inputs as K increases. In our case, Gaussian noise has spectral components in all frequencies, which means that the injected noise tends to be amplified with K . Figure 4 illustrates the steady state standard deviation of the output σ_y for different K and for $\delta = 0.01$. Clearly, the closest the system approaches an instability region, the largest the output noise becomes. Note that a specific level of privacy can be obtained without adding any privacy mechanism. For instance, if $K > 1.5$, (0.1, 0.001)-differential privacy can be preserved.

On the other hand, Figure 4 (bottom) depicts the steady state ϵ_y for different K and δ . Clearly, since δ relaxes the privacy definition such that less noise is required, it is possible to ensure different

Figure 3: Supply-demand mismatch (output Y) without differential privacy mechanism for $K = 1$ (top) and the standard deviation of the inherent noise (bottom)

levels of privacy for a given K without injecting any additional noise.

Figure 4: Effect of the control parameter K in the amount of inherent noise σ_Y and the inherent privacy level ϵ_Y .

In order to achieve a desired level of privacy, it may be necessary to add a differential privacy mechanism $\eta(k)$. The intensity of the noise that has to be injected can be obtained according to (8). We aim to maintain $(0.1, 0.01)$ -differential privacy, such that, according to Definition 1.4, $\sigma_d \geq 0.354$. Without DP, σ_Y tends to 0.133, such that additional noise is necessary. Figure 5 depicts the evolution of the standard deviation $\sigma_{\eta}(k)$ of the injected noise that maintains a constant $\sigma_Y = \sigma_d$ with minimum noise, such that $(0.1, 0.001)$ -differential privacy is preserved for all time.

Figure 5: Comparison of the supply-demand mismatch with and without DP mechanism for $K = 1$ (Top). To achieve a desired $(0.1, 0.01)$ - differential privacy the standard deviation of the output must be $\sigma_d = 0.354$. The time-varying differential privacy mechanism $\eta(k)$ with standard deviation $\sigma_{\eta}(k)$ that ensures $\sigma_Y = \sigma_d$ is depicted (bottom).

4 NOISELESS PRIVACY MECHANISMS

As we studied before, adding noise enables some desired levels of privacy by making outputs indistinguishable for changes in at most one user. However, some systems are very susceptible to external noise and even convergence to a desired steady state can be affected. For instance, Figure 6 depicts a distributed frequency control system with a differential privacy mechanism. The control objective is to maintain frequency of all nodes in the grid (i.e., distributed generators, loads) synchronized to 60 Hz. Unfortunately, the addition of noise caused the frequency to fluctuate and reach unsafe states (i.e., frequency should not fluctuate more than 1 Hz). In particular, distributed control strategies that rely on the exchange of local information are largely affected by the injection of noise [19]. For instance, the well known consensus algorithm tends to a Brownian motion that never reaches a steady state. In the rest of this work we focus on distributed multi-agent strategies, for which DP might not be a practical solution, and thus require different privacy enhancing mechanisms.

4.1 Privacy of Random Events

A different domain of privacy consists on focusing only on specific changes in the dataset that we call *events*. Let us consider the system in Eq. (1) with output $y(k) = Cx(k)$. Let $\Omega_e = \{t_{e1}, t_{e2}, \dots\}$ be the set of time stamps at which events that need to remain hidden occur. We consider only one event at each $t_{ek} \in \Omega_e$ and it changes the i^{th} system state by a magnitude $\Delta x_i(t_{ek})$, such that $y(t_{ek}) = Cx(t_{ek}) + \Delta y(t_{ek})$, where $\Delta y(t_{ek}) = C[0, \dots, 0, \Delta x_i(t_{ek}), \dots, 0]$ is the change in the output caused by the event. Notice that $\Delta x_i(t_{ek})$ can be considered as an impulse function that suddenly changes

Figure 6: Frequency synchronization with DP mechanism for $\gamma = 0.1$ and $\epsilon = 0.3$. The injected Laplace noise induces a random walk and the system never reaches the virtual leader.

the states only at the instant t_{ek} and is zero otherwise. For instance, if the states are already in the equilibrium, the event will take the system to new initial states and then they will converge again to the equilibrium. Thus, we call *event duration* to the time it takes to the system to recover from the event, i.e., time to converge to the steady state. The magnitude of the event affects the event duration, such that it is necessary to keep both of them private.

For example, imagine that the employees in a company do not care about maintaining privacy of their salaries so there is no need to add noise to a query that uses their information. However, at a point, one of them may not want to share a salary increase, so she prefers just to share her previous salary. This kind of changes that may or not be given to a query can be represented as an uncertainty in the states similar to $\omega(k)$ in Eq. (1). As another example, we can consider the power electricity consumption of a military base. In order to ensure a correct power scheduling and control, it is necessary to share specific consumption information with peering micro-grids so that they remain synchronized; however, at the times of military tests they can share outdated information and hide the current consumption state to prevent malicious adversaries from inferring these events (see Figure 7).

We call this type of privacy *event-based privacy*, where only specific events are meant to remain private. This type of privacy for control systems was first introduced in our previous work [8], where we analyzed the consensus algorithm with sampled information such that the amount of data transmitted to some neighbors is minimized, hiding the presence of some events. In particular, our previous work proposed the notion of *discretionary sampling* as a new privacy enhancing technology for control systems.

Figure 7: Example of events in a smart grid. It is necessary to share information about events in order to ensure a correct behavior of the system, but some of these events need to remain private.

4.2 Synchronization of Multi-Agent Systems

The type of systems that we want to analyze are known as multi-agent systems (MAS). MAS can model interactions among agents that share local information in order to achieve a common goal. MAS have been widely studied in several research fields, such as communications, social networks, vehicle coordination, and neuroscience, just to name a few. One of the most popular MAS models is known as the consensus algorithm, where a group of agents update their state (e.g., X,Y position of a vehicle) by comparing their current states with a set of neighbors [15]. Let $\mathcal{G}_p(\mathcal{V}, \mathcal{E}_p, \mathcal{A}_p)$ be the physical graph that describes the natural information flow among N agents with adjacency matrix $\mathcal{A}_p$. The continuous-time MAS model is described by

Figure 8: Physical and communication graphs. Even though the physical topology is disconnected, including a communication network leads to a connected graph.

$$\dot{x}_i = - \sum_{k=1}^N a_{ik} (x_i - x_k) + u_i, \quad (13)$$

where a_{ij} are the elements of matrix $\mathcal{A}_p$, and u_i is a control action that helps to achieve specific objectives. Typically, when $u_i = 0$ for all i , consensus is only achieved if the graph representing the interaction between agents is connected [15]. Based on the results in [10] and [8], a communication network can be included in order to allow synchronization to a desired state even if the physical graph $\mathcal{G}_p$

is disconnected. Thus, the communication network compensates the lack of connectivity. Let $\mathcal{G}_c(\mathcal{V}, \mathcal{E}_c, \mathcal{B}_c)$ be the communication graph that describes the exchange of information through a communication network. Figure 8 illustrates both topologies and how they interact to form a joint graph that compensates the lack of connectivity.

The control action that drives the states to a desired reference x^r relies on the information transmitted through the communication network, an is described by

$$u_i = -K_i \sum_{k=1}^N b_{ij} (x_i - x_j) + b_{i(N+1)} (x_i - x^r). \quad (14)$$

Using the properties of the Laplacian matrices, the system dynamics can be described in compact form by

$$\dot{\mathbf{x}} = -(\mathcal{L}_p + K\mathcal{L}_c)\mathbf{x} \quad (15)$$

where $\mathcal{L}_p, \mathcal{L}_c$ are the Laplacian matrices of the physical topology and the communication network, respectively. Let $L_T = \mathcal{L}_p + K\mathcal{L}_c$. We then have that

$$\dot{\mathbf{x}} = -L_T \mathbf{x}.$$

Stability of this type of systems depends completely on L_T . If all eigenvalues of L_T are positive and it only contains one zero eigenvalue (the union of both graphs is connected), all the states converge to x^r . We want to extend our previous results [8] and take advantage of some properties of consensus-like systems that let us to manipulate the amount of information that is transmitted through the communication network, in order to preserve the privacy of events.

First, let us define the synchronization algorithm with sampled information.

4.3 Consensus with Sampled Information

We can define $y_j^\tau(t)$ as the data received at the instant t , which is held using a zero-order hold. This value can be described by a piecewise function that remains constant in the interval $t \in [k\tau, (k+1)\tau)$, for $\tau > 0$ (the sampling period). We make two main assumptions: i) the information that each controller i knows about x_i is not sampled and it is continuous; ii) the information received from neighbors is sampled and corresponds to $y_j^\tau(t)$. By definition, the Laplacian matrix can be divided into $\mathcal{L}_c = D_c - \mathcal{B}_c$, and we can rewrite the consensus algorithm as follows

$$\dot{\mathbf{x}} = -(\mathcal{L}_p + K D_c)\mathbf{x} + K \mathcal{B}_c \mathbf{y}^\tau \quad (16)$$

where $\mathbf{y}^\tau = [y_1^\tau, \dots, y_N^\tau]^\top$.

LTI dynamic systems with sampled information and a zero-order hold can be described in a discrete-time fashion by defining the transformation matrices $\Phi = e^{-L_T \tau}$ and $\Gamma = \int_0^\tau e^{-L_T s} ds K \mathcal{B}_c$, such that

$$\mathbf{x}(k+1) = \Phi \mathbf{x}(k) + \Gamma \mathbf{y}(k). \quad (17)$$

Clearly, the information that is being transmitted corresponds to the system states, such that $\mathbf{y}(k) = \mathbf{x}(k)$. As a consequence, conditions for synchronization are given by the eigenvalues of $P = \Phi + \Gamma$, according Theorem A.1, and stability is preserved independently of the sampling period τ if the union of the physical and the communication graphs forms a connected graph; i.e., the eigenvalues of L_T are positive and L_T contains only one zero eigenvalue.

4.4 Convergence Time

Even though the proposed synchronization system achieves synchronization in finite time independently of the sampling period, convergence time increases with τ . It is possible to find some bounds on the time it takes for all agents to reach the desired reference x^r based on the discrete-time representation.

Let S be the matrix formed by the eigenvectors of P and Λ be the diagonal matrix of eigenvalues. Then, using the diagonalization transformation we can rewrite the discrete-time system as $\mathbf{x}(k+1) = S\Lambda S^{-1}\mathbf{x}(k)$. Solving the difference equations we obtain

$$\mathbf{x}(k+1) = P^k \mathbf{x}(0) = S\Lambda^k S^{-1} \mathbf{x}(0). \quad (18)$$

We prove in Theorem A.1 that the eigenvalues of P $\mu_i(P)$ are inside the unitary disk, such that the sequence is a Cauchy sequence, and for $m, n > k^*$,

$$\|\mathbf{x}(m) - \mathbf{x}(n)\| = \|S(\Lambda^m - \Lambda^n)S^{-1}\mathbf{x}(0)\| < \varsigma,$$

for $\varsigma > 0$. $\Lambda = \text{diag}(\mu_1, \dots, \mu_{N+1})$ corresponds to the eigenvalue matrix. Therefore, from (18) we have that $(\Lambda^m - \Lambda^n) = \Lambda^{k^*}(\Lambda^{m-k^*} - \Lambda^{n-k^*})$, which yields to convergence when the elements of Λ^{k^*} tend to zero, except for the eigenvalue equal to 1. We define $\mu_1 = 1 > |\mu_2| > \dots > |\mu_{N+1}|$ the ordered eigenvalues of P and $|\hat{\mu}|$ the largest eigenvalue different from 1. The convergence state is reached after k^* steps such that

$$k^* = \{k : \max_{i \in Y} |\mu_i^k| \leq \varsigma\},$$

for $\varsigma > 0$ significantly small. It is easy to verify that $|\mu_i^k| = |\mu_i|^k$, which leads to

$$k^* \leq \ln(\varsigma)/\ln(\max |\mu_i|) = \ln(\varsigma)/\ln |\hat{\mu}|,$$

and that the convergence time t^* is bounded by $\tau \cdot k^*$. Note that $\hat{\mu}$ depends on the sampling period τ , since P depends on τ .

Now that we have defined the main properties of synchronization under sampled information, we can introduce two event-based privacy mechanisms.

4.5 Privacy Mechanisms

Using the results obtained above, we can define two privacy mechanisms that aim to hide information about the magnitude and duration of an event. Event though an adversary can note that some information is missing, it would be hard to infer specific properties about the event.

4.5.1 Privacy by Data Minimization. As it was described above, an event consists on a sudden change in the system states that causes oscillations while the controller drives the system back to stability. Figure 9 depicts an example of an event whose information is transmitted with different sampling periods. Note that with a small sampling period, it is more likely that information about the event will be transmitted. On the other hand, increasing the sampling period increases the chances that information from an event will be hidden. As we shown above, our proposed control strategy maintains stability independent of the sampling period. Therefore, we can hide information by minimizing the amount of data shared between agents. However, since the transmission of information is periodic, it is not possible to ensure that all events

Figure 9: Example of privacy by data minimization. Increasing sampling period leads to less transmitted information.

will remain hidden. Therefore, it is possible to select a sampling period τ that minimizes the amount of sensitive information that is shared, but is not possible to ensure complete privacy. See for instance Figure 12 which depicts the number of shared events with respect to the sampling period for the case study introduced below.

One of the main advantages of data minimization is that it is not necessary to know the specific time of an event.

4.5.2 Discretionary Sampling: Selecting When to Sample and When to Lie. In discretionary sampling, each agent is able to lie about each of the events that needs to remain private.

Notice that if an agent decides not to send information during an event, an adversary can notice that the agent is silent and therefore it can infer it is trying to hide an event. Therefore, instead of not sending updates during the event, we force agents to repeatedly send the last sampled data before the sensitive event started. For instance, if it is necessary to hide the information between times k_1 to k_2 , the transmitted data in that interval will correspond to $y(k_1 - 1)$, which is the last piece of information shared before k_1 . Using some results from switched control systems with packet losses [22], we are able to obtain the conditions that achieve synchronization while *hiding any desired information*. Let $z(k) = [x(k), y(k-1)]^\top$ be the augmented state variable. Let $\alpha(k) \in \mathbb{R}^{N+1}$ be the matrix whose diagonal elements $\alpha_i(k) = 1$ if a packet in node i is transmitted and 0 if it is sending outdated data in the instant k .

$$z(k+1) = \underbrace{\begin{bmatrix} \phi + \Gamma\alpha(k) & \Gamma(I - \alpha(k)) \\ \alpha(k) & (I - \alpha(k)) \end{bmatrix}}_{A(k)} z(k) \quad (19)$$

The stability of these types of systems can be determined by the set of matrices A_j , such that $A_T = A_1 A_2 \dots A_k$. Therefore, if A_T contains a spanning tree, synchronization is achieved as stated in the following Lemma.

LEMMA 4.1. *Let $A_1, A_2, \dots, A_k$ be the sequence of time-varying matrices. If A_i is stochastic and if the union of the set of directed graphs has a spanning tree, the matrix product $A_T = A_k A_{k-1} \dots A_1$ also possesses a spanning tree such that there exists a vector v that satisfies $A_T v = v$.*

Recall from Theorem A.1 that P is right stochastic. Since $\phi + \Gamma\alpha(k) + \Gamma(I - \alpha(k)) = P$ and the summation of $\alpha(k), (I - \alpha(k))$ is I ,

matrix $A(k)$ in Eq. 19 is right always stochastic. The following Theorem establishes conditions for synchronization with discretionary sampling.

THEOREM 4.1. *Let us consider the controlled system with sampled information in Eq. (19) with the proposed discretionary privacy mechanism. Synchronization is asymptotically achieved in finite time if conditions in Theorem A.1 are satisfied and each agent hides its information for at most $T_i < \infty$ steps.*

Sketch of the proof:

According to Lemma 4.1, it is necessary to ensure that the union of the sequence of directed graphs possesses a spanning tree. Since $P = \Phi + \Gamma$ is stable and, T_i is finite, updating the information at least after T_i iterations ensures the existence of a spanning tree in $A(k)$. Thus, synchronization to the reference state $z(\infty) \rightarrow x^r 1_{2N}$ is asymptotically achieved $\blacksquare$.

While large sampling intervals might miss some events, they cannot provide strong privacy guarantees. Discretionary sampling (i.e., lying) can, however, achieve perfect secrecy if required (it can hide all events if the operator wants) at the cost of longer consensus settling times. Besides, to execute a discretionary sampling strategy it is necessary to have prior knowledge of all the event times and their intended durations, which is not always possible.

4.5.3 Impact of Discretionary Sampling. Similarly to the increasing sampling mechanism, we can find some bounds on the convergence time for the discretionary sampling case. Recall that discretionary sampling transmits data at a base sampling period τ and then each agent may choose to hide some specific information for at most T_i steps. Let $T = \max T_i$ be the maximum number of steps any agent can hide information. Since hiding information delays the system convergence to the consensus state, the worst case occurs if all agents start hiding information simultaneously at an instant k_h during T steps such that their next state update occurs at $k_h + T$. Clearly, this is equivalent to sampling at a sampling period of $\bar{\tau} = \tau T$. Thus, we can define $\bar{\Phi} = e^{-L\tau T}$, $\bar{\Gamma} = \int_0^{\tau T} e^{-L\tau s} ds K \mathfrak{B}_c$, and $\bar{P} = \bar{\Phi} + \bar{\Gamma}$. The maximum time of convergence is then bounded by $t^* \leq \tau T (\ln(\zeta) / \ln(\bar{\mu}))$ for $\bar{\mu}$ the largest eigenvalue different from 1 of $\bar{P}$.

4.6 Comparison with Differential Privacy for Events

The discretionary sampling mechanism can be analyzed from the differential privacy framework by making some assumptions about Δx and taking advantage of its uncertainty. We can assume that Δx is drawn from a Uniform random distribution. In other words, the likelihood of occurrence of an event $\Delta x \in [\Delta_{\min}, \Delta_{\max}]$ is uniform and the PDF is of the form

$$pU(t) = \begin{cases} \frac{1}{\Delta_{\max} - \Delta_{\min}}, & \text{if } \Delta_{\min} \leq t \leq \Delta_{\max} \\ 0, & \text{otherwise} \end{cases}.$$

Let us consider the states $x(k_e)$ and its adjacent state $x'(k_e) = x(k_e) + \Delta x(k_e)$, where $\Delta x(k_e) = [0, \dots, \Delta x_l(k_e), \dots, 0]$ such that an event occurs in the state l at a time k_e .

Notice that $y_i(k_e)$ is a linear combination of the dataset $x(k_e)$. Therefore, in the presence of an event in state l and, since there is

only one event at each time instant,

$$y'_i(k_e) = \sum_{j=1}^n C_{ij}x_j(k_e) + C_{il}\Delta x_l(k_e).$$

When the event is hidden, we assume $\Delta x_l(k_e) = 0$, i.e., $y_i(k_e) = \sum_{j=1}^n C_{ij}x_j(k_e)$.

The relationship between the likelihood of an output with an event or hiding the event is then given by

$$\frac{pY(y'_1, y'_2, \dots, y'_p)}{pY(y_1, y_2, \dots, y_p)} = \prod_{i=1}^p \frac{pU(\eta_i = C_{il}\Delta x_l(k_e))}{pU(\eta'_i = 0)} = 1.$$

The probability over the entire domain of y is equal with and without an event, such that it is not possible to distinguish whether or not the event was hidden. This is $\epsilon = 0$ Differential privacy, which is the highest level of privacy. As a consequence, for the discretionary sampling mechanism, even if the transmitted information consists of a sequence of the same data during a time T , it does not affect the privacy of the system states. In other words, an adversary observing the same data for several iterations, may suspect that some information is being hidden, but the magnitude and duration of such information is never revealed.

4.7 Experiment: Microgrid Synchronization

Microgrids can be defined as a group of interconnected loads and distributed energy sources that can interact with other microgrids or with the main grid. One of the main properties of microgrids is that they can supply power even if they are disconnected from the main grid. The problem of frequency synchronization among microgrids is important to ensure a correct connection and disconnection and avoid undesired oscillations or even instability. We model the power network using a multi-agent approach. We define the set of N agents as $\mathcal{V}$, where each one can be modeled as synchronous generators or DC-sources with inverters [5]. We assume that the power AC network is described by the *purely inductive* admittance matrix $Y \in j\mathbb{R}^{N \times N}$, the nodal voltage magnitude E_i for all $i \in \mathcal{V}$. The physical interaction of the power grid can be modeled as a graph $\mathcal{G}_p = \{\mathcal{V}, \mathcal{E}, \mathcal{A}_p\}$, where the elements of $\mathcal{A}_p$ are $a_{ij} = |E_i||E_j||Y_{ij}|$ and correspond to the maximum real power transfer between nodes i and j .

Let $x_i(t)$ corresponds to the frequency of each node. The dynamics of the system can be described by

$$D_i \dot{x}_i(t) = - \sum_{j=1}^N a_{ij}(x_i - x_j) + u_i \quad i \in \mathcal{V} \quad (20)$$

where $D_i > 0$ is a damping coefficient for synchronous generators, and it is also related with the slope of the droop controller of DC-sources with inverters. The control action is the one described in Eq. (14).

We consider the IEEE 30 nodes benchmark in Figure 10, where Area 2 is disconnected such that two physically isolated microgrids are formed. A communication network ensures synchronization by exchanging information between the two microgrids.

Figure 11 depicts the convergence time for different sampling periods and different control parameters K . Note that if K is small the convergence time increases considerably. Moreover, while large sampling periods lead to better privacy, as illustrated in Figure 12, they also increase the time to achieve synchronization, which

Figure 10: IEEE 30 nodes benchmark.

affects the system performance. For instance, if it is necessary to reconnect the two microgrids and they are not synchronized, it could cause undesired oscillations or instability of both microgrids.

Similarly, discretionary sampling ensures perfect privacy without affecting synchronization in finite time. However, if the duration of the events that need to remain private is large, the convergence time will also increase. Figure 13 depicts the system behavior for $\tau = 5$, $T = 20$ for both privacy mechanisms.

Note that discretionary sampling prevents the transmission of any important event, however with periodic sampling the presence of the event is still evident.

Figure 11: K vs. t^* for different sampling periods.

5 CONCLUSIONS

Privacy in CPS introduces new challenges due to the interaction of a physical system with transmitted data. In particular, we formulated the inherent privacy level of feedback control systems from the differential privacy perspective by taking advantage of the system uncertainties by using tools from stochastic control systems theory. We found that there is a relationship between the inherent privacy level and the feedback control parameter K which is related to the system capacity to attenuate/amplify noise. We also proposed

Figure 12: Comparison of the proposed privacy mechanisms for the case study of 8 events for node 14. We consider an event was successfully transmitted if the frequency $59.985 \leq x_{14} \leq 60.03$. Increasing the amount of data transmitted decreases the privacy level; however, discretionary sampling guarantees complete privacy.

Figure 13: Frequency excursion of node 14 for $K = 5$ and $\tau = 5$ for both privacy mechanisms.

a novel privacy mechanism that injects the minimum amount of noise in order to achieve a desired (ϵ, δ) -differential privacy. Finally, we introduced a novel privacy definition focused on events and we proposed a noiseless privacy strategy that exploits some properties of multi-agent systems. We also investigated the trade-offs between minimizing data transmission via periodic sampling and with discretionary sampling.

The inherent properties of CPS investigated in this paper are an example of why classic security tools developed for information technology problems need to be reconsidered when we want to apply them to CPS. The science of CPS security and privacy requires the development of new models and tools beyond the direct application of traditional security and privacy technologies.

ACKNOWLEDGMENTS

This work was partially supported by NSF awards CNS-1553683, CNS-1111529, CNS-1228198, and CICI-1547324, and by the Department of Commerce through NIST 70NANB16H019.

REFERENCES

- [1] Gergely Ács and Claude Castelluccia. 2011. I have a dream!(differentially private smart metering). In *Information Hiding*. Springer, 118–132.
- [2] Attiye Alaeddini. 2016. *Observability-Based Approach to Design, Analysis and Optimization of Dynamical Systems*. Ph.D. Dissertation.
- [3] Goong Chen, Guanrong Chen, and Shih-Hsun Hsu. 1995. *Linear stochastic control systems*. Vol. 3. CRC press.
- [4] Jorge Cortés, Geir E Dullerud, Shuo Han, Jerome Le Ny, Sayan Mitra, and George J Pappas. 2016. Differential privacy in control and network systems. In *Decision and Control (CDC), 2016 IEEE 55th Conference on*. IEEE, 4252–4272.
- [5] Florian Dörfler and Francesco Bullo. 2012. Exploring synchronization in complex oscillator networks. In *Proceedings of 51th IEEE Conference on Decision and Control (CDC)*. Maui, HI, USA, 7157–7170.
- [6] Cynthia Dwork, Aaron Roth, and others. 2014. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science* 9, 3–4 (2014), 211–407.
- [7] Stein-Erik Fleten and Erling Pettersen. 2005. Constructing bidding curves for a price-taking retailer in the Norwegian electricity market. *IEEE Transactions on Power Systems* 20, 2 (2005), 701–708.
- [8] Jairo Giraldo, Alvaro Cardenas, Eduardo Mojica-Nava, Nicanor Quijano, and Roy Dong. 2014. Delay and sampling independence of a consensus algorithm and its application to smart grid privacy. In *Proceedings of the 53rd IEEE Conference on Decision and Control (CDC)*. 1389–1394.
- [9] Jairo Giraldo, Alvaro Cárdenas, and Nicanor Quijano. 2016. Integrity Attacks on Real-Time Pricing in Smart Grids: Impact and Countermeasures. *IEEE Transactions on Smart Grid* (2016).
- [10] Jairo Giraldo, Eduardo Mojica-Nava, and Nicanor Quijano. 2014. Tracking of Kuramoto oscillators with input saturation and applications in smart grids. In *Proceedings of the 2014 American Control Conference (ACC)*. 2656–2661.
- [11] Roger A Horn, Noah H Rhee, and So Wasin. 1998. Eigenvalue inequalities and equalities. *Linear Algebra Appl.* 270, 1 (1998), 29–44.
- [12] Zhenqi Huang, Sayan Mitra, and Geir Dullerud. 2012. Differentially private iterative synchronous consensus. In *Proceedings of the 2012 ACM workshop on Privacy in the electronic society*. ACM, 81–90.
- [13] Mark G Lijesen. 2007. The real-time price elasticity of electricity. *Energy economics* 29, 2 (2007), 249–258.
- [14] Andrés Molina-Markham, Prashant Shenoy, Kevin Fu, Emmanuel Cecchet, and David Irwin. 2010. Private memoirs of a smart meter. In *Proceedings of the 2nd ACM workshop on embedded sensing systems for energy-efficiency in building*. ACM, 61–66.
- [15] Reza Olfati-Saber, J. Alex Fax, and Richard M. Murray. 2007. Consensus and cooperation in networked multi-agent systems. *Proc. IEEE* 95, 1 (2007), 215–233.
- [16] Raguathan Raj Rajkumar, Insup Lee, Lui Sha, and John Stankovic. 2010. Cyber-physical systems: the next computing revolution. In *Proceedings of the 47th Design Automation Conference*. ACM, 731–736.
- [17] Mardavij Roozbehani, Munther A Dahleh, and Sanjoy K Mitter. 2012. Volatility of power grids under real-time pricing. *IEEE Transactions on Power Systems* 27, 4 (2012), 1926–1940.
- [18] Kevin P Schneider, Yousu Chen, David P Chassin, Robert Pratt, Dave Engel, and Sandra Thompson. 2008. Modern grid initiative distribution taxonomy final report. *PNNL-18035, Pacific Northwest National Laboratory, Richland, Washington* (2008).
- [19] Victor Solo and Marc Piggott. 2016. What to do about noisy consensus?. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 4836–4839.
- [20] Carol L Stimmel. 2014. *Big data analytics strategies for the smart grid*. CRC Press.
- [21] Rui Tan, Varun Badrinath Krishna, David KY Yau, and Zbigniew Kalbarczyk. 2013. Impact of integrity attacks on real-time pricing in smart grids. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM, 439–450.
- [22] Jing Wu and Tongwen Chen. 2007. Design of networked control systems with packet dropouts. *IEEE Trans. Automat. Control* 52, 7 (2007), 1314–1319.

A PROOF: SYNCHRONIZATION WITH SAMPLED INFORMATION

Firsts, let us introduce the following definition.

Definition A.1. The Perron matrix $\mathcal{P} \in \mathbb{R}^{n \times n}$ of a graph $\mathcal{G}$ is a right stochastic matrix satisfying $\mathcal{P}\mathbf{1} = \mathbf{1}$, such that 1 is one of its eigenvalues. For $\mu_1 > \mu_2 \dots > \mu_n$ its eigenvalues, if $\mu_1 = 1$ and $|\mu_i| < 1$ for $i=2, \dots, N+1$, then the discrete-time system converges to a consensus state [15].

THEOREM A.1. *Consider the synchronization model with distributed sampled measurements in (16). We define $A = K\mathcal{D}_c + \mathcal{Q}_p$ and $B = K\mathcal{B}_c$ and we obtain the discretized model in Eq. (17), where $P = \Phi + \Gamma$. If P is a Perron matrix then the states exponentially converge to the reference state x^r independent of the sampling period.*

Proof

Recall that $\Phi = e^{-A\tau}$ and $\Gamma = \int_0^\tau e^{-As} ds B = A^{-1} (I - e^{-A\tau}) B$. Let us define $\mathcal{L}_d$ as a matrix associated to P by

$$P = I - A^{-1} \mathcal{L}_d. \quad (21)$$

First, we have to show that $\mathcal{L}_d$ is a Laplacian matrix. As $P = \Phi + \Gamma$, from (21) we obtain

$$L_d = A(I - P) = A - Ae^{-A\tau} + e^{-A\tau} B - B,$$

which leads to

$$L_d = K\mathcal{D}_c + \mathcal{Q}_p - K\mathcal{D}_c e^{-A\tau} - \mathcal{Q}_p e^{-A\tau} + e^{-A\tau} K\mathcal{B}_c - K\mathcal{B}_c. \quad (22)$$

Recall that $\mathcal{Q}_c = \mathcal{D}_c + \mathcal{B}_c$ and $A - B = \mathcal{L}_p + K\mathcal{L}_c = L_T$. Since $\mathcal{G}_p$ is an undirected graph, then A is symmetric and we can group the terms in (22) such that $L_d = (I - e^{-A\tau}) L_T$.

A is positive definite and diagonally dominant, such that $I - e^{-A\tau}$ is symmetric and positive. L_T is a Laplacian matrix and L_d is then also a Laplacian matrix. Now, we have that

$$P = I - (K\mathcal{D}_c + \mathcal{Q}_p)^{-1} (I - e^{-A\tau}) L_T.$$

Let λ_A^{max} be the maximum eigenvalue of $K\mathcal{D}_c + \mathcal{Q}_p$, η^{max} the maximum eigenvalue of $\mathcal{Q}_p$ and d^{max} the maximum value of the diagonal matrix $K\mathcal{D}_c$. Due to the symmetry of both terms, using the Weyl's inequality [11] we have that $\lambda_A^{max} \leq \eta^{max} + d^{max}$. Therefore, the matrix P is bounded by

$$P \leq I - \frac{(I - e^{-A\delta}) L_T}{\lambda_A^{max}}.$$

Note that the diagonal elements of the Laplacian matrix L_T / λ_T^{max} are positive and less than one and the term $I - e^{-L_T \tau} \leq I$ for any $\tau > 0$, such that P is a Perron matrix independently of τ . Consequently, the states reach a consensus, which corresponds to x^r . ■