Differential Privacy for Network Identification

Vaibhav Katewa, Aranya Chakrabortty, *Senior Member, IEEE*, and Vijay Gupta, *Senior Member, IEEE*

***Abstract*—We consider a multiagent linear time-invariant system whose dynamical model may change from one disturbance event to another. The system is monitored by a control center that collects output measurements from the agents after every event and estimates the eigenvalues of the model to keep track of any adverse impact of the disturbance on its spectral characteristics. Sharing measurements in this way, however, can be susceptible to privacy breaches. If an intruder gains access to these measurements, she may estimate the values of sensitive model parameters and launch more severe attacks. To prevent this, we employ a differential privacy framework by which agents can add synthetic noise to their measurements before sending them to the control center. The noise is designed carefully by characterizing the sensitivity of the system so that it limits the intruder from inferring any incremental change in the sensitive parameters, thereby protecting their privacy. Our numerical results show that the proposed design results in marginal degradation in eigenvalue estimation when compared to the error incurred by the intruder in identifying the sensitive parameters.**

***Index Terms*—Cyber-security, differential privacy (DP), monitoring, multiagent system, topology identification.**

I. Introduction

DISTRIBUTED dynamical systems are pervasive in today's society. Examples of such systems include electric power systems, transportation systems, robotic networks, teams of unmanned air vehicles, communication networks, and so on [1]. The underlying dynamic model of these systems may change over time due to variations in operating conditions such as the number of agents, the interaction rules between the agents, or the underlying model parameters. Thus, a system administrator needs to continuously monitor the system to manage it more effectively. One possible way to perform this for a linear time-invariant system is to collect the measured outputs of the agents at a control center after a disturbance event, and estimate the eigenvalues of the model to check if the disturbance caused any serious degradation to its spectral properties. Depending on the result of this estimation, the administrator may decide to retune the controllers in each agent so that the system regains a desired closed-loop spectrum. This, in fact, is commonly used in electric power systems, where measurements from geographically dispersed sensors are used for estimating eigenvalues of the oscillation modes [2].

Although these periodic checks help in keeping track of the health of the system, they also open up possibilities for severe data privacy breaches. The state-space model of a multiagent system, for example, is often parameterized by parameters associated with each agent that may contain sensitive information about the agent behavior, and should not be revealed to external entities. The model of a consensus network, for instance, is parametrized by the weights associated with its nodes and the edges that represent the coupling strength between the agents [3]. Agents in the network may not prefer these coupling strengths to be known to others. A common example, again, is the electric power system where these weights may reflect the loading profiles of various utility companies that they do not want other companies to know. Another example is a multiagent linear quadratic regulator (LQR) control problem wherein the agents are dynamically coupled but have individual cost matrices, resulting in an overall decoupled cost function. In this case, the optimal control inputs depend on the state and input cost weighing matrices of agents. In many applications, such as biological systems, it has been shown that the cost matrices of an agent (human) represent the intent of a human (see [4], [5], and references therein) that she would like to keep private. In many economic applications, such as price determination, welfare planning, and resource allocation, the agents solve an LQR problem to obtain optimal results [6]. In such cases, the cost matrices represent the pricing and welfare strategies of the agents, which they may not want to reveal to their competitors.

In our problem setting, if an intruder hacks into the control center and gains access to the output measurements, she may use these measurements to estimate the model of the system and infer the sensitive parameters of the system from this estimate [7], [8]. This can cause a major privacy breach, enabling the intruder in gaining critical system-level information that she can further use to plan a more severe attack on the system. Therefore, protecting the privacy of these sensitive model parameters is a crucial task.

To address this issue, in this paper we present a differential privacy (DP) mechanism by which agents can add synthetic noise to their measurements before sending them to the control

Manuscript received February 11, 2019; revised May 10, 2019; accepted May 21, 2019. Date of publication June 11, 2019; date of current version March 18, 2020. This work was supported by the NSF under Grant CNS-1544871, Grant CNS-1739295, and Grant CNS-1544724. Recommended by Associate Editor S. Dey. *(Corresponding author: Vaibhav Katewa.)*

V. Katewa is with the Department of Mechanical Engineering, University of California at Riverside, Riverside, CA 92507 USA (e-mail: vkatewa@engr.ucr.edu).

A. Chakrabortty is with the Department of Electrical and Computer Engineering, North Carolina State University, Raleigh, NC 27606 USA (e-mail: achakra2@ncsu.edu).

V. Gupta is with the Department of Electrical Engineering, University of Notre Dame, Notre Dame, IN 46556 USA (e-mail: vgupta2@nd.edu).

Digital Object Identifier 10.1109/TCNS.2019.2922169

center. DP, as proposed by Dwork in [9], has been used widely in computer science and statistics. The guarantees provided by a DP mechanism are relative (not absolute), implying that there is no significant privacy loss of the sensitive parameters when they are changed within some specified thresholds, regardless of any potential side information that the intruder might have [9], [10]. In the setting of this paper, ensuring DP of the system parameters means that for any two sets of parameters that are "not very different," the outputs of the corresponding linear dynamical systems will also be "statistically similar (within a multiplicative factor)." To ensure this property, each agent adds a synthetic noise to its state measurements in such a way that the intruder cannot identify a "differential change" in the parameters from the noisy measurements, thereby making them private. Of course, the addition of noise will also adversely affect the eigenvalue estimation, which is the goal of the control center. However, estimating the characteristic polynomial of the system model, from which the eigenvalues are computed, is a pseudolinear regression problem, whereas identifying the system model or any sensitive parameter embedded in it is a nonlinear problem. Pseudolinear regression methods are much more robust to noise than their nonlinear counterparts [11]. The premise of our design rests upon this robustness property. Accordingly, our numerical results show that the proposed DP design results in marginal degradation in eigenvalue estimation when compared to the error incurred by the intruder in identifying the sensitive parameters.

An alternative way to guarantee privacy would be to encrypt the measurements with decryption keys available only to the control center. However, such schemes may not be feasible in all scenarios. First, it may not be possible to encrypt the measurements in real time. For instance, the phasor measurement units (PMUs) in power network do not have encryption capabilities. Further, encryption is typically computationally expensive, introduces delays, and can be breached if other system protocols are not strong. Second, encryption can protect the privacy in case the intruder is eavesdropping on the measurements, but not in the case when it hacks into the control center.

Several recent papers have presented privacy mechanisms for dynamical systems. In [12], the DP framework was extended to discrete-time dynamical systems, which we have also been used in this paper. In [13] and [14], the authors propose a DP mechanism to keep the initial states private for the consensus problem. In [15], a DP mechanism is developed to keep a reference trajectory private for a general distributed control system. In [16], a privacy mechanism involving careful noise addition and removal is presented to keep the initial state private and achieve exact consensus. Some papers have addressed privacy issues in optimization problems. In [7], [8] and [17], [18], the authors present noisy update algorithms to protect sensitive constraints, states and cost functions, respectively. Further, [19] and [20] present mechanisms for privately solving optimization problems with linear and piecewise affine objectives, respectively. In [21], the authors present a stochastic gradient algorithm to solve a convex optimization problem wherein the noisy updates preserve DP. Differentially private linear quadratic Gaussian (LQG) control is studied in [22]–[24], where the goal is to keep the states private over time. While these works present privacy mechanisms to protect the initial conditions, states, inputs, reference trajectories, and cost functions, mechanisms to protect the *parameters embedded inside the system dynamics* are lacking. In this paper, we develop a privacy mechanism for this fundamentally different problem. There have been some works [25], [26] on the privacy of parameters of transfer functions. However, the mapping in [25] between the parameters and output is linear and [26] considers a simple case of single-input and single-output transfer functions. In contrast, we consider a more general problem of privacy of parameters embedded in the state matrix. The mapping from initial condition, states, input, or reference trajectory to the output is linear whereas the mapping from the state matrix (and the parameters) to the output is nonlinear. Characterizing this nonlinear mapping poses additional challenges for our problem.

There have also been many studies on using DP for static graphs while releasing queries such as counts of different types of subgraphs [27], degree distribution [28], [29], and spectral properties like eigenvalues and eigenvectors [30]. In contrast, we consider privacy in dynamical systems and aim to make parameters related to the system dynamics private. Our privacy mechanism depends fundamentally on the dynamics of the system. Finally, the authors in [31] study system identification when a subset of sensors are under attack, and characterize the class of systems that are indistinguishable under such attacks. In contrast, our problem setup and approach are different. We explicitly design the privacy noise and add it to all the outputs of the system. Moreover, we are concerned with characterizing the estimation error of the system matrix and its eigenvalues.

A conference version of this paper has been presented in [32]. The DP mechanism in [32], however, is specific to protecting the topology of first-order consensus networks only, whereas the approaches in this paper apply to any general class of linear time-invariant systems.

The main contributions of the paper are as follows.

1) We study a privacy problem for a multiagent linear time-invariant (LTI) system to protect the sensitive parameters embedded in the dynamics of the agents. We show that the parameters can be leaked as the agents share their measurements with a control center. To prevent such privacy breaches, we design a DP mechanism that adds suitable noises to the measurements before sharing them. We determine the noise level by obtaining an analytical upper bound on the sensitivity of the system. The computation of bound involves a nonlinear mapping from the parameters to the measurements, which poses additional challenges.
2) We present two important and diverse applications of our framework, namely, a mode estimation in power systems and an LQR control for multiagent systems. The DP mechanism protects the privacy of the topology of the power system network in the first case, and that of the cost matrix used in the quadratic cost function in the second. These examples highlight the practical applicability of our proposed approach.
3) We consider a specific scenario where the goal of the control center is to estimate the eigenvalues of the system.

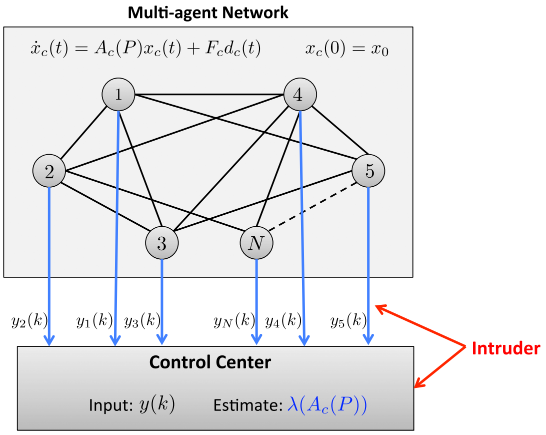

Fig. 1. Dynamical system architecture.

Using numerical simulations, we show that the privacy mechanism has marginal effect on eigenvalue estimation performance when compared to the parameter identification error incurred by the intruder. Thus, for this specific scenario, the proposed mechanism achieves the desired privacy goal with marginal performance degradation.

The paper is organized as follows. In Section II, we introduce the problem setup and present the eigenvalue and parameter identification procedures. Section III presents two application that fit our setup. In Section IV, we present a noise-adding DP mechanism and analyze its effect on the parameter and eigenvalue estimation methods. We design the privacy noise by deriving an analytical bound on the system sensitivity. In Section V, using the power network second-order consensus example, we numerically illustrate the effect of the DP mechanism. Section VI concludes the paper.

Mathematical Notation: $\|.\|_p$ and $\|.\|_F$ denote the (induced) p-norm and Frobenius norm of a vector/matrix, respectively. Let $\{y(k)\}_{k\geq 0}$ denote a sequence/trajectory. The truncated version of y until time T is donated by $y[0:T]$. Without loss of generality, we also treat a truncated sequence as a vector of appropriate dimension. For a square matrix A, $\lambda(A)$ denotes its set of eigenvalues, $\mu(A) = \max\{\text{Real}(\lambda(A))\}$ denotes its spectral abscissa, and $\kappa(A) = ||A||_2||A^{-1}||_2$ denotes its condition number. I_N denotes the $N \times N$ identity matrix, $\mathbf{1}_N = [1, 1, \ldots, 1]^T \in \mathbb{R}^N$, $\mathbf{0}_N = [0, 0, \ldots, 0]^T \in \mathbb{R}^N$, and $\mathbf{0}_{N\times M} \in \mathbb{R}^{N\times M}$ denote an all-zero matrix. The Kronecker product of two matrices is denoted by $\otimes$. $\delta(t)$ and $\delta(k)$ denote the Dirac and Kronecker delta functions, respectively. $Lap(0, b)^N$ denotes an N-dimensional Laplace distribution with i.i.d. components, each with p.d.f. $f(x) = \frac{1}{2b}e^{\frac{-|x|}{b}}$. $Re(.)$ and $Im(.)$ denote the real and imaginary parts of a complex number, respectively.

II. Problem Setup

We consider a networked system consisting of N linear dynamical agents as shown in Fig. 1. The coupling between the states of the agents can either result from the inherent physical connection between them (for example, transmission lines connecting generators in a power system) or from communication of states among each other (for example, in a team of robots trying to achieve consensus). The evolution of the whole system is represented as the following continuous-time LTI system

$$\dot{x}_c(t) = A_c(P)x_c(t) + F_c d_c(t) \quad t \geq 0,\ P \in \mathcal{P} \qquad (1)$$

where $x_c \in \mathbb{R}^n$ is a vector of aggregated states of the individual agents $x_i \in \mathbb{R}^{n_i}$, $i \in \mathcal{N} \triangleq \{1, \ldots, N\}$ with $\sum_{i=1}^N n_i = n$. The term $d_c(t) \in \mathbb{R}$ denotes a disturbance which can occur due to a fault in the network. Since faults are isolated and do not occur simultaneously, we assume $d_c(t)$ to be a scalar. The state matrix A_c can either represent an open-loop system or a closed-loop system where a possible state-feedback or output-feedback gain matrix may already be embedded inside A_c. Let $P = \{P_1, P_2, \ldots, P_N\}$ denote a collection of parameters P_i of individual agent dynamics that they wish to keep private. For instance, P_i may contain sensitive parameters that can reveal the coupling structure of agent i with other agents or its control preferences. We will illustrate this later using some examples. $\mathcal{P}$ denotes the set of all possible parameters P. Without loss of generality, we treat the sets P_i and P as matrices and/or vectors, depending on the application. The output of the system is given by

$$y_c(t) = Cx_c(t) \qquad (2)$$

where y_c is a vector of aggregated outputs $\{y_i\}_{i=1}^N$ of all agents. We make the following assumptions regarding (1) and (2).

A1: The state matrix $A_c(P)$ is stable or marginally stable for all $P \in \mathcal{P}$. Further, $\lim_{t\to\infty} x_c(t)$ exists and is finite.

A2: The disturbance $d_c(t)$ vanishes at $t = t_0 \triangleq i_0 T_s$ for some $i_0 \in \mathbb{N}$, i.e., $d_c(t) = 0$ for $t > t_0$. Further, $d_c(t_0) = d_0$.

A3: The system is fully observable, i.e., matrix C is square and invertible ($y_c \in \mathbb{R}^n$).

A4: The state matrix $A_c(P)$ is diagonalizable for all $P \in \mathcal{P}$.

A5: The steady-state value $\lim_{t\to\infty} x_c(t)$ is independent of the sensitive parameters P.

Stability Assumption **A1** is standard for physical systems. It implies that $Re(\lambda_i(A_c(P))) \leq 0$ for $i \in \{1, 2, \ldots, n\}$ and in case the system is marginally stable, there is only a single pole on the imaginary axis which is located at the origin. Assumption **A2** captures sudden faults in the system which vanish after some transient period. Also, the time period between consecutive faults is typically large and the transients introduced by a fault settle down during this time period. Our setup and analysis focus on one such time period, and therefore, we assume that there is a single disturbance over time. Assumption **A3** is needed for the identification of $A_c(P)$ using the measurements. This represents a worst case scenario in which an intruder can infer the sensitive parameters P, and this results in a privacy breach (see Section II-A). Thus, our goal is to prevent such privacy breaches even in this worst case scenario.

Assumption **A4** is required later in the paper for the analysis of the privacy mechanism (see Section IV-B). Further, it is trivially satisfied when the eigenvalues of $A_c(P)$ are distinct. Assumption **A5** is trivially satisfied for stable systems since the steady state is zero. For marginally stable systems, this assumption guarantees that the privacy noise level remains bounded. The details are provided in Section IV-B.

The agents periodically sample their outputs at discrete time instants with a sampling time $T_s > 0$. The sampled output is denoted as

$$y(k) \triangleq y_c(kT_s) = Cx_c(kT_s),\ k \geq i_0 \tag{3}$$

where $y \in \mathbb{R}^n$. These measurements are sent to a control center as shown in Fig. 1. Since the measurements are available at discrete time instants, we convert the continuous-time model (1) to discrete time. To describe the response for $t \geq t_0$, we treat $d_c(t) = d_0\delta(t - t_0)$ as an impulse (see Assumption **A2**) and use the impulse invariant transform for this purpose. The state evolution for (1) at time instants $t = kT_s, \quad k = i_0 + 1, i_0 + 2, \ldots$ can be described by

$$x(k{+}1) = A_d(P)x(k) + F_d(P)d(k),\ k \geq i_0 \tag{4}$$

where $x(k) \triangleq x_c(kT_s),\ \ A_d(P) \triangleq e^{A_c(P)T_s}, F_d(P) \triangleq A_d(P)$ F_c, and $d(k)$ is a shifted impulse in discrete time, i.e., $d(k) = d_0\delta(k - i_0)$. The output equation becomes $y(k) = Cx(k)$, $k \geq i_0$. We make the following assumptions regarding the discrete-time system.

A6: The pair $(A_d(P), F_d(P))$ is controllable for all $P \in \mathcal{P}$.

A7: The sampling time period satisfies the following property:

$$T_s < \bar{T}_s \triangleq \frac{\pi}{\sup\limits_{P\in\mathcal{P},\, i\in\{1,2,\ldots,n\}} |Im(\lambda_i(A_c(P)))|}. \tag{5}$$

Assumptions **A6** and **A7** are required for accurate identification of private parameters and to avoid aliasing, respectively (see Sections II-A and II-B).

Sharing the measurements with the control center can enable intruders to gain access to $y(k)$. An intruder can use the measurements to get information about the sensitive parameters P embedded in $A_c(P)$, thereby resulting in a privacy breach (see Section II-A for details). The primary goal of the paper is to develop a mechanism that protects the privacy of the sensitive parameters P. Further, we consider a specific scenario where the aim of the control center is to estimate the eigenvalues of the system ($A_c(P)$) to verify that it is functioning as desired. We emphasize that while the objectives of the intruder and control center can vary depending on the application, the design of the privacy mechanism is independent of these objectives.

We emphasize that the objectives of the intruder and the control center are different, although they have access to the same measurements. This is a typical scenario in power networks where it is important to estimate the characteristics of the global dynamics, such as the eigenvalues, and there is rarely a need to identify the global topology. On the other hand, the intruder needs to identify the topology to infer the sensitive parameters, which can be potentially used to perform attacks. The sole knowledge of the eigenvalues will not be sufficient toward this aim. We explain the parameter and eigenvalue estimation procedures in the next sections.

A. Privacy Issues in Dynamical Systems

An intruder may have information about the system that is obtained from external resources. Towards this, we make the following assumption.

A8: The intruder knows $x(t_0)$, F_c, the disturbance $d(t)$, and sampling time period T_s. Further, it may also have information about some nonsensitive parameters of the system.

This assumption implies that the intruder knows that the system is excited by a disturbance that vanishes at $t = t_0$ and how the disturbance affects the system. Also, some nonsensitive parameters of the system are known to the intruder. Collectively, all the information about the system except the sensitive parameters P is referred to as *side information* or auxiliary information.

Remark 1 (Side information): Assumption **A8** is not restrictive, as in reality, there always exists a possibility that side information about the system operation is known to an intruder. In fact, one of the main motivations behind the DP framework (which we will introduce shortly) is to develop a privacy mechanism that abstracts away from arbitrary side information that the intruder might possess [9] (also see Remark 4). Moreover, Assumption **A8** represents the worst case scenario. The capabilities of the intruder will be further limited in case she does not have any side information. ■

Since $d(k) = d_0\delta(k - i_0)$, the output of (4) can be written as

$$y(k) = CA_d(P)^{k-i_0}(x(i_0) + F_c d_0),\ k > i_0. \tag{6}$$

The aim of the intruder is to infer the sensitive parameters P using $y(k)$ and the side information. Since $y(k)$ depends on P in a nonlinear manner, a natural choice for the intruder will be to use nonlinear least squares (NLS) to identify A_d and infer P from that estimate. The method is described as follows.

Identification of Sensitive Parameters

P1) *Using (6), solve the following NLS problem:*

$$\hat{A}_{d,T}(\hat{P}_T) = \arg\min_{Z\in\mathbb{R}^{n\times n}} \sum_{k=i_0+1}^{T+i_0} \left\| y(k) - CZ^{k-i_0}(x(i_0) + F_c d_0) \right\|_2^2 \tag{7}$$

where T is the identification time horizon.

P2) *Estimate the continuous-time state matrix from the NLS solution as $\hat{A}_{c,T}(\hat{P}_T) = \frac{1}{T_s} log(\hat{A}_{d,T}(\hat{P}_T))$.*

P3) *Extract the sensitive parameter $\hat{P}_T$ from $\hat{A}_{c,T}(\hat{P}_T)$.*

As (4) is controllable (Assumption **A6**) and fully observable (Assumption **A3**), the discrete-time state matrix $A_d(P)$ can be uniquely and accurately identified by the intruder in step **P1** [33]. Since condition (5) on the sampling time guarantees that there is no aliasing, the intruder can thereafter use A_d to compute the continuous-time state matrix $A_c(P)$ in step **P2**. Finally, from $A_c(P)$, the intruder can identify the parameters P. The subscript T in the estimates indicate that the accuracy of the estimation is dependent on the choice of the estimation horizon. For sufficiently large values of T the intruder can obtain a fairly reliable estimate of P. This is clearly undesirable, and calls in for a privacy mechanism.

Remark 2 (Non-identifiability of the sensitive parameters): In some cases, it may be possible that the mapping from P to $A_c(P)$ is many to one. Also, the intruder may not have enough side information to infer P from $A_c(P)$. In all such cases, the intruder will not be able to uniquely identify the parameters from the measurements. ■

B. Eigenvalue Identification by Control Center

Over time or following a disturbance, the system characteristics and the parameters P may change, as a result of which A_c may change. The control center may not know the new value of A_c. To monitor if this change adversely affects the system dynamics, the control center collects $y(t)$ after the disturbance $d_c(t)$ and estimates the eigenvalues of the new state matrix $A_c(P)$. We assume that the control center knows the time instant when the disturbance $d_c(t)$ vanishes[1] (i.e., t_0), and uses the measurements collected after $t = t_0$ to perform this estimation. The eigenvalues of the discrete-time model are related to that of the continuous-time model as

$$\lambda(A_d(P)) = e^{T_s \lambda(A_c(P))}. \tag{8}$$

The easiest way to estimate $\lambda(A_c(P))$ is to first estimate $\lambda(A_d(P))$ using $y(k)$, and then use (8) to compute $\lambda(A_c(P))$. Assumptions **A3** and **A6** guarantee accurate identification of $\lambda(A_d(P))$ [33]. Further, Assumption **A7** prevents aliasing and enables accurate identification of $\lambda(A_c(P))$. Next, we describe the identification procedure in detail.

Since $d(k)$ is a vanishing signal, we treat it as an impulse input, and denote the transfer function matrix from d to y as

$$H(z) = C(zI_n - A_d(P))^{-1} F_d \quad \text{with} \tag{9}$$

$$[H]_i = \frac{\sum_{k=1}^{n} b_i^{(k)} z^{-k}}{1 + \sum_{k=1}^{n} a^{(k)} z^{-k}} \tag{10}$$

where $a_d(z) \triangleq 1 + \sum_{k=1}^{n} a^{(k)} z^{-k}$ is the characteristic polynomial of $A_d(P)$. Let $a = [a^{(1)}, a^{(2)}, \ldots, a^{(n)}]^T$, $b_i = [b_i^{(1)}, b_i^{(2)}, \ldots, b_i^{(n)}]^T$. The outputs can be written as

$$y(k) = \varphi^T(k)\theta, \quad k > i_0 \quad \text{where} \tag{11}$$

$$\varphi(k) = \begin{bmatrix} -y^T(k-1) \\ -y^T(k-2) \\ \vdots \\ -y^T(k-n) \\ I_n \otimes \begin{bmatrix} d(k-1) \\ d(k-2) \\ \vdots \\ d(k-n) \end{bmatrix} \end{bmatrix}_{(n+n^2)\times n} \qquad \theta = \begin{bmatrix} a \\ b_1 \\ b_2 \\ \vdots \\ b_n \end{bmatrix}_{(n+n^2)\times 1}. \tag{12}$$

The data vector φ is constructed from the output measurements and the impulse disturbance. Thereafter, θ is estimated using the instrumental-variable (IV) method[2] as

$$\hat{\theta}_T = sol \left\{ \frac{1}{T} \sum_{k=i_0+1}^{T+i_0} \zeta(k)[y(k) - \varphi(k)^T \theta] = 0 \right\} \tag{13}$$

where sol denotes the solution of an equation, T is the time horizon for identification, and $\zeta(k)$ are appropriately chosen instruments that are correlated with $\varphi(k)$. We skip the details of the IV method for brevity, and refer the interested reader to [11]. Let $\hat{a}_T$ denote the estimate of a that is extracted from $\hat{\theta}_T$. The eigenvalue identification procedure can be summarized as follows.

Eigenvalue Identification Procedure

E1) *Obtain the estimate $\hat{\theta}_T$ using the IV method in (13).*

E2) *Extract $\hat{a}_T$ from $\hat{\theta}_T$ using its structure given in (12).*

E3) *Estimate $\hat{\lambda}(A_d(P))$ by solving the roots of the characteristic polynomial with coefficients contained in $\hat{a}_T$.*

E4) *Compute $\hat{\lambda}(A_c(P)) = \frac{1}{T_s}\log(\hat{\lambda}(A_d(P)))$.*

As pointed out earlier, Assumptions **A3**, **A6**, and **A7** guarantee accurate identification, i.e., $\hat{\lambda}(A_c(P)) = \lambda(A_c(P))$.

Besides the IV method for eigenvalue estimation and the NLS method for parameter estimation, one can also use other advanced methods for these problems. However, we focus on these methods since finding the best method is not the primary focus of the paper.

To protect the privacy of P, we next develop a noise-adding privacy mechanism using the notion of DP. Before presenting the details of the privacy mechanism in Section IV, we present two applications that are suitable to our framework.

III. Application Examples

A. Wide-Area Monitoring in Power Systems

Consider an electric power system network consisting of N generators. Let ω_i, $E_i \angle \delta_i$, M_i, and D_i denote the speed, internal voltage phasor, inertia, and damping factor of the ith generator, respectively. Let x_{ij} denote the reactance of the line-connecting generators i and j. Assume the line resistance to be negligible. Considering the classical model for the generators, assume E_i to be a constant. The small-signal electromechanical dynamics of the network in the so-called *Kron reduced* form can then be modeled as [34]

$$\dot{x}_c(t) = \underbrace{\begin{bmatrix} \mathbf{0}_{N\times N} & I_N \\ -M^{-1}L & -M^{-1}D \end{bmatrix}}_{A_c(L)} x_c(t) + \underbrace{\begin{bmatrix} \mathbf{0}_N \\ M^{-1}\bar{F}\mathbf{1}_N \end{bmatrix}}_{F_c} d_c(t) \tag{14}$$

where $x_c = [\Delta\delta_1, \ldots, \Delta\delta_N, \Delta\omega_1, \ldots, \Delta\omega_N] \in \mathbb{R}^{2N}$ consists of the deviation of the phase angles and speeds of all generators from their equilibrium; $M = \text{diag}(M_1, \ldots, M_N)$ and $D = \text{diag}(D_1, \ldots, D_N)$ denote the diagonal matrices containing

[1] This assumption is typically valid in systems that are continuously monitored in a distributed manner. For instance, in electrical power system networks, relays and circuit breakers sense when a disturbance occurs or vanishes by monitoring the transients introduced in the currents in the network. This information can be easily sent to the control center via communication links.

[2] In the presence of privacy noise in the outputs (see Section IV), (11) corresponds to the output error model. Estimation of θ for this model is a pseudolinear problem for which the IV method is a standard identification technique [11].

generator inertias and dampings; and $L \in \mathbb{R}^{N \times N}$ is a symmetric Laplacian matrix with $L_{ij} = (E_i E_j)/x_{ij} \cos(\delta_{i0} - \delta_{j0})$, $i \neq j$, and $L_{ii} = -\sum_{\substack{k=1 \\ k \neq i}}^{N} L_{ik}$. The elements of L can be viewed as the edge weights of the underlying graph. For monitoring purposes, the measurements from the power system are collected by PMUs according to (2) and transmitted to a control center. The model (14) can be viewed as a second-order consensus where: 1) $\lim_{t \to \infty} |\Delta\delta_i(t) - \Delta\delta_j(t)| = 0$ for $i, j \in \mathcal{N}$ and 2) $\lim_{t \to \infty} |\Delta\omega_i(t)| = 0$ for all $i \in \mathcal{N}$ indicating that all generators must synchronize with each other over time at the common synchronous frequency of 60 Hz.

These types of second-order dynamics are also used in modeling transportation systems and robot motions [35]. A standard result in consensus literature is that synchronization is achieved if the underlying network graph is connected [1]. In this case, all eigenvalues of $-L$ have negative real parts except for one, which is zero. As a result, the state matrix $A_c(L)$ in (14) has one eigenvalue at zero, one real eigenvalue that is negative, and the remaining all complex conjugate pairs of eigenvalues with negative real parts [36]. Therefore, it satisfies Assumption **A1**. We will later derive the explicit expression of the consensus value reached by the dynamical system (14). As indicated above, the elements of L contain sensitive information about the generator's internal voltages, the line parameters, as well as the loads. These information should not be revealed to extraneous sources [32]. Comparing (14) with the general dynamical system (1), for this example, we get $P_i = \{L_{ij} : j \in \mathcal{N}_i\}$, $P = L$, and $A_c(P) = A_c(L)$.

B. LQR Control

Consider an LTI system given as

$$\dot{x}_c(t) = A x_c(t) + B u(t) + F_c d(t) \tag{15}$$

where $x_c = [x_1, \ldots, x_N]^T \in \mathbb{R}^n$ represent the state vector of N coupled subsystems and $u(t) = [u_1, \ldots, u_N]^T \in \mathbb{R}^m$ is the control input vector. The control objective is to design a state-feedback law $u(t) = -K x_c(t)$ to minimize the following infinite horizon quadratic cost function:

$$J(x, u) = \int_0^{\infty} (x_c^T(t) Q x_c(t) + u^T(t) R u(t)) dt \tag{16}$$

where $Q = \text{diag}(Q_1, \ldots, Q_N) \in \mathbb{R}^{n \times n}$ and $R = \text{diag}(R_1, \ldots, R_N) \in \mathbb{R}^{m \times m}$ with $Q_i \in \mathbb{R}^{n_i \times n_i} \geq 0$, $R_i \in \mathbb{R}^{m_i \times m_i} > 0$ with $\sum_{i=1}^{N} m_i = m$. The cost matrices (Q_i, R_i) may contain private information about the control strategies of the agents. For example, in biological systems, the cost matrices can represent human intent [4], [5] and in economic applications, they can represent pricing and welfare strategies [6]. This is the classical LQR optimal control problem for which we make the following assumptions.

A9.1: The pair (A, B) is controllable.

A9.2: The input cost matrix $R = I_m$.

A9.3: Rank(B) = n.

Assumption **A9.1** is standard in LQR problems. Assumption **A9.2** means that the input cost matrix R is specified for the system and the agents wish to keep only the state cost matrix Q private. Further, without loss of generality, we assume R to be an identity (otherwise, one can always do a coordinate transformation s.t. $R = I_m$). This is a standard assumption in most of the inverse control problems [37], [38] and also simplifies the analysis (c.f. Lemma 7). Assumption **A9.3** ensures that the cost matrix Q can be uniquely identified from the optimal control gain defined as follows (see [38, Theorem 7]).

The optimal control law that minimizes the cost in (16) depends on Q, and is given by

$$u(t) = -K(Q) x_c(t), \; K(Q) = R^{-1} B^T V(Q) \tag{17}$$

where $V(Q)$ is the unique positive definite solution of the following Riccati equation:

$$A^T V(Q) + V(Q) A - V(Q) B R^{-1} B^T V(Q) + Q = 0. \tag{18}$$

Since (A, B) is controllable, the closed loop matrix $A - BK(Q)$ is stable (cf. Assumption **A1**). Comparing the closed-loop evolution of the LQR system with the general dynamical system (1), we get $P_i = Q_i, P = \text{diag}(P_i) = Q$, and $A_c(Q) = A - BK(Q)$.

The above-mentioned examples illustrate the generality of our problem formulation for a wide range of practical applications. Next, we present a mechanism to protect the privacy of P.

IV. DP Mechanism

In this paper, we use the notion of DP to develop the privacy mechanism. We first present an intuitive explanation of DP. In order to prevent the intruder from accurately identifying the sensitive parameters, the agents add noise to the measurements before transmitting them to the control center. The privacy noise ensures the following *"differential"* property—if the sensitive parameters are *"changed within some specified limits,"* then the corresponding measurements appear *"probabilistically similar (within a mutiplicative factor)"* to the intruder. In other words, the DP noise makes it difficult for the intruder to distinguish between two "adjacent" parameters with a high confidence level, thereby preserving their privacy. Next, we provide formal definitions of DP.

Definition 1 (Adjacency): Given $\beta \geq 0$, two parameters P and P' (which in our case are matrices) are β-adjacent (denoted by $adj(\beta)$) if

$$\|P - P'\|_2 \leq \beta. \tag{19}$$

Remark 3 (Generalized adjacency): In the DP definition for static databases [9] and for dynamical systems [12], adjacency is defined with respect to the change of data or input of one agent while keeping that of other agents unchanged. In contrast, our definition of adjacency is more general and allows changes in the parameters of one or more agents. ■

As mentioned before, the agents add noise to the measurements according to the following DP mechanism

$$\mathcal{M} : \tilde{y}(k) = y(k) + w(k) \tag{20}$$

where $w(k) \in \mathbb{R}^n$ is the noise. We will specify the properties of the noise w in Section IV-A.

Definition 2 (DP): Let $\tilde{y}$ and $\tilde{y}'$ denote the noisy measurements corresponding to any two β-adjacent parameters P and $P^{'}$. The mechanism $\mathcal{M}$ in (20) is ϵ-differentially private until time T if for all measurable $S \subset \mathbb{R}^{n(T+1)}$

$$\mathbb{P}[\tilde{y}[0:T] \in S] \leq e^{\epsilon}\mathbb{P}[\tilde{y}'[0:T] \in S] \tag{21}$$

where $\epsilon > 0$ is the privacy level. The definition says that if the parameter changes from P to $P^{'}$ that is β-adjacent to P, then the corresponding measurement probabilities change only within a factor of e^{ϵ}. A smaller value of ϵ implies a higher level of privacy and vice versa.

Remark 4 (Differential versus absolute privacy): We would like to emphasize that the DP mechanism does not guarantee absolute privacy of the parameters. In some scenarios, the intruder may be able to obtain an extensive amount of side information and it can identify some of the sensitive parameters even without using the measurements. The only guarantee DP provides is that there will only be a marginal privacy loss (determined by ϵ) due to any changes in the sensitive parameters of the agents within a specified limit (determined by β). Thus, it abstracts away from any possible side information that the intruder may have [10]. ■

Remark 5 (DP via parameter perturbation): An alternative approach is to ensure that DP would perturb (add noise to) the parameters P directly, instead of adding noise to the measurements (for a similar approach for function perturbations, see [18]). In this case, since the measurements are a function of the perturbed parameters, the postprocessing property of DP [12] will ensure that the privacy of the parameters is guaranteed to the intruder which has access to the measurements. However, it is not always feasible to randomly change the parameters that represent physical components in the system. A possible option to circumvent this issue would be to determine the noise in the measurements resulting from potential perturbation of the parameters, and then, add that specific noise to the measurements. However, it is difficult to determine how the nonlinear mapping from P to y transforms the noise, and thus, obtaining the statistics of the measurement noise is not trivial in this case. ■

With DP mechanism, step $\mathbf{P1}$ of the parameter identification procedure can be rewritten as [for comparison, see (7)]

$$\begin{aligned}&\hat{A}_{d,T}(\hat{P}_T)\\&= \underset{Z \in \mathbb{R}^{n\times n}}{\arg\min} \sum_{k=i_0+1}^{T+i_0} \left\|\tilde{y}(k) - CZ^{k-i_0}(x(i_0)+F_c d_0)\right\|_2^2\end{aligned} \tag{22}$$

while steps $\mathbf{P2}$ and $\mathbf{P3}$ remain the same. The presence of w in $\tilde{y}$ will make the estimate inaccurate, and thereby, preserve the privacy of the parameters. The parameter identification error suffered by the intruder can be quantified as

$$\mathcal{E}_P = \mathbb{E}[\|\hat{P}_T - P\|_F] \tag{23}$$

where the expectation $\mathbb{E}$ is taken with respect to the noise. We will present numerical simulation results regarding the identification error in Section V, and show that it increases with increase in the privacy level (i.e., decrease in the parameter ϵ).

Remark 6 (Robustness to other identification techniques): In this paper, we consider one specific technique for parameter identification, which is NLS. There may be alternative and possibly more advanced techniques that the intruder may employ to obtain better estimates of P from the noisy measurements. However, a fundamental property of DP mechanism is its resilience to postprocessing, which means that one cannot weaken the DP guarantee by processing the differentially private outputs using *any* technique [12]. Finding the optimal identification technique is not the central premise and is beyond the scope of this paper. ■

While the privacy noise prevents accurate identification by the intruder, it also adversely affects the eigenvalue identification procedure by the control center, which we characterize next. Using (9), the noisy measurements can be written as [for comparison, see (11)]

$$\tilde{y}(k) = \tilde{\varphi}^T(k)\theta + v(k), \quad k > i_0 \tag{24}$$

where $v(k) = \sum_{i=1}^{n} a^{(i)} w(k-i)$ is a weighted noise, and the data vector $\tilde{\varphi}$ has a structure similar to φ in (12) except that it is constructed using noisy measurements $\tilde{y}$. These noisy measurements are used in the IV identification step (step **E1** of the eigenvalue identification method in Section II-B) which can be written as [see (13) for comparison]

$$\begin{aligned}\hat{\theta}_T &= \text{sol}\left\{\frac{1}{T}\sum_{k=i_0+1}^{T+i_0} \tilde{\zeta}(k)[\tilde{y}(k) - \tilde{\varphi}(k)^T\theta] = 0\right\}\\&= \left(\frac{1}{T}\sum_{k=i_0+1}^{T+i_0} \tilde{\zeta}(k)\tilde{\varphi}(k)^{T+i_0}\right)^{-1} \frac{1}{T}\sum_{k=i_0+1}^{T} \tilde{\zeta}(k)\tilde{y}(k)\\&= \theta + \left(\frac{1}{T}\sum_{k=i_0+1}^{T+i_0} \tilde{\zeta}(k)\tilde{\varphi}(k)^{T+i_0}\right)^{-1} \frac{1}{T}\sum_{k=i_0+1}^{T} \tilde{\zeta}(k)v(k)\end{aligned} \tag{25}$$

where $\tilde{\zeta}(k)$ are appropriately chosen instruments that are uncorrelated with the noise $v(k)$. The remaining steps **E2–E4** remain same as in the noiseless case. Observe that due to DP noise, eigenvalue identification will be inaccurate. Let a_c denote the coefficients of the characteristic polynomial of $A_c(P)$ and let $\hat{a}_{c,T}$ denote its estimate. We evaluate the eigenvalue identification performance in terms of the estimation error of these coefficients

$$\mathcal{E}_{a_c} = \mathbb{E}\left[\|\hat{a}_{c,T} - a_c\|_2\right] \tag{26}$$

where the expectation is taken with respect to the noise.

A. Noise Design for DP

Next, we present the properties of the privacy noise $w(k)$ that guarantee that the mechanism $\mathcal{M}$ satisfies the DP criterion of (21). As standard in the literature [12], [39], the noise provides DP if it satisfies the following two conditions: 1) w is Laplacian, zero-mean, and independent[3] over time and 2) the noise level is calibrated according to the sensitivity of the system.

[3] Although independence over time is not a necessary requirement, it is the standard approach in DP literature (for example, see [12], [25]). Since optimal noise design is not the primary focus of the paper, we use independent noises in the privacy mechanism.

Since we wish to protect the privacy the parameters P using the measurements y, we compute the sensitivity from P to y and calibrate the noise level proportional to the sensitivity. Next, we present the formal definition of sensitivity. For the remainder of this section, we assume that the impulse disturbance vanishes at $t_0 = 0$. This is merely for the ease of representation, and the analysis can be readily generalized for any $t_0 > 0$.

Definition 3 (Sensitivity): Let y and y' denote the measurements corresponding to parameters P and $P^{'}$, respectively. The system sensitivity at time instant $k \geq 0$ is defined as

$$\Delta(k) = \sup_{P, P^{'}:\text{adj}(\beta)} \|y(k) - y'(k)\|_1 . \tag{27}$$

The sensitivity characterizes the maximum possible difference (in terms of 1-norm) between the measurements resulting from any two possible adjacent parameters. It depends on a number of system parameters that we will characterize later in this section. Next, we show that we can use the sensitivity to design the privacy noise.

Theorem 1 (Noise design for DP): The mechanism $\mathcal{M}$ in (20) is ϵ-differentially private until time T if $w(k)$ is drawn independently at each k according to

$$w(k) \sim \text{Lap}(0, c_k)^p \quad \text{and} \quad \sum_{k=0}^{T} \frac{\Delta(k)}{c_k} \leq \epsilon \tag{28}$$

for any $c_k > 0,\ k = 0, \ldots, T$.

Proof: Let $\tilde{y}$ and $\tilde{y}'$ denote the noisy measurements corresponding to any two β-adjacent parameters P and $P^{'}$. From the mechanism $\mathcal{M}$, we have $\tilde{y}[0:T] = y[0:T] + w[0:T]$. Further, let $z = [z_0^T, z_1^T, \ldots, z_T^T]^T \in \mathbb{R}^{n(T+1)}$ denote the integration variable. Then

$$\begin{aligned}
\mathbb{P}[\tilde{y}[0:T] \in S] &\overset{(a)}{=} \int_S \prod_{k=0}^{T} \frac{1}{(2c_k)^n} e^{\frac{-\|z_k - y(k)\|_1}{c_k}} dz_k \\
&\overset{(b)}{\leq} \int_S \prod_{k=0}^{T} \frac{1}{(2c_k)^n} e^{\frac{-\|z_k - y'(k)\|_1}{c_k}} e^{\frac{\|y(k) - y'(k)\|_1}{c_k}} dz_k \\
&\overset{(c)}{\leq} e^{\left(\sum_{k=0}^{T} \frac{\Delta(k)}{c_k}\right)} \int_S \prod_{k=0}^{T} \frac{1}{(2c_k)^n} e^{\frac{-\|z_k - y'(k)\|_1}{c_k}} dz_k \\
&\overset{(d)}{\leq} e^{\epsilon} \mathbb{P}[\tilde{y}'[0:T] \in S]
\end{aligned}$$

where (a) follows from the joint Laplacian distribution of $w[0:T]$, (b) follows from the triangle inequality: $-\|x - p\|_1 \leq -\|x - p^{'}\|_1 + \|p - p^{'}\|_1$, (c) follows from the definition of sensitivity, and (d) follows from the condition given in the theorem. Thus, the DP condition (21) is satisfied. ■

From the preceding theorem, it is clear that in order to design the noise w, we need to characterize the sensitivity of the system. However, it is difficult to obtain an exact expression for the sensitivity. Therefore, we next obtain an upper bound on the sensitivity which can be used to design the noise level.

B. Upper Bound on Sensitivity

Let $\bar{x}(P)$ denote the steady-state value of the continuous-time system (1) (c.f. Assumption **A1**). If $A_c(P)$ is stable, then $\bar{x}(P)$ is trivially zero. For marginally stable systems [i.e., a single eigenvalue on the imaginary axis located at zero (see Assumption **A1**)], using Assumption **A4**, the steady-state value can be readily obtained as (assuming $d_c(t)$ as an impulse at $t = 0$)

$$\begin{aligned}
&\bar{x}(P) = \bar{A}_c(P)(x_c(0) + F_c d_0) \quad \text{where} \\
&\bar{A}_c(P) \triangleq \lim_{t \to \infty} e^{A_c(P)t} = \nu_0(A_c(P))\tilde{\nu}_0(A_c(P)) \quad \text{with} \\
&\qquad \tilde{\nu}_0(A_c(P))\nu_0(A_c(P)) = 1
\end{aligned} \tag{29}$$

where $\nu_\lambda(\cdot)$ and $\tilde{\nu}_\lambda(\cdot)$ denote the right and left eigenvectors associated with the eigenvalue λ, respectively. Next, Assumption **A5** implies that $\bar{A}_c(P)$ and $\bar{x}(P)$ do not depend on P. Therefore, we drop their dependence on P and denote them by $\bar{x}$ and $\bar{A}_c$, respectively.[4]

For marginally stable systems, we use the fact that the system describing the evolution of the error defined as $e(t) \triangleq x_c(t) - \bar{x}$ is stable. Towards this, we have the following lemma.

Lemma 1 (Error evolution): The dynamics of the error satisfies

$$\dot{e}(t) = \tilde{A}_c(P)e(t) \quad \text{for} \quad t \geq 0 \quad \text{where}$$

$$\tilde{A}_c(P) = \begin{cases} A_c(P) & \text{if } A_c(P) \text{ is stable} \\ A_c(P) - \eta\bar{A}_c & \text{if } A_c(P) \text{ is marginally stable} \end{cases} \tag{30}$$

and η is any positive constant.

Proof: The statement is trivial when the dynamics is stable. Let $x_0 = x_c(0) + F_c d_0$. For marginally stable cases, from the definition of $\bar{A}_c$ in (29), it can be easily seen that $A_c(P)\bar{A}_c = \bar{A}_c A_c(P) = 0$ and $\bar{A}_c^2 = \bar{A}_c$. Further, we have

$$\begin{aligned}
\bar{A}_c x_c(t) = \bar{A}_c e^{A_c(P)t} x_0 = \bar{A}_c \left[\sum_{k=0}^{\infty} \frac{A_c^k(P) t^k}{k!}\right] x_0 \\
\overset{(a)}{=} \bar{A}_c x_0 = \bar{x}
\end{aligned} \tag{31}$$

where (a) follows from $\bar{A}_c A_c(P) = 0$. Next, we have

$$\tilde{A}_c(P)e(t) = (A_c(P) - \bar{A}_c)(x_c(t) - \bar{x}) \overset{(b)}{=} A_c(P)x_c(t) = \dot{e}(t)$$

where (b) follows from $A_c(P)\bar{x} = 0, \bar{A}_c\bar{x} = \bar{x}$ and (31). ■

Next, we derive the eigenvalues of the $\tilde{A}_c(P)$ for the marginally stable case and show that it is stable.

Lemma 2 (Error eigenvalues): For marginally stable systems, the set of eigenvalues of $\tilde{A}_c(P)$ is

$$\lambda(\tilde{A}_c(P)) = \{-\eta, \{\lambda_i(A_c(P)) : \lambda_i(A_c(P)) \neq 0\}_{i=1}^n\}. \tag{32}$$

[4] $\bar{x}$ is also the steady-state value of the discrete-times system in (4). Thus, it can also be written as $\bar{x} = \bar{A}_d(P)(x(0) + F_c d_0)$, where $\bar{A}_d(P) \triangleq \lim_{k \to \infty} A_d^k(P) = \nu_1(A_d(P))\tilde{\nu}_1(A_d(P))$, with $\tilde{\nu}_1(A_d(P))\nu_1(A_d(P)) = 1$.

Proof: Let λ_i and $\nu_{\lambda_i}, \tilde{\nu}_{\lambda_i}$ denote the eigenvalues and corresponding eigenvectors of $A_c(P)$. Since $\bar{A}_c = \nu_0 \tilde{\nu}_0$ with $\tilde{\nu}_0 \nu_0 = 1$, we have

$$\tilde{A}_c(P)\nu_0 = A_c(P)\nu_0 - \eta \bar{A}_c \nu_0 = -\eta \nu_0.$$

Thus, $-\eta$ is an eigenvalue of $\tilde{A}_c(P)$. Next, using $\bar{A}_c A_c(P) = 0$, for $\lambda_i \neq 0$ we have

$$0 = \bar{A}_c A_c(P) \nu_{\lambda_i} = \bar{A}_c \lambda_i \nu_{\lambda_i} \Rightarrow \bar{A}_c \nu_{\lambda_i} = 0.$$

Thus, we get $\tilde{A}_c(P)\nu_{\lambda_i} = A_c(P)\nu_{\lambda_i} - \bar{A}_c \nu_{\lambda_i} = \lambda_i \nu_{\lambda_i}$. Thus, all the eigenvalues $\lambda_i \neq 0$ are also the eigenvalues of $\tilde{A}_c(P)$. ■

The above result shows that the eigenvalues of $\tilde{A}_c(P)$ are the same as that of $A_c(P)$, except the eigenvalue zero which is shifted to $-\eta$. Thus, $\tilde{A}_c(P)$ is stable.

Corollary 1: $A_c(P)$ and $\tilde{A}_c(P)$ are simultaneously diagonalizable.

Proof: From the proof of Lemma 2, we can observe that both $A_c(P)$ and $\tilde{A}_c(P)$ have the same eigenvectors and thus, they are simultaneously diagonalizable. ■

With a slight abuse of notation, we define the global spectral abscissa as

$$\begin{aligned}\mu_{\max} &\triangleq \sup_{P \in \mathcal{P}, i \in \{1,2,\ldots,n\}} \{Re(\lambda_i(A_c(P))) : \lambda_i(A_c(P) \neq 0\} \\ &= \sup_{P \in \mathcal{P}} \mu(\tilde{A}_c(P)) < 0 \end{aligned} \quad (33)$$

where $\mu(\cdot)$ denotes the spectral abscissa and the second equality is valid for $\tilde{A}_c(P) = A_c(P) - \eta \bar{A}_c$ with sufficiently large η (using Lemma 2). By Corollary 1, let $A_c(P) = X^{-1}(P)\Lambda_c(P)X(P)$ and $\tilde{A}_c(P) = X^{-1}(P)\tilde{\Lambda}_c(P)X(P)$, where $\Lambda_c(P)$ and $\tilde{\Lambda}_c(P)$ are diagonal matrices consisting of the eigenvalues of $A_c(P)$ and $\tilde{A}_c(P)$, respectively. Then, let

$$\kappa_{\max} \triangleq \sup_{P \in \mathcal{P}} \kappa(X(P)) = \sup_{P \in \mathcal{P}} \|X(P)\|_2 \|X^{-1}(P)\|_2. \quad (34)$$

Lemma 3: For $\tilde{A}_c(P) = A_c(P) - \eta \bar{A}_c$ with $\bar{A}_c$ defined in (29), we have

$$e^{\tilde{A}_c(P)t} = e^{A_c(P)t} - (1 - e^{-\eta t})\bar{A}_c. \quad (35)$$

Proof: Using $A_c(P)\bar{A}_c = \bar{A}_c A_c(P) = 0$ and $\bar{A}_c^2 = \bar{A}_c$, it follows that

$$\tilde{A}_c^k(P) = A_c^k(P) + (-1)^k \eta^k \bar{A}_c, \quad k \geq 1. \quad (36)$$

Next, we have

$$\begin{aligned} e^{\tilde{A}_c(P)t} &= \sum_{k=0}^{\infty} \frac{\tilde{A}_c^k(P)t^k}{k!} \\ &\overset{(36)}{=} I + \sum_{k=1}^{\infty} \frac{A_c^k(P)t^k}{k!} + \bar{A}_c \left[\sum_{k=1}^{\infty} \frac{(-1)^k \eta^k t^k}{k!} \right] \\ &= e^{A_c(P)t} - (1 - e^{-\eta t})\bar{A}_c. \end{aligned}$$

■

Using (35), the response of (4) for the impulse $d(k) = d_0 \delta(k)$ can be written as

$$\begin{aligned} x(k) &= A_d^k(P)(x_0 + F_c d_0) = e^{A_c(P)T_s k}(x_0 + F_c d_0) \\ &= [e^{\tilde{A}_c(P)T_s k} + (1 - e^{-\eta T_s k})\bar{A}_c](x_0 + F_c d_0). \end{aligned} \quad (37)$$

Next, we provide a known result which will be used to derive the upper bound to the sensitivity.

Lemma 4 (Perturbation of matrix exponential [40]): Let A_1 and A_2 be two $n \times n$ diagonalizable matrices for which $\mu(A_1) < 0$ and $\mu(A_2) < 0$. Let X_1 and X_2 be eigenvector matrices of A_1 and A_2, respectively. Then

$$\begin{aligned} \|e^{A_1 t} - e^{A_2 t}\|_1 &\leq \kappa(X_1)\kappa(X_2) t \sqrt{n} \|A_1 - A_2\|_F \\ &\times \max\{e^{\mu(A_1)t}, e^{\mu(A_2)t}\}. \end{aligned}$$

Proof: The proof follows from [40, Corollary 2.4] using $f(s) = e^{st}$. ■

Next, we derive the upper bound on the sensitivity of the system.

Theorem 2 (Sensitivity bound): The sensitivity $\Delta(k)$ in (27) can be upper bounded as

$$\Delta(k) \leq \bar{\Delta}(k) \triangleq \alpha k \rho_{\max}^k \quad (38)$$

where $\alpha = \kappa_{\max}^2 \sqrt{n} T_s \delta \|C\|_1 \|(x(0) + F_c d_0)\|_1$ and $\rho_{\max} = e^{\mu_{\max} T_s}$ $\delta = \sup_{P,P':\text{adj}(\beta)} \|A_c(P) - A_c(P^{'})\|_F$.

Proof: Let y and y' denote the measurements corresponding to any two β-adjacent parameters P and $P^{'}$. For the impulse input $d(k)$, we have

$$\begin{aligned} &\|y(k) - y'(k)\|_1 \\ &\overset{(a)}{=} \|Ce^{\tilde{A}_c(P)T_s k}(x(0) + F_c d_0) - Ce^{\tilde{A}_c(P')T_s k}(x(0) + F_c d_0)\|_1 \\ &\overset{(b)}{\leq} \|C\|_1 \|(x(0) + F_c d_0)\|_1 \|e^{\tilde{A}_c(P)T_s k} - e^{\tilde{A}_c(P')T_s k}\|_1 \\ &\overset{(c)}{\leq} \|C\|_1 \|(x(0) + F_c d_0)\|_1 \kappa_{max}^2 \sqrt{n} T_s k e^{\mu_{max} T_s k} \\ &\quad \times \|A_c(P) - A_c(P^{'})\|_F \end{aligned}$$

where (a) follows from (37), (b) follows from the submultiplicative property of matrix norm, (c) follows from Lemma 4, definitions of $\kappa_{\max}$ and $\mu_{\max}$ in (33) and (34) and the fact that $\tilde{A}_c(P) - \tilde{A}_c(P') = A_c(P) - A_c(P')$. The theorem then follows using the definition of δ. ■

It is easy to observe that the sensitivity bound $\bar{\Delta}(k)$ can be used in (28) to determine the noise level and guarantee DP. Observe that the sensitivity bound $\bar{\Delta}(k)$ in (38) decays exponentially for large k. This is a direct consequence of Assumption **A5**. Without this assumption, the sensitivity (and the sensitivity bound) would remain constant asymptotically. As a result, the noise level required to ensure DP would increase with T. Clearly, such unbounded noise in the system is undesirable. Further, the sensitivity bound becomes loose if the global condition number $\kappa_{\max}$ of the matrices $\{A_c(P)\}_{P \in \mathcal{P}}$ is large or the global spectral abscissa $|\mu_{\max}|$ is small. Utilizing the decaying behavior of the sensitivity bound, we next show that DP can be guaranteed using an exponentially decaying noise.

Lemma 5 (DP through decaying noise): The mechanism $\mathcal{M}$ in (20) is ϵ-differentially private until time T if $w(k)$ is drawn independently at each k according to $w(k) \sim Lap(0, ck\gamma^k)^n$ with $0 < \rho_{\max} < \gamma < 1$, and

$$c \geq \underline{c} \triangleq \frac{\alpha}{\epsilon} \frac{1 - (\frac{\rho_{\max}}{\gamma})^{T+1}}{1 - \frac{\rho_{\max}}{\gamma}}. \tag{39}$$

Proof: It can be easily verified that (39) implies that the condition in (28) is satisfied. ■

The lower bound $\underline{c}$ in (39) is bounded for all T. Thus, the privacy noise level also remains bounded.

Remark 7 (Effect of time horizon T): As T increases, the number of measurements available to the intruder increases, and therefore, it is able to identify the sensitive parameters with increasing accuracy. To prevent this, the noise level in the privacy mechanism also increases with T, as evident in (39). Since the system is stable (marginally stable), the transient introduced by the disturbance will vanish (settle) after a period of time and privacy noise is not required after that period. This is also evident from (39), where $\underline{c}$ saturates as T increases, and the fact that the noise level decays with time. ■

In order to compute the upper bound given in (38), we need to further characterize the sensitivity δ from the sensitive parameters to the continuous-time closed-loop state matrix, as contained in the constant α (see Theorem 2). This sensitivity depends on the structure of $A_c(P)$, and is specific to the application for which privacy is being designed. We characterize this quantity for the two examples considered earlier in Section III.

1) Sensitivity for Power System Models: Consider the second-order consensus example for power system models in (14). Since L is symmetric and $L\mathbf{1}_N = \mathbf{0}_N$, we have $\nu_1(A_d(L)) = \nu_0(A_c(L)) = \gamma_1 [\begin{smallmatrix} \mathbf{1}_N \\ \mathbf{0}_N \end{smallmatrix}]$ and $\tilde{\nu}_1(A_d(L)) = \tilde{\nu}_0(A_c(L)) = \gamma_2 [\mathbf{1}_N^T D \; \mathbf{1}_N^T M]$ where $\gamma_1, \gamma_2 \in \mathbb{R}$ are some scalars. Using (29), the steady-state value for consensus can be obtained as

$$\bar{x} = \frac{1}{\sum_{i=1}^{N} d_i} \begin{bmatrix} \mathbf{1}_N \\ \mathbf{0}_N \end{bmatrix} [\mathbf{1}_N^T D \quad \mathbf{1}_N^T M] (x(0) + F_c d_0). \tag{40}$$

Note that $\bar{x}$ does not depend on L, indicating that Assumption **A5** is satisfied.

Lemma 6 (Sensitivity bound for consensus): The sensitivity δ for (14) is upper bounded by

$$\delta_{\text{CONS}} \leq \bar{\delta}_{\text{CONS}} \triangleq \|M^{-1}\|_F \beta.$$

Proof: Using the structure of $A_c(L)$ in (14), we have

$$\|A_c(L) - A_c(L^{'})\|_F = \|M^{-1}(L - L^{'})\|_F \overset{(a)}{\leq} \|M^{-1}\|_F \beta$$

where (a) follows from the property $\|AB\|_F \leq \|A\|_F \|B\|_2$ and the fact that L and $L^{'}$ are adjacent. ■

2) Sensitivity for LQR Control: We define separation of an $n \times n$ matrix A as

$$\text{sep}(A) = \min\{\|SA + A^T S\|_2 : S = S^T \in \mathbb{R}^{n \times n}, \|S\|_2 = 1\}.$$

Next, we present a known result regarding the perturbation of the Riccati equation.

Lemma 7 (Perturbation of Riccati equation [41]): Consider two Riccati equations of the form (18) that differ only with respect to the parameters, given as Q and Q'

$$\text{If} \quad \|Q - Q'\|_2 < \frac{\text{sep}^2(A_c(Q))}{4\|BR^{-1}B^T\|_2} \quad \text{then} \tag{41}$$

$$\|V(Q) - V(Q')\|_2 \leq \frac{2\|Q - Q'\|_2}{\text{sep}(A_c(Q))}. \tag{42}$$

Proof: The proof follows from [41, Theorem 2.1]. ■

Next, we present the sensitivity bound for the LQR problem.

Lemma 8 (Sensitivity bound for LQR control): If

$$\beta < \frac{\mu_{\max}^2}{\kappa_{\max}^4 \|BR^{-1}B^T\|_2} \tag{43}$$

then sensitivity δ for the LQR control is upper bounded by

$$\delta_{\text{LQR}} \leq \frac{\|BR^{-1}B^T\|_2 \kappa_{\max}^2 \beta}{|\mu_{\max}|}. \tag{44}$$

Proof: From [41, Theorem 3.2], and the definitions of $\kappa_{\max}$ [in (34)] and $\mu_{\max}$ [in (33)], we have

$$\text{sep}(A_c(Q)) \geq \frac{2|\mu_{\max}|}{\kappa_{\max}^2} \quad \forall \quad Q \in \mathcal{P}. \tag{45}$$

Further, since Q and Q' are β-adjacent, we have

$$\|Q - Q'\|_2 \leq \beta < \frac{\mu_{\max}^2}{\kappa_{\max}^4 \|BR^{-1}B^T\|_2} \leq \frac{\text{sep}^2(A_c(Q))}{4\|BR^{-1}B^T\|_2}$$

and thus, the condition (41) in Lemma 7 is satisfied. Next

$$\|A_c(Q) - A_c(Q^{'})\|_F = \|BR^{-1}B^T(V(Q) - V(Q'))\|_F$$
$$\overset{(a)}{\leq} \|BR^{-1}B^T\|_F \frac{2\|Q - Q'\|_2}{\text{sep}(A_c(Q))} \overset{(b)}{\leq} \frac{\|BR^{-1}B^T\|_F \kappa_{\max}^2 \beta}{|\mu_{\max}|}$$

where (a) follows from the submultiplicative property of the norm and Lemma 7 and (b) follows from the the β-adjacency of Q and Q' and (45). ■

An intuitive interpretation of Lemma 8 is that the sensitivity bound in (44) is large if the systems are closer to instability ($|\mu_{\max}|$ is small), or they have a large condition number ($\kappa_{\max}$ is large). Both these quantities dictate the magnitude of the transient response of the system. Thus, systems with large transient responses have larger sensitivity.

We conclude this section by emphasizing that the privacy guarantees provided by the noise are valid when the model (1) and (2) accurately represents the underlying physical system. If the system modeling errors are large, then computation of sensitivity can be erroneous and privacy guarantees might be violated.

V. Numerical Example

In this section, we present a numerical example for our DP-based design. We consider a power system network with $N = 3$

TABLE I
COMPARISON OF ESTIMATION ERRORS FOR DIFFERENT NOISE (PRIVACY) LEVELS

Noise level (c)	Average $\Delta\mathcal{E}_{a_c}$	Average $\Delta\mathcal{E}_{\bar{L}}$
0	3.113×10^{-11}	6.263×10^{-10}
0.2	1.615×10^{-3}	1.416×10^{-1}
0.4	1.924×10^{-3}	1.756×10^{-1}
0.6	2.350×10^{-3}	1.887×10^{-1}
0.8	3.615×10^{-3}	2.115×10^{-1}
1	5.167×10^{-3}	2.499×10^{-1}

generators for which the Laplacian matrix is given as

$$L = \begin{bmatrix} 10 & -6 & -4 \\ -6 & 9.5 & -3.5 \\ -4 & -3.5 & 7.5 \end{bmatrix}.$$

The other system parameters are considered as $M = 10I_3$, $D = 10I_3$, $\bar{F} = \text{diag}(1, 0, 0)$, $x(0) = \mathbf{0}_6$, and the disturbance is an impulse with $d_0 = 1$. The coefficients of the characteristic polynomial of A_c are $a_c = [1, 3, 5.7, 6.4, 4.47, 1.77, 0]^T$ with $||a_c||_2 = 10.32$. For the above system, $\bar{T}_s = 2.72$ [see (5)]. The state measurements are sampled at $T_s = 1$ s to satisfy Assumption **A7**. The privacy mechanism $\mathcal{M}$ in (20) is used to add Laplacian noise to the measurements before sending them to the control center.

To design the privacy noise, we use the sensitivity upper bounds obtained from Theorem 2 and Lemma 6. Here, $\kappa(X(L)) = 2.41$ [see (34)] and $\mu(\tilde{A}_c(L)) = -0.5$ [see (30)]. In order to compute $\kappa_{\max}$ in (34) and $\mu_{\max}$ in (33), we need a characterization of the set $\mathcal{P}$ for all possible sensitive parameters. Without loss of generality, we avoid this explicit characterization and choose $\kappa_{\max} = 3$ and $\mu_{\max} = -0.4$. In scenarios where the set $\mathcal{P}$ is given explicitly, $\kappa_{\max}$ and $\rho_{\max}$ can be easily computed. We choose the adjacency parameter $\beta = 0.1$ (see Definition 1). With the given choice of parameters, we have $\alpha = 0.038$ and $\rho_{\max} = 0.67$ [see (38) and Lemma 6]. We design the decaying Laplacian noise using Lemma 5. We choose the noise decay factor $\gamma = 0.8$ and simulate the system for $T = 1000$ time steps. Using these parameter values, the noise lower bound in (39) becomes $\underline{c} = 0.236\epsilon^{-1}$. We, therefore, choose $c = \underline{c}$ with the range of the noise level being $c = [0, 1]$.

We perform Monte Carlo simulations by choosing 1000 random topologies L of size $N = 3$. The nondiagonal weights of each symmetric L are sampled uniformly between the range $[-7, -1]$ and we use these nondiagonal weights to compute the diagonal weights. The results are reported in Table I.

The second column of Table I shows the effect of the noise on eigenvalue identification. The identification is performed in MATLAB using the `iv4` function in its system identification toolbox, which implements a four stage IV identification method [11]. The estimation quality is characterized in terms of expected relative coefficient error $\Delta\mathcal{E}_{a_c} \triangleq \frac{\mathcal{E}_{a_c}}{||a_c||_2}$ [see (26)]. For each L, we approximate the expected error by averaging the error values over 1000 iterations for each noise level and report the average error over all the topologies. The table shows that the identification error increases with increase in the noise (or privacy) level. Similarly, the third column of Table I shows the effect of the noise on the identification of L. We assume that the intruder does not have access to M and D, and thus, it can only identify $\bar{L} \triangleq M^{-1}L$ instead of L using the state measurements (see Remark 2). Therefore, it extracts $\hat{\bar{L}}_T$ from the NLS estimate $\hat{A}_{c,T}(\hat{\bar{L}}_T)$ [see structure of $A_c(L)$ in (14)] in step **P3** of the parameter identification procedure. The relative parameter identification error is defined as $\Delta\mathcal{E}_{\bar{L}} \triangleq \frac{\mathbb{E}\left[\left\|\hat{\bar{L}}_T - M^{-1}L\right\|_F\right]}{\|M^{-1}L\|_F}$ and we report the average error over all the topologies. As seen in Table I the identification error is close to zero for $c = 0$. Thus, the weighted topology is accurately identified if there is no privacy mechanism. As the noise (or privacy) level increases, the identification error also increases making it more difficult for the intruder to identify the topology. Comparing columns two and three of Table I, we can observe that the average relative error of eigenvalue identification is two times the order of magnitude less than that of topology identification. In other words, the privacy noise promotes significantly larger performance degradation for the intruder as compared to the control center, implying that our mechanism can provide privacy without significant performance loss for system monitoring. As mentioned in Section I, this performance difference is not due to the use of different algorithms (IV versus NLS), but due to the fundamentally different nature of the two problems, i.e., pseudolinear estimation for eigenvalue identification and nonlinear estimation for topology identification.

VI. CONCLUSION

In this paper, we develop a noise-adding DP mechanism to protect the privacy of the sensitive parameters associated with the dynamics of multiagent LTI systems. We derive an upper bound for the sensitivity of the system, and use it to design the privacy noise. We present two applications of our privacy framework, namely, an electromechanical model of a power system and an LQR control model of a generic linear system. We consider a specific scenario where the goal of the control center is to estimate the eigenvalues of the system. Our numerical simulations show that for this scenario, the performance degradation suffered by the intruder due to privacy noise is significantly higher compared to that of the system-level objective at the control center. Our future research directions will be to obtain a tighter bound on the sensitivity, and to extend this method for the privacy of nonlinear system identification, and also to a game-theoretic framework where the intruder and the system operator may challenge each other in terms of finding more accurate ways for estimating their respective design objectives.

REFERENCES

[1] M. Mesbahi and M. Egerstedt, *Graph Theoretic Methods in Multiagent Networks*. Princeton, NJ, USA: Princeton Univ. Press, 2010.

[2] S. Nabavi, J. Zhang, and A. Chakrabortty, "Distributed optimization algorithms for wide-area oscillation monitoring in power systems using interregional PMU-PDC architectures," *IEEE Trans. Smart Grid*, vol. 6, no. 5, pp. 2529–2538, Sep. 2015.

[3] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proc. IEEE*, vol. 95, no. 1, pp. 215–233, Jan. 2007.

[4] M. C. Priess, R. Conway, J. Choi, J. M. Popovich, and C. Radcliffe, "Solutions to the inverse LQR problem with application to biological systems analysis," *IEEE Trans. Control Syst. Technol.*, vol. 23, no. 2, pp. 770–777, Mar. 2014.

[5] M. Monfort, A. Liu, and B. Ziebart, "Intent prediction and trajectory forecasting via predictive inverse linear-quadratic regulation," in *Proc. AAAI Conf. Artif. Intell.*, 2015, pp. 3672–3678.

[6] D. Kendrik, "Control theory with applications to economics," in *Handbook of Mathematical Economics*, vol. 1, K. Arrow, M. D. Intriligator, Eds., Amsterdam, The Netherlands: North Holland, 2000, ch. 4, pp. 111–158.

[7] S Han, U. Topcu, and G. J. Pappas, "Differentially private distributed constrained optimization," *IEEE Trans. Autom. Control*, vol. 62, no. 1, pp. 50–64, Jan. 2017.

[8] M. T. Hale and M. Egerstedt, "Cloud-enabled differentially private multiagent optimization with constraints," *IEEE Trans. Control Netw. Syst.*, vol. 4, no. 4, pp. 1693–1706, Dec. 2018.

[9] C. Dwork, "Differential privacy," in *Proc. Int. Colloq. Automata, Lang., Program.*, 2006, vol. 4052, pp. 1–12.

[10] C. Dwork, "A firm foundation for private data analysis," *Commun. ACM*, vol. 54, no. 1, pp. 86–95, 2011.

[11] L. Ljung, *System Identification: Theory for the User*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1999.

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

[13] Z. Huang, S. Mitra, and G. Dullerud, "Differentially private iterative synchronous consensus," in *Proc. ACM Workshop Privacy Electron. Soc.*, 2012, pp. 81–90.

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

[15] Z. Huang, Y. Wang, S. Mitra, and G. Dullerud, "On the cost of differential privacy in distributed control systems," in *Proc. ACM Int. Conf. High Confidence Netw. Syst.*, 2014, pp. 105–114.

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

[17] Z. Huang, S. Mitra, and N. Vaidya, "Differentially private distributed optimization," in *Proc. Int. Conf. Distrib. Comput. Netw.*, 2015, Art. no. 4.

[18] E. Nozari, P. Tallapragada, and J. Cortes, "Differentially private distributed convex optimization via functional perturbation," *IEEE Trans. Control Netw. Syst.*, vol. 5, no. 1, pp. 395–408, Mar. 2018.

[19] J. Hsu, A. Roth, T. Roughgarden, and J. Ullman, "Privately solving linear programs," in *Proc. Int. Conf. Automata, Lang., Program.*, 2014, pp. 612–624.

[20] S. Han, U. Topcu, and G. Pappas, "Differentially private convex optimization with piecewise affine objectives," in *Proc. IEEE Conf. Decis. Control*, 2014, pp. 2160–2166.

[21] S. Song, K. Chaudhuri, and A. D. Sarwate, "Stochastic gradient descent with differentially private updates," in *Proc. IEEE Global Conf. Signal Inf. Process.*, 2013, pp. 245–248.

[22] M. Hale, A. Jones, and K. Leahy, "Privacy in feedback: The differentially private LQG," in *Proc. Amer. Control Conf.*, 2018, pp. 3386–3391.

[23] H. Zhang, Y. Shu, P. Cheng, and J. Chen, "Privacy and performance tradeoff in cyber-physical systems," *IEEE Netw.*, vol. 30, no. 2, pp. 62–66, Mar./Apr. 2016.

[24] T. Tanaka, M. Skoglund, H. Sandberg, and K. H. Johansson, "Directed information and privacy loss in cloud-based control," in *Proc. Amer. Control Conf.*, 2017, pp. 1666–1672.

[25] G. Bottegal, F. Farokhi, and I. Shames, "Preserving privacy of finite impulse response systems," *IEEE Control Syst. Lett.*, vol. 1, no. 1, pp. 128–133, Jul. 2017.

[26] J. Le Ny and G. J. Pappas, "Privacy-preserving release of aggregate dynamic models," in *Proc. Int. Conf. High Confidence Netw. Syst.*, 2013, pp. 49–56.

[27] V. Karwa, S. Raskhodnikova, A. Smith, and G. Yaroslavtsev, "Private analysis of graph structure," *ACM Trans. Database Syst.*, vol. 39, no. 3, 2014, Art. no. 22.

[28] W. Day, N. Li, and M. Lyu, "Publishing graph degree distribution with node differential privacy," in *Proc. Int. Conf. Manage. Data*, 2016, pp. 123–138.

[29] S. P. Kasiviswanathan, K. Nissim, S. Raskhodnikova, and A. Smith, "Analyzing graphs with node differential privacy," in *Theory Cryptography (Lecture Notes in Computer Science, vol. 7785)*. Berlin, Germany: Springer, 2013.

[30] Y. Wang, X. Wu, and L. Wu, "Differential privacy preserving spectral graph analysis," in *Advances in Knowledge Discovery and Data Mining (Lecture Notes in Computer Science, vol. 7819)*. Berlin, Germany: Springer, 2013.

[31] M. Showkatbakhsh, P. Tabuada, and S. Diggavi, "System identification in the presence of adversarial outputs," in *Proc IEEE Conf. Decis. Control*, 2016, pp. 7177–7182.

[32] V. Katewa, A. Chakrabortty, and V. Gupta, "Protecting privacy of topology in consensus networks," in *Proc. Amer. Control Conf.*, 2015, pp. 2476–2481.

[33] R. Bellman and K. J. Astrom, "On structural identifiability," *Math. Biosci.*, vol. 7, pp. 329–339, 1970.

[34] P. Kundur, *Power System Stability and Control*. New York, NY, USA: McGraw-Hill, 1993.

[35] W. Ren and R. W. Beard, *Distributed Consensus in Multi-Vehicle Cooperative Control*. London, U.K.: Springer-Verlag, 2008.

[36] F. Tisseur and K. Meerbergen, "The quadratic eigenvalue problem," *SIAM Rev.*, vol. 43, no. 2, pp. 235–286, 2001.

[37] H. Kong, G. Goodwin, and M. Seron, "A revisit to inverse optimality of linear systems," *Int. J. Control*, vol. 85, no. 10, pp. 1506–1514, 2012.

[38] B. Molinari, "The stable regulator problem and its inverse," *IEEE Trans. Autom. Control*, vol. 18, no. 5, pp. 454–459, May 1973.

[39] C. Dwork *et al.*, "Calibrating noise to sensitivity in private data analysis," in *Proc. Theory Cryptography Conf.*, 2006, pp. 265–284.

[40] M. I. Gil, "Perturbations of functions of diagonalizable matrices," *J. Linear Algebr. Electron.*, vol. 20, pp. 303–313, 2010.

[41] S. F. Xu, "Sensitivity analysis of the algebraic Riccati equations," *Numerische Math.*, no. 75, pp. 121–134, 1996.

Vaibhav Katewa received the bachelor's degree from the Indian Institute of Technology, Kanpur, India, in 2007, and the M.S. and Ph.D. degrees from the University of Notre Dame, Notre Dame, IN, USA, in 2012 and 2016, respectively, all in electrical engineering.

He is a Postdoctoral Scholar with the Department of Mechanical Engineering, University of California, Riverside, CA, USA. His research interests include analysis and design of security and privacy methods for cyber-physical systems and complex networks, decentralized and sparse feedback control, and protocol design for networked control systems.

Aranya Chakrabortty (SM'15) received the B.E. degree from Jadavpur University, Kolkata, India, in 2004, and the M.S. and Ph.D. degrees from Rensselaer Polytechnic Institute, Troy, NY, USA, in 2005 and 2008, respectively, all in electrical engineering.

Since 2010, he has been a Faculty Member with the Electrical and Computer Engineering Department, North Carolina State University, Raleigh, NC, USA, where he is currently an Associate Professor. His research interests are in all branches of control theory with applications to electric power systems.

Dr. Chakrabortty is an Associate Editor for the IEEE Transactions on Control System Technology and IEEE Transactions on Power Systems. He was the recipient of the NSF CAREER Award in 2011.

Vijay Gupta (SM'18) received the B.Tech. degree from the Indian Institute of Technology, Delhi, India, in 2001, and the M.S. and Ph.D. degrees from the California Institute of Technology, Pasadena, CA, USA, in 2002 and 2007, respectively, all in electrical engineering.

He is currently a Professor with the Department of Electrical Engineering, University of Notre Dame, having joined the faculty in January 2008. Prior to joining Notre Dame, he also served as a Research Associate with the Institute for Systems Research, University of Maryland, College Park, MD, USA. His research and teaching interests are broadly at the interface of communication, control, distributed computation, and human decision making.

Prof. Gupta was the recipient of the 2018 Antonio Ruberti Award from IEEE Control Systems Society, the 2013 Donald P. Eckman Award from the American Automatic Control Council, and a 2009 National Science Foundation (NSF) CAREER Award.