Preprints of the 20th World Congress
The International Federation of Automatic Control
Toulouse, France, July 9-14, 2017

Differential Privacy and Minimum-Variance Unbiased Estimation in Multi-agent Control Systems

Yu Wang, Sayan Mitra and Geir E. Dullerud

Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, USA
{*yuwang8, mitras, dullerud*}*@illinois.edu*

Abstract: In a discrete-time linear multi-agent control system, where the agents are coupled via an environmental state, knowledge of the environmental state is desirable to control the agents locally. However, since the environmental state depends on the behavior of the agents, sharing it directly among these agents jeopardizes the privacy of the agents' profiles, defined as the combination of the agents' initial states and the sequence of local control inputs over time. A commonly used solution is to randomize the environmental state before sharing – this leads to a natural trade-off between the privacy of the agents' profiles and the variance of estimating the environmental state. By treating the multi-agent system as a probabilistic model of the environmental state parametrized by the agents' profiles, we show that when the agents' profiles is ε-differentially private, there is a lower bound on the ℓ_1 induced norm of the covariance matrix of the minimum-variance unbiased estimator of the environmental state. This lower bound is achieved by a randomized mechanism that uses Laplace noise.

Keywords: ε-differential privacy, minimum-variance unbiased estimation, multi-agent control systems, Laplace-noise-adding mechanisms

1. INTRODUCTION

In a multi-agent system, where the dynamics of the agents are coupled via an environmental state, a knowledge on the environmental state can greatly improve the performance of individual agents. For instance, using mobile navigation applications informing the traffic condition, such as Google Maps and Waze, vehicles can avoid road congestion more accurately (Andrés et al., 2013; Xue et al., 2014). However, since the environmental state depends on the agents' states, these benefits come with the risk on privacy. For example, researchers have shown that Waze can be used to follow a user's movements (Wang et al., 2016); and even with anonymized data such as Google Maps (Andrés et al., 2013), the inherent structure of location data can lead to de-anonymization (Ma et al., 2013; Shokri et al., 2011). Similar risks and benefits arise in two-way coordination between consumers' demands and electric power utility companies: on one hand, sharing information can prevent over-provisioning through "peak-shaving" and reduce energy costs (Farhangi, 2010; Koufogiannis et al., 2014; Masoum et al., 2011), and on the other hand, expose the consumers' personal habits.

To reconcile this conflict between privacy and sharing information, a commonly used approach is to randomize the environmental state before sharing (Dwork, 2006; McSherry and Talwar, 2007; Li et al., 2010; Dwork et al., 2006). This leads to a natural trade-off – the more the users' data are randomized, the more privacy, and the less accurate the estimation on the environmental state can be, and vice versa. In this article, we study this trade-off between privacy and the variance of estimation in an idealized discrete-time, linear, multi-agent control system derived from (Wang et al., 2014, 2017; Huang et al., 2014).

The data to keep private are the agents' profiles, defined as the combination of the agents' initial states and the sequence of local control inputs over time. For example, in a routing problem, they are the initial locations and the sequence of maneuvering commands of the agents. To formally measure privacy, we adopt the concept of ε-differential privacy, which has been proposed from the literature on databases and theoretical computer science (Dwork, 2006), and has proven to be popular both theoretically and practical (Apple Inc, 2016). Informally, a differentially private statistical query on a database ensures that the probability distribution of the output does not change substantially with changes in the private data. Thus, an adversary cannot learn much about the participants by querying the database. More recently, this concept has been extended to dynamical systems and applied to various problems, such as distributed consensus protocols (Huang et al., 2012; Mo and Murray, 2014; Nozari et al., 2015), distributed optimization (Hale and Egerstedty, 2015; Huang et al., 2015) and filtering (Le Ny and Pappas, 2014).

Since its original development, several variations on the formal definition of differential privacy have been proposed in (Le Ny and Pappas, 2014; Chatzikokolakis et al., 2013). In this work, we are interested in real-valued continuously changing data (modeling physical quantities like position, energy consumption, etc.), and therefore, the

definition of differential privacy used here is the one similar to (Chatzikokolakis et al., 2013), which uses a ℓ_p-type metric on the user data.

For the estimation part, we consider the minimum-variance unbiased estimator (MVUE). This distinguishes this work from (Wang et al., 2014, 2017), in which the measure of variance is entropy. The MVUE are frequently involved in designing stochastic controllers, when the cost is quadratic. Since the definition of ε-differential privacy requires that the probability distribution function of the randomized environmental states does not change much when the agents' profile changes, we can expect that there is a lower bound, for example, the famous Cramer-Rao bound on the variance of the MVUE. Actually, we will show that the Cramer-Rao bound is not tight in this situation, and can derive a tight minimax lower bound on the ℓ_1 induced norm of the covariance matrix of the MVUE. In addition, we will give explicitly a randomized mechanism that achieves the tight lower bound.

In (Geng et al., 2015), the authors study the optimal noise-adding mechanisms that minimize certain ℓ_1 cost function while keeping the query ε-differential private. The solution is a staircase mechanism. This work is different in three aspects. First, while they use the common definition of ε-differential privacy, we adopt a stronger definition here. Second, the performance measure in this work is the ℓ_1 induced norm of the covariance matrix, as opposed to the ℓ_1 cost function. Finally, our problem is set upon multi-agent control systems where communication happens at every time step instead of a single query.

The rest of the paper is organized as follows: Section 2 gives basic definitions and notations used throughout the paper. Section 3 studies the trade-off between ε-differential privacy and mean-variance unbiased estimation in a parametric probabilistic model. Section 4 extends the results derived in Section 3 and establishes the trade-off between ε-differential privacy and mean-variance unbiased estimation in a multi-agent control system. Finally, the conclusions are presented in Section 5.

2. PRELIMINARIES

In this paper, we denote the sets of natural numbers and real numbers by $\mathbb{N}$ and $\mathbb{R}$, respectively. The set of n-dimensional real vectors and the set of $m \times n$ real matrices are denoted respectively by $\mathbb{R}^n$ and $\mathbb{R}^{m\times n}$. For a vector $x \in \mathbb{R}^n$, we denote the i^{th} component by x_i; for a matrix $A \in \mathbb{R}^{m\times n}$, we denote the (i, j) component by A_{ij}. For $n \in \mathbb{N}$, let $[n] = \{1, 2, \ldots, n\}$. Let I_n be the $n \times n$ identity matrix, and $1_{m\times n}$ be the $m \times n$ matrix with all elements being 1. We may also write I when the dimension of the identity matrix is clear in the context. The absolute value of a real number is denoted by $|\cdot|$. For a vector $x \in \mathbb{R}^n$, its ℓ_1-norm is

$$\|x\| = \sum_{i=1}^{n} |x_i| . \tag{1}$$

The ℓ_1 induced norm of a matrix $A \in \mathbb{R}^{m\times n}$ is

$$\|A\| = \sup_{x\in\mathbb{R}^n, x\neq 0} \frac{\|Ax\|}{\|x\|} = \max_{j\in[n]} \sum_{i\in[m]} |A_{ij}| \tag{2}$$

For two matrices $A = [a_{ij}]_{m\times n}$ and $B = [b_{ij}]_{p\times q}$, the Kronecker product is

$$A \otimes B = \begin{bmatrix} a_{11}B & \ldots & a_{1n}B \\ \vdots & \ddots & \vdots \\ a_{m1}B & \ldots & a_{mn}B \end{bmatrix} . \tag{3}$$

A scalar random variable X obeys Laplace distribution with parameter λ, written as $x \sim \text{Lap}(\lambda)$, if its probability distribution function is

$$f_X(x) = \frac{1}{2\lambda} \exp(-\frac{|x|}{\lambda}). \tag{4}$$

This extends to the n-dimensional case by using the ℓ_1-norm, namely, $x \sim \text{Lap}(\lambda, n)$ if

$$f_X(x) = \left(\frac{1}{2\lambda}\right)^n \exp(-\frac{\|x\|}{\lambda}). \tag{5}$$

Obviously, the components $X_i \sim \text{Lap}(\lambda)$ are independent 1-dimensional Laplace noise.

3. DIFFERENTIAL PRIVACY AND ESTIMATION IN PARAMETRIC PROBABILISTIC MODELS

Consider a parametric probabilistic model $\mathcal{M}$, in which the probability distribution $f_X(x; \theta)$ of the observation $X \in \mathbb{R}^n$ depends on some parameters $\theta \in \mathbb{R}^n$ of the same dimension, which we want to keep differentially private. The requirement of differential privacy gives an upper bound on how much the probability distribution function $f_X(x; \theta)$ can change with the private data θ. According to the classic result of Cramer-Rao lower bound, this leads to a lower bound on the variance of unbiased estimation on θ from X. In this section, we will formally study this trade-off between ε-differential privacy and minimum-variance unbiased estimation.

For further discussion, we introduce the following assumptions.

Assumption 1. In the parametric probabilistic model $\mathcal{M}$,

(A1) $f_X(x; \theta)$ is absolutely continuous in both x and θ;
(A2) $f_X(x; \theta)$ is supported on $\mathbb{R}^{2n}$;
(A3) $\mathbb{E}[X] = \theta$, namely,

$$\int_{\mathbb{R}^n} x f_X(x; \theta) \mathrm{d}x = \theta. \tag{6}$$

3.1 Differential Privacy in Differential Form

Unlike most other literature where the ℓ_0 norm is adopted to measure the change in the parameter θ, we adopt the ℓ_1 norm in this work, as with the study in location privacy Andrés et al. (2013); Chatzikokolakis et al. (2013), for the following two reasons. First, the ℓ_1 norm is a finer metric on θ, which we believe is more natural when the variables are real-valued. Consequently, Definition 1 is stricter than the usual definition of ε-differential privacy using the ℓ_0 norm. In addition, using ℓ_1 allows us to treat the components in vector θ separately (see Section 2).

Definition 1. The parametric probabilistic model $\mathcal{M}$ is ε-differentially private, if for any $\theta, \theta', x \in \mathbb{R}^n$ and $O \subseteq \mathbb{R}^n$, the probability distribution function satisfies

$$\int_O f_X(x; \theta) \mathrm{d}x \leq e^{\varepsilon\|\theta-\theta'\|} \int_O f_X(x; \theta') \mathrm{d}x. \tag{7}$$

The degree of privacy increases as ε decreases. For $\varepsilon \to \infty$, any parametric probabilistic model becomes ε-differentially private; for $\varepsilon = 0$, only the parametric probabilistic models that are independent of θ are ε-differentially private.

As shown in Wang et al. (2014, 2017), when the probability distribution function $f_X(x;\theta)$ is smooth enough, the definition of ε-differential privacy can be written in differential form by letting $\theta' \to \theta$ in Definition 1.

Lemma 1. When Assumption (A1) holds, the parametric probabilistic model $\mathcal{M}$ is ε-differentially private, if and only if for any $\theta \in \mathbb{R}^n$ and almost any $x \in \mathbb{R}^n$, the probability distribution function satisfies

$$\left| \frac{\partial}{\partial \theta_i} \ln f_X(x;\theta) \right| \leq \varepsilon. \tag{8}$$

From Lemma 8 in the Appendix, we know that

$$\text{supp}_x f_X(x;\theta) = \{x | f_X(x;\theta) \neq 0\} \tag{9}$$

is independent of the value of θ except for a measure-zero set.

3.2 Performance Bound on Minimum-variance Unbiased Estimator

When Assumption (A3) holds, since X is a complete sufficient statistic for estimating θ, the minimum-variance unbiased estimator of θ is $\hat{\theta}(X) = X$. Lemma 1 shows that when the ℓ_1 norm is used, the condition of ε-differential privacy gives an upper bound on the absolute value of the components of the score $\ln f_X(x;\theta)$. This implies a connection with the Cramer-Rao lower bound of the minimum-variance unbiased estimator (MVUE). When $n = 1$, namely, X and θ are scalars, the fisher information satisfies

$$I_X(\theta) = \mathbb{E}\left[\left(\frac{\partial}{\partial \theta} \ln f_X(x;\theta) \right)^2 \middle| \theta \right] \leq \varepsilon^2. \tag{10}$$

Therefore, by the Cramer-Rao bound, the variance of the MVUE $\hat{\theta}(X) = X$ is bounded by

$$\text{Var}_\theta(\hat{\theta}) \geq \frac{1}{\varepsilon^2}. \tag{11}$$

However, the Cramer-Rao lower bound is not tight, since its attainment condition cannot be satisfied. For the scalar case $n = 1$, we can give a tight minimax lower bound on the variance of $\hat{\theta}$ together with the attainment condition below.

Proposition 2. When Assumptions (A1)-(A3) hold and X and θ are scalars, if θ is ε-differentially private in the probabilistic model $\mathcal{M}$, then the variance of MVUE $\hat{\theta}$ satisfies

$$\max_{\theta \in \mathbb{R}} \text{Var}_\theta(\hat{\theta}) \geq \frac{2}{\varepsilon^2}. \tag{12}$$

The equality holds when $\hat{\theta}(X) = X$ and the probabilistic model is of the form

$$f_X(x;\theta) = \frac{\varepsilon}{2} e^{-\varepsilon |x-\theta|}. \tag{13}$$

We note that the attainment condition in Proposition 2 is sufficient and may not be necessary. Proposition 2 can be generalized to vectors.

Proposition 3. When Assumptions (A1)-(A3) hold and X and θ are scalars, if θ is ε-differentially private in the probabilistic model $\mathcal{M}$, then the variance of each component $\hat{\theta}_i$ of the MVUE $\hat{\theta}$ satisfies

$$\max_{\theta \in \mathbb{R}^n} \text{Var}_\theta(\hat{\theta}_i) \geq \frac{2}{\varepsilon^2}. \tag{14}$$

The equality holds when $\hat{\theta}(X) = X$ and the probabilistic model is of the form

$$f_X(x;\theta) = \frac{\varepsilon}{2} e^{-\varepsilon \|x-\theta\|}. \tag{15}$$

The optimal probabilistic model given by Proposition 3 can be written as

$$X = \theta + w \tag{16}$$

where $w \sim \text{Lap}(1/\varepsilon, n)$. Namely, X is derived by adding independently identically distributed Laplace noise to each component of θ.

Proposition 3 can be extended to the case $\theta = \mathbb{E}[g(X)]$ with a change of variable $Y = g(X)$.

Theorem 4. When Assumptions (A1)-(A3) hold and $\theta = \mathbb{E}[g(X)]$, if θ is ε-differentially private in the probabilistic model $\mathcal{M}$, then the variance of each component $\hat{\theta}_i$ of the MVUE $\hat{\theta}$ satisfies

$$\max_{\theta \in \mathbb{R}^n} \text{Var}_\theta(\hat{\theta}_i) \geq \frac{2}{\varepsilon^2}. \tag{17}$$

The equality holds when $\hat{\theta}(X) = g(X)$ and the probabilistic model is of the form

$$f_X(x;\theta) = \frac{\varepsilon |J(g)|}{2} e^{-\varepsilon \|g(x)-\theta\|}. \tag{18}$$

where $J(g)$ is the determinant of the Jacobian matrix of $g(x)$.

In particular, when $g(x) = Kx$ for $K \in \mathbb{R}^{n \times n}$, Assumption (A2) implies that K is invertible. Then the optimal probabilistic model is given by

$$X = K^{-1}(\theta + w) \tag{19}$$

Theorem 4 shows that there is a trade-off between privacy and the variance of the MVUE. In the following section, we will extend this trade-off to multi-agent control systems.

4. DIFFERENTIAL PRIVACY AND ESTIMATION IN MULTI-AGENT SYSTEMS

Consider a discrete-time linear system with N agents, for example, a simple power grid in which N power generators provide power supply to a single load. The dynamics of the agents are discrete-time linear and coupled via some environmental state $x(t)$,

$$y_i(t+1) = a y_i(t) + b\theta_i(t) + cx(t), \tag{20}$$

where $y_i(t) \in \mathbb{R}^n$ is the state of the i^{th} agent, $\theta_i(t) \in \mathbb{R}^n$ is its local input, and a, b, and c are matrices of proper dimensions. Each agent has an influence on the environmental state $x(t)$. In this work, for simplicity, we assume the environmental state to be the average of all the agents' state.

$$x(t) = \frac{1}{N} \sum_{i=1}^{N} y_i(t). \tag{21}$$

The results presented later can be generalized to the case of $x(t)$ being a linear combination of the individual agents'

state. In the power grid example, when $a = 1$, we may view (20) as the discretized swing equation for the power generators (Bergen, 2009). Specifically, $y_i(t)$ is the angular velocity of the rotor of the i^{th} generator; $\theta_i(t)$ is the mechanical torque driving the rotor; and $x(t)$ is the average angular velocity of all the generators, which is proportional to the current in the circuit.

To control each agent, for example, to stabilize the angular velocity of the rotors at a given value, some knowledge on the environment state $x(t)$ is desirable. However, sharing $x(t)$ directly between the agents jeopardizes privacy – the initial state $y_i(0)$ and the local control inputs $\theta_i(t)$ of an agent can be inferred by observing the environmental states $x(t)$. This situation is undesirable when, for example, the generators are belonged to different power companies who would like to keep the data of their power generators private.

To keep the profiles of the agents

$$D = \{D_i\}_{i\in[N]} = \{(y_i(0), \theta_i(0), \dots, \theta_i(T-1))\}_{i\in[N]} \quad (22)$$

private from observing the sequence of the environmental states, a commonly used approach is to randomize the environmental state $x(t)$ by adding a mean-zero noise $n(t)$ before sharing,

$$X(t) = x(t) + n(t), \quad (23)$$

which we refer to as the randomized environmental state.

Using Kronecker product, we can write the aggregated dynamics of the system as

$$y(t+1) = Ay(t) + B\theta(t) + Cx(t), \quad (24)$$

where $y(t) = (y_1(t), \dots, y_N(t))$ and $\theta(t) = (\theta_1(t), \dots, \theta_N(t))$ are the aggregated state and the aggregated input, respectively, and

$$\begin{aligned} &A = I_N \otimes a, \quad B = I_N \otimes b, \\ &C = 1_{N\times 1} \otimes c, \quad K = 1_{N\times N} \otimes \frac{c}{N}. \end{aligned} \quad (25)$$

From (20) and (23), given the noise $\{n(t)\}_{t\le T}$, the probability distribution of the sequence of the environmental states $X = \{X(t)\}_{t\le T}$ is uniquely determined by the agents' profiles D. Therefore, we can treat the multi-agent control system as a probabilistic model parametrized by the agents' profiles D as discussed in Section 3.

4.1 Differential Privacy in Multi-agent Systems

To be consistent with other literature, we first introduce the notion of adjacency. Given two profiles of the agents D and D', they are called *adjacent* if there exists $i \in [N]$ such that for all $j \neq i \in [N]$, $D_j = D'_j$; namely, if they differ only in the i^{th} agent's profile. With the notion of adjacency, the ε-differential privacy of the agents' profiles in the multi-agent control system is given below.

Definition 2. Given a time horizon $T > 0$ and $\varepsilon > 0$, the multi-agent control system is ε-differentially private up to time T, if for any adjacent profiles D' and $\mathcal{X} \subseteq \mathbb{R}^{nN(T+1)}$,

$$\mathbb{P}\left[X_D \in \mathcal{X}\right] \le e^{\varepsilon \|D-D'\|} \mathbb{P}\left[X_{D'} \in \mathcal{X}\right], \quad (26)$$

where X_D is the sequence of environmental states corresponding to the agents' profiles D.

Since the ℓ_1 norm is used, the adjacency condition in Definition 2 may be removed. The following remark follows directly from the Definition 2.

Remark 5. If the multi-agent system is ε-differentially private up to time T, then it is ε-differentially private up to any time $S \le T$.

Definition 2 can also be written in differential form by letting $D' \to D$.

Lemma 6. For $T > 0$ and $\varepsilon > 0$, the agents' profiles D is kept ε-differentially private up to time T, if for any agents' profiles D' and almost every $x \in \mathbb{R}^{nN(T+1)}$,

$$f_X(x; D) \le e^{\varepsilon \|D-D'\|} f_X(x; D), \quad (27)$$

where $f_X(x; D)$ is the probability distribution of the sequence of environmental states X for the agents' profiles D.

Lemma 6 is consistent with Lemma 1, when we view the multi-agent control system as a probabilistic model parametrized by the agents' profiles D.

4.2 Performance Bound on Estimating the Environment State

In this section, consider the problem of estimating the average profile

$$\bar{D} = \big(\bar{x}(0), \bar{\theta}(0), \dots, \bar{\theta}(T-1)\big) = \frac{\sum_{i=1}^N D_i}{N}, \quad (28)$$

while keeping the agent's profiles D ε-differentially private. As we will show later in (29)-(30), the average profile $\bar{D}$ determines the evolution of the environmental state $x(t)$, hence is desirable for designing a local control law for the agents. Let $\hat{D}(X(0), \dots, X(T))$ be the minimum-variance unbiased estimator of the average profile $\bar{D}$ from the sequence of randomized environmental states $X \in \mathbb{R}^{n(T+1)}$.

Unlike the case in Section 3.2, here the private data, namely the agents' profiles $D \in \mathbb{R}^{nN(T+1)}$, and the observation, namely the sequence of randomized environmental states $X \in \mathbb{R}^{n(T+1)}$ are of different size. Also, the observation X is generated by a multi-agent control system.

Multiplying $1_{1\times n}/N$ on both sides of (24) gives

$$X(0) = \bar{x}(0) + n(0), \quad (29)$$

$$\begin{aligned} X(t+1) - (a+c)X(t) &= \frac{b}{N}\bar{\theta}(t) \\ &+ n(t+1) - (a+c)n(t), \quad t \le T-1. \end{aligned} \quad (30)$$

For further discussion, we introduce the following assumption.

Assumption 2. In the multi-agent control system,

(A4) the matrix b in (20) is invertible;

Assumption (A4) is adopted to ensure that the average profile can be recovered precisely from the sequence of environmental states X, when there is no noise.

By Definition 2 and noting that

$$\|\bar{y}(0)\| \le \frac{1}{N}\|y(0)\|, \quad (31)$$

$$\|\bar{\theta}(t)\| \le \frac{1}{N}\|\theta(t)\|, \quad t \le T-1, \quad (32)$$

if the agents' profiles D is ε-differentially private, then the average profile $\bar{D}$ is $(N\varepsilon)$-differentially private.

As shown in (29)-(30), protecting the ε-differential privacy of $\bar{y}(0)$ and $\bar{\theta}(0), \ldots, \bar{\theta}(T-1)$ are decoupled: $\bar{y}(0)$ is randomized by $n(0)$, and $\bar{\theta}(t)$ is randomized by $n(t+1) - (a+c)n(t)$ for $t < T$. Therefore, applying Theorem 4 iteratively gives the following theorem.

Theorem 7. When Assumption (A4) holds, if the agents' profiles D is ε-differentially private in the multi-agent control system, then each component $\hat{D}_i$ of the minimum-variance unbiased estimator $\hat{D}$ of the average profile $\bar{D}$ satisfies

$$\max_{D \in \mathbb{R}^{nN(T+1)}} \mathrm{Var}_D(\hat{D}_i) \geq \frac{2}{N^2\varepsilon^2}. \tag{33}$$

The minimum is achieved when in (23),

$$n(0) = w(0), \tag{34}$$

$$n(t+1) = (a+c)n(t) + bw(t), \tag{35}$$

where $w(t) \sim \mathrm{Lap}(1/\varepsilon, n)$ for $t = 0, \ldots, T-1$.

In Theorem 7, as the number of agents N or the parameter ε grows, the estimation on the average profile becomes increasingly accurate, while the profiles of all agents are still kept ε-differentially private. This is in accordance with other literature in differential privacy. Similar to Wang et al. (2014, 2017), the lower bound is independent of T. This is because ℓ_1 norm is used in the definition of differential privacy.

5. CONCLUSION

In this work, we study the trade-off between ε-differential privacy and the minimum-variance unbiased estimation first in a parametric probabilistic model. Then, we extend the results to a discrete-time linear multi-agent control system, where the agents are coupled via an environmental state. We show that when the agents' profiles, which is the combination of the agents' initial states and the sequence of local control inputs over time, is ε-differentially private, there is a lower bound on the variance of the minimum-variance unbiased estimator of the environmental state. This lower bound is achieved by a randomized mechanism that uses Laplace noise.

REFERENCES

Andrés, M.E., Bordenabe, N.E., Chatzikokolakis, K., and Palamidessi, C. (2013). Geo-indistinguishability: Differential privacy for location-based systems. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, 901–914. ACM.

Apple Inc (2016). What's new in iOS 10.0. URL `https://developer.apple.com/library/prerelease/content/releasenotes/general/whatsnewinios/articles/ios10.html`.

Bergen, A.R. (2009). *Power systems analysis*. Pearson Education India.

Chatzikokolakis, K., Andrés, M.E., Bordenabe, N.E., and Palamidessi, C. (2013). Broadening the scope of differential privacy using metrics. In *Privacy Enhancing Technologies*, 82–102. Springer.

Dwork, C. (2006). Differential privacy. In *Automata, languages and programming*, 1–12. Springer.

Dwork, C., Kenthapadi, K., McSherry, F., Mironov, I., and Naor, M. (2006). Our data, ourselves: Privacy via distributed noise generation. In S. Vaudenay (ed.), *Advances in Cryptology - EUROCRYPT 2006*, volume 4004 of *Lecture Notes in Computer Science*, 486–503. Springer Berlin Heidelberg.

Farhangi, H. (2010). The path of the smart grid. *IEEE power and energy magazine*, 8(1), 18–28.

Geng, Q., Kairouz, P., Oh, S., and Viswanath, P. (2015). The staircase mechanism in differential privacy. *Selected Topics in Signal Processing, IEEE Journal of*, 9(7), 1176–1184.

Hale, M. and Egerstedty, M. (2015). Differentially private cloud-based multi-agent optimization with constraints. In *American Control Conference (ACC), 2015*, 1235–1240.

Huang, Z., Mitra, S., and Dullerud, G. (2012). Differentially private iterative synchronous consensus. In *Proceedings of the 2012 ACM workshop on Privacy in the electronic society*, WPES '12, 81–90. ACM, New York, NY, USA.

Huang, Z., Mitra, S., and Vaidya, N. (2015). Differentially private distributed optimization. In *Proceedings of the 2015 International Conference on Distributed Computing and Networking*, 4. ACM.

Huang, Z., Wang, Y., Mitra, S., and Dullerud, G.E. (2014). On the cost of differential privacy in distributed control systems. In *Proceedings of the 3rd international conference on High confidence networked systems*, 105–114. ACM.

Koufogiannis, F., Han, S., and Pappas, G.J. (2014). Computation of privacy-preserving prices in smart grids. In *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, 2142–2147. IEEE.

Le Ny, J. and Pappas, G. (2014). Differentially private filtering. *Automatic Control, IEEE Transactions on*, 59(2), 341–354.

Li, C., Hay, M., Rastogi, V., Miklau, G., and McGregor, A. (2010). Optimizing linear counting queries under differential privacy. In *Proceedings of the twenty-ninth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, PODS '10, 123–134. ACM, New York, NY, USA.

Ma, C.Y., Yau, D.K., Yip, N.K., and Rao, N.S. (2013). Privacy vulnerability of published anonymous mobility traces. *IEEE/ACM Transactions on Networking*, 21(3), 720–733.

Masoum, A.S., Deilami, S., Moses, P., Masoum, M., and Abu-Siada, A. (2011). Smart load management of plug-in electric vehicles in distribution and residential networks with charging stations for peak shaving and loss minimisation considering voltage regulation. *IET generation, transmission & distribution*, 5(8), 877–888.

McSherry, F. and Talwar, K. (2007). Mechanism design via differential privacy. In *Foundations of Computer Science, 2007. FOCS '07. 48th Annual IEEE Symposium on*, 94 –103.

Mo, Y. and Murray, R. (2014). Privacy preserving average consensus. In *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, 2154–2159.

Nozari, E., Tallapragada, P., and Cortés, J. (2015). Differentially private average consensus with optimal noise selection. *IFAC-PapersOnLine*, 48(22), 203–208.

Shokri, R., Theodorakopoulos, G., Le Boudec, J.Y., and Hubaux, J.P. (2011). Quantifying location privacy. In *2011 IEEE Symposium on Security and Privacy*, 247–262. IEEE.

Wang, G., Wang, B., Wang, T., Nika, A., Zheng, H., and Zhao, B.Y. (2016). Defending against sybil devices in crowdsourced mapping services. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*, 179–191. ACM.

Wang, Y., Huang, Z., Mitra, S., and Dullerud, G. (2014). Entropy-minimizing mechanism for differential privacy of discrete-time linear feedback systems. In *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, 2130–2135.

Wang, Y., Huang, Z., Mitra, S., and Dullerud, G.E. (2017). Differential privacy in linear distributed control systems: Entropy minimizing mechanisms and performance tradeoffs. *IEEE Transactions on Control of Network Systems*, 4(1), 118–130.

Xue, M., Wang, W., and Roy, S. (2014). Security concepts for the dynamics of autonomous vehicle networks. *Automatica*, 50(3), 852–857.

Appendix A. PROOF OF PROPOSITION 2

Let's first look at the case $n = 1$; namely, x and θ are scalars. To minimize $\max_{\theta\in\mathbb{R}^n} \mathrm{Var}_\theta(\hat{\theta})$, the probability distribution function $f_X(x;\theta)$ of the probabilistic model $\mathcal{M}$ should satisfy the following properties.

Lemma A.1. The maximum variance of $\hat{\theta}$ is minimized when

$$f_X(x;\theta) = f_X(2\theta - x;\theta). \tag{A.1}$$

Proof. If not, consider the probabilistic model $\mathcal{M}'$ with the probability distribution function $f'(x;\theta)$ given by

$$f'_X(x;\theta) = \begin{cases} f_X(x;\theta) & \text{if } x > \theta, V_+ \le V_- \\ & \text{or } x \le \theta, V_+ > V_-, \\ f_X(2\theta - x;\theta) & \text{if } x > \theta, V_+ > V_- \\ & \text{or } x \le \theta, V_+ \le V_-, \end{cases} \tag{A.2}$$

where

$$V_+ = \max_{\theta\in\mathbb{R}} \int_\theta^\infty x^2 f_X(x;\theta)\mathrm{d}x, \tag{A.3}$$

$$V_- = \max_{\theta\in\mathbb{R}} \int_{-\infty}^\theta x^2 f_X(x;\theta)\mathrm{d}x. \tag{A.4}$$

By construction, the probabilistic model $\mathcal{M}'$ is ε-differentially private and satisfies Assumption (A1)-(A3). Noting that $\max_{\theta\in\mathbb{R}^n} \mathrm{Var}_\theta X$ is smaller in $\mathcal{M}'$, the lemma holds.

Lemma A.2. The maximum variance of $\hat{\theta}$ is minimized when for any $\alpha \in \mathbb{R}$,

$$f_X(x;\theta) = f_X(2\alpha - x; 2\alpha - \theta). \tag{A.5}$$

Proof. If not, consider the probabilistic model $\mathcal{M}''$ with the probability distribution function $f''(x;\theta)$ given by

$$f''_X(x;\theta) = \begin{cases} f_X(x;\theta) & \text{if } \theta \le \alpha, W_+ > W_- \\ & \text{or } \theta > \alpha, W_+ \le W_-, \\ f_X(2\alpha - x; 2\alpha - \theta) & \text{if } \theta > \alpha, W_+ > W_- \\ & \text{or } \theta \le \alpha, W_+ \le W_-. \end{cases} \tag{A.6}$$

where

$$W_+ = \max_{\theta>\alpha} \int_{\mathbb{R}} x^2 f_X(x;\theta)\mathrm{d}x, \tag{A.7}$$

$$W_- = \max_{\theta\le\alpha} \int_{\mathbb{R}} x^2 f_X(x;\theta)\mathrm{d}x. \tag{A.8}$$

By Lemma A.1, we assume that $f_X(x;\theta) = f_X(2\theta - x;\theta)$. Then, the probabilistic model $\mathcal{M}''$ is ε-differentially private and satisfies Assumption (A1)-(A3). Noting that $\max_{\theta\in\mathbb{R}^n} \mathrm{Var}_\theta X$ is smaller in $\mathcal{M}''$ for all θ, the lemma holds.

Combining Lemma A.1 and Lemma A.2, we know that $\max_{\theta\in\mathbb{R}} \mathrm{Var}_\theta(\hat{\theta})$ is minimized when

$$\begin{aligned} f_X(x;\theta) &= f_X(2\theta - x;\theta) \\ &= f_X(x + 2\alpha - 2\theta; 2\alpha - \theta) = f_X(x + \beta; \theta + \beta), \end{aligned} \tag{A.9}$$

for any $\alpha \in \mathbb{R}$ and $\beta = 2\alpha - 2\theta$. Therefore, we can write

$$f_X(x;\theta) = f(x - \theta) \tag{A.10}$$

and accordingly

$$\max_{\theta\in\mathbb{R}} \mathrm{Var}_\theta(\hat{\theta}) = 2\int_0^\infty x^2 f(x)\mathrm{d}x. \tag{A.11}$$

To minimize $\max_{\theta\in\mathbb{R}} \mathrm{Var}_\theta(\hat{\theta})$, $f(x)$ should decrease on $[0,\infty)$ at the fastest speed, namely, $f'(x) = -\varepsilon f(x)$. In this case, we have

$$f_X(x;\theta) = \frac{\varepsilon}{2} e^{-\varepsilon|x-\theta|} \tag{A.12}$$

and

$$\min_{\mathcal{M}} \max_{\theta\in\mathbb{R}} \mathrm{Var}_\theta(\hat{\theta}) = \frac{2}{\varepsilon^2}. \tag{A.13}$$

Appendix B. PROOF OF PROPOSITION 3

For simplicity, let $n = 2$. The probability distribution function of the parametric probabilistic model $\mathcal{M}$ is $f_{X_1,X_2}(x_1,x_2;\theta_1,\theta_2)$. For any fixed θ_2, the marginal probability distribution function of X_1 satisfies that

$$\left| \frac{\partial}{\partial\theta_1} \ln f_{X_1}(x_1;\theta_1,\theta_2) \right| \le \varepsilon. \tag{B.1}$$

Therefore, by Proposition 2, the MVUE $\hat{\theta}_1(X_1,X_2) = X_1$ satisfies

$$\max_{\theta_1,\theta_2\in\mathbb{R}} \mathrm{Var}_{\theta_1,\theta_2}(\hat{\theta}_1) \ge \frac{2}{\varepsilon^2}. \tag{B.2}$$

Similarly, the MVUE $\hat{\theta}_2(X_1,X_2) = X_2$ satisfies

$$\max_{\theta_1,\theta_2\in\mathbb{R}} \mathrm{Var}_{\theta_1,\theta_2}(\hat{\theta}_2) \ge \frac{2}{\varepsilon^2}. \tag{B.3}$$

Therefore, we can derive (14), in which the equality holds if and only if the equalities in (B.2)-(B.3) hold and $\mathrm{Cov}(X_1,X_2) = 0$. This is attained when (15) holds.