On the Cost of Differential Privacy in Distributed Control Systems

Zhenqi Huang Yu Wang Sayan Mitra Geir Dullerud
{zhuang25, yuwang8, mitras, dullerud}@illinois.edu
Coordinate Science Laboratory
University of Illinois at Urbana Champaign
Urbana, IL 61801

ABSTRACT

Individuals sharing information can improve the cost or performance of a distributed control system. But, sharing may also violate privacy. We develop a general framework for studying the cost of differential privacy in systems where a collection of agents, with coupled dynamics, communicate for sensing their shared environment while pursuing individual preferences. First, we propose a communication strategy that relies on adding carefully chosen random noise to agent states and show that it preserves differential privacy. Of course, the higher the standard deviation of the noise, the higher the cost of privacy. For linear distributed control systems with quadratic cost functions, the standard deviation becomes independent of the number agents and it decays with the maximum eigenvalue of the dynamics matrix. Furthermore, for stable dynamics, the noise to be added is independent of the number of agents as well as the time horizon up to which privacy is desired. Finally, we show that the cost of ϵ-differential privacy up to time T, for a stable system with N agents, is upper bounded by $O(\frac{T^3}{N\epsilon^2})$.

Keywords

Differential Privacy, Distributed Control, Cyber-physical Security

1. INTRODUCTION

In distributed control systems there is a trade-off between privacy and cost. A vehicle with a smart navigation device may provide some information about its trajectory for crowd sourced traffic data collection [7]. Aggregates of this data can then be used by the navigation device for traffic-aware routing. Similarly, a consumer in a power-grid may share some information about her energy demands to then use the aggregate demands for deciding her own consumption plans and save energy costs. At one extreme is the completely private society in which agents neither share nor receive any information through communication. They only interact through their coupled dynamics. The other extreme is the completely non-private or broadcast society. Agents share complete information which at least in principle allows, all agents to make accurate predictions (e.g., traffic or electricity demands) and to make optimal decisions. Between these two extremes lie a multitude of other possible communication strategies. The privacy-cost trade-off can formalized as the *cost of privacy* measured by the difference between the cost achieved through a given communication strategy and the cost achieved by the completely non-private strategy.

In this paper, we present a general framework for studying cost of privacy for distributed control systems in which a collection of agents pursue individual goals and communicate for the purpose of sensing their shared environment. Each agent i has a *preference* p_i—an infinite sequences of points that it wants to visit in a Euclidean space. These preference capture, for example, a sequence of waypoints for a vehicle or the electric power demand of a household. The evolution of an agent depends on (a) its dynamics, (b) the control action it takes, and also (c) the environment or the aggregate state of the other agents. If the communication strategy shares more information about its preference, then all agents in the society can estimate the environment more accurately, and therefore, make better control decisions. On the other hand, such a communication strategy may leak information about agent preference. For a given communication strategy (r), the difference between the actual sequence of states visited by an agent following r and the preferred trajectory p_i defines the cost incurred by the agent i.

In our formulation, once the underlying dynamics of the system, the individual preferences, and the communication strategy are fixed the overall system is deterministic[1]. We show (Proposition 1) that knowing the preference for the agents and an observation of the system allows an adversary to uniquely infer the complete state trajectory of an agent over time. That is, there is a one-to-one correspondence between the observation sequences and the state trajectories. Therefore, protecting the privacy of the state trajectories is tantamount to protecting the preferences.

The notion of privacy we adopt in this paper is differential privacy [1, 2] as applied to continuous bit streams in [4]. We have to make two technical adjustments to the earlier

[1]If the communication strategy uses randomization, the the overall system is purely probabilistic.

definition. First, since preferences are infinite sequences, we define adjacency of preferences over a time horizon, say T. Secondly, we define a pair of agent preferences p and p' to be *adjacent up to time* T, if they differ about the preference of at most one agent, say agent i, and for any time before T, the L^1-norm of difference $p_i - p_i'$ at that time is bounded. The resulting notion of ϵ-differential privacy then ensures that an adversary with access to all the communication in the system—we call this an *observation sequence*—cannot gain information about the preference of any agent up to time T with any significant probability.

Contributions. (1) We present an ϵ-differential privacy preserving communication strategy for distributed control systems. For a given privacy parameter ϵ, and a given agent state x_i, the strategy adds noise to x_i from a Laplace distribution with zero mean and standard deviation proportional to the sensitivity of a particular function to changes in agent preferences. The function in question maps observations to sets of executions η^{-1} (see Section 4).

(2) We show that for linear distributed systems with quadratic cost functions, the sensitivity of η^{-1}, and therefore the required standard deviation of the noise, is independent of the number of participating agents (see Theorem 4). Roughly, the sensitivity of η^{-1} with respect to the changes in an individual agent's preference is influenced by the number of agents N in two opposing directions. As N increases, a larger number of agents are influenced by the changes in the preference of an individual. In contrast, with larger N, the influence of i on another individual agent through the environment weakens in an environment which aggregates the state of all the agents. In the linear case, these two effects roughly cancel out making the sensitivity independent of N.

(3) We show that the required standard deviation of noise decreases with the stability of the dynamics and with the weakening of the environment's influence on an individual. When the modulus of the maximum eigenvalue of the dynamics matrix is smaller, the effect of changes in an individual's preference on the system's trajectory decays faster over time. Furthermore, as the time horizon goes to infinity, the sensitivity converges to a constant for stable systems. Thus, the amount of noise to be added for differentially private communication is independent of the number of participating agents as well as the time horizon up to which privacy is desired. For unstable dynamics, on the other hand, the sensitivity can grow exponentially with the time bound T.

(4) We establish that the cost of ϵ-differential privacy using our communication strategy up to time T for a system with N agents is at most $O(\frac{T^3}{N\epsilon^2})$ for stable systems and the cost can also grow exponentially with T for unstable systems. This suggests that the proposed strategy is more likely to be useful for stable systems with short-lived participants (e.g., drivers with short commutes), and further research is needed for strategies that scale better with time.

Organization. Section 2 discusses related research. Section 3 defines Markov chains with observation maps and their semantics. Section 4 develops the general framework. Section 5 presents the Laplace communication mechanism for general distributed control systems and Section 6 develops specialized results for linear systems. Several proofs are omitted because of limited space; details can be found in the online version of the paper [9].

2. RELATED WORK ON DIFFERENTIAL PRIVACY

While there are several notions of data privacy in the Computer Science literature, the quantitative and statistical nature of differential privacy makes it suitable for adoption in distributed control. The notion used in this paper follows the definition of differential privacy introduced originally in the context of statistical queries on databases [1] (see [2] for a survey). Differential privacy requires that the change of an individual agent's data can only result in *unsubstantial* changes in the statistics of any output. It follows that an adversary looking at the output of any analysis cannot reason with high confidence about the individual's data. Various mechanisms for achieving differential privacy have been studied in the literature [3, 12, 13]. The Laplace mechanism requires adding a Laplace noise to the query output and was proposed in [1]. In the recent paper [5], a staircase mechanism is shown to be the optimal noise-adding mechanism in terms of maximizing the accuracy of a query. In [4], the notion of differential privacy is expanded to include streaming and online computations in which the adversary can look at the entire sequence of outputs from the analysis algorithm.

Our work is concerned with protecting the privacy of the agent's states and preferences instead of its participation status. Consequently, like the definitions presented in [6, 15], we define differential privacy in terms of adjacent preferences that are identical for all agents excepting one agent whose preferences are close as measured by the L^1 norm. Our notion of ϵ-differential privacy ensures that an adversary with access to all the communication in the system cannot gain information about the preference of any agent up to time T with any significant probability.

In the paper [10], the authors develop a notion of differential privacy which ensures that an adversary cannot tell the exact input to a dynamical system by looking at its output stream. Laplace and Gaussian mechanisms are presented for converting an ordinary dynamical system to a differentially private one. Unlike our message-based and distributed implementation schemes, here the privacy-preserving implementation consists of a filter and an estimator which are designed to minimize the mean-squared error from the outputs of the ideal system. In the follow-up work [11], a Kalman filter is designed to estimate the states of differentially private systems with minimized error. The sufficient condition of the minimization problem is established in the form of linear matrix inequalities. The problem studied in this paper is different from the ones introduced in these two papers in several ways. First, in the class of systems studied here, an agent's dynamics is coupled with the environment which depends on the aggregate of all other agents' states. Secondly, these systems are "closed loop" and the noise added for privacy in one round affects all future states of the system.

The results in this paper generalize our previous work on differentially private iterative consensus [8] where the agent are required to converge to a common value while preserving the privacy of their initial values. Setting the coupling to zero, it is possible to recover the mechanism in [8] from the communication strategy proposed here.

3. PRELIMINARIES

For a set S, S^k and S^ω denote the k-ary and infinite Cartesian products of S. For a natural number $N \in \mathbb{N}$, we denote

the set $\{1, \ldots, N\}$ by $[N]$.

For a set S, $\mathcal{P}(S)$ is the set of probability distributions over S. For a random variable X taking values in S with distribution $\mu \in \mathcal{P}(S)$, we write $X \sim \mu$. The mechanisms presented in this paper rely on random real numbers drawn according to the Laplace distribution. $Lap(b)$ denotes the Laplace distribution with probability density function $p_L(x|b) = \frac{1}{2b} e^{-|x|/b}$. This distribution has mean 0 and variance $2b^2$.

For a vector v of length N, the i^{th} component is denoted by v_i. For a vector v in $\mathbb{R}^n$, $|v|_p$ $(1 \leq p \leq \infty)$ stands for the standard L^p-norm for v. Without a subscript, $|\cdot|$ stands for L^1-norm by default. For a matrix $A \in \mathbb{R}^{m \times n}$, $|A|_p = \max_{|v|_p=1} |Av|_p$ stands for the standard induced p-norm of matrix A. Without a subscript, $|A|$ stands for induced 1-norm of A.

A matrix $K \in \mathbb{R}^{n \times n}$ is said to be stable, if the modulus of all the eigenvalues of K are smaller than unity. The smaller the maximum modulus of eigenvalues of K, the more stable K is. If some of the eigenvalues are larger than 1, it is said to be unstable. A property of a stable matrix A is that for any x, $A^t x \to \mathbf{0}$ vector as $t \to 0$.

The behavior of the complete system in this paper is modeled as a Markov chain parametrized by a quantity p, and which produces some observations. This Markov chain $\mathcal{M}(p) = \langle Q, Y, q_0, \mu, \eta \rangle$, where each of the following components may depend on p: (i) Q is a measurable set of states, (ii) Y is a set of observations, (iii) $q_0 \in Q$ is the initial state, (iv) $\mu : Q \to \mathcal{P}(Q)$ is a *probabilistic state transition function*, and (v) $\eta : Q \to Y$ is the *observation function*. We will denote the components of $\mathcal{M}(p)$ by $Q_{\mathcal{M}(p)}, Y_{\mathcal{M}(p)}, \mu_{\mathcal{M}(p)}$, etc.

An *execution* of length k of $\mathcal{M}(p)$ is a sequence of states $\alpha = q_0, q_1, \ldots, q_{k-1}$, such that for each $i \in [k]$, $\mu(q_{k-1}, q_k) > 0$. The probability measure over the space of executions $\mathbb{P}_{\mathcal{M}(p)}$ is defined in the standard way by first defining a σ-algebra of cones over the space of executions, and then by defining the probability of the cones by integrating over μ (see for example [14]).

For the execution α of length k, the corresponding *observation* is a sequence in Y^k obtained by point-wise application the observation function to α, that is, $\eta_{\mathcal{M}(p)}(\alpha) \triangleq \eta_{\mathcal{M}(p)}(q_0), \ldots, \eta_{\mathcal{M}(p)}(q_{k-1})$. For a given observation sequence $\beta \in Y^k$, the corresponding set of executions $\eta^{-1}_{\mathcal{M}(p)}(\beta)$ is defined as the set $\{\alpha \mid \eta_{\mathcal{M}(p)}(\alpha) = \beta\}$ and the functions $\eta_{\mathcal{M}(p)}$ and $\eta^{-1}_{\mathcal{M}(p)}$ are extended to sets of executions and observations in the usual way.

4. DISTRIBUTED CONTROL SYSTEMS

We begin by defining a distributed control system abstractly (see Figure 1) ; Section 6 provides more concrete instantiations of these definitions in terms of linear models. A control system consists of N agents operating in a shared environment. Agent i, $i \in [N]$, has a preference p_i. The agent's behavior consists of a physical part which evolves according to some deterministic dynamics and a controller which computes the control inputs for the physical dynamics. The agent uses a communication strategy r to broadcast some noisy version of its state to the other agents. The broadcasts are noisy to preserve privacy and are used to estimate the state of the environment. These estimates are used by the i's controller for computing the inputs (along side its own state).

Fixing the vector of preferences p for all agents, the evolution of the complete system becomes a stochastic process, specifically a Markov chain with observations $\mathcal{M}(p)$ with the stochasticity arising from the noise values used in the communication strategy of the individual agents.

The Markov chain modeling the distributed control system is specified by the following parameters: (a) Euclidean spaces $\mathcal{X}, \mathcal{U}$ and $\mathcal{Z}$ which define an individual agent's state space, its control input space, and the state space of the environment, respectively. (b) The preferences of each agent i $(p_i \in \mathcal{X}^\omega)$ consisting of a sequence of points in the individual agent's state space $\mathcal{X}$ which defines a path agent i wants to follow. (c) A *dynamics function* $f : \mathcal{X} \times \mathcal{U} \times \mathcal{Z} \to \mathcal{X}$ which defines the next state of an agent as a function of its current state, control input and the environment's state. (d) An *aggregation function* $h : \mathcal{X}^N \to \mathcal{Z}$ which defines the state of the environment as a function of the agent states. (e) A *control function* $g : \mathcal{X} \times \mathcal{Z} \times \mathbb{N} \to \mathcal{U}$ which defines the agent's controller output as a function of its state, the environment state and the current time. And finally (f) A *probabilistic observation map* $r : \mathcal{X} \times \mathcal{Z} \times \mathbb{N} \to \mathcal{P}(\mathcal{X})$ which selects a noisy state observation for an agent as a function of its actual state, its knowledge of the environment and the current time.

The dynamics functions f and the aggregation function h capture the physical behavior of the system and the coupling between agents—as control designers, we cannot change them. In this paper, we assume that the controller function g is obtained through existing control theoretic techniques (see Section 2 for a discussion of related work). The only component up for design is the observation map r. In defining the Markov chain below, we will use r to probabilistically update a state component of the agent (called $\tilde{x}_i$ below) which is produced as an observation. This simplifies our model by keeping the observation function η deterministic.

A state of agent i is a point in $\mathcal{X}^2 \times \mathcal{U}$ and its three components are the true agent state (denoted by x_i), the observed agent state $(\tilde{x}_i)$, and the control input (u_i), respectively. The state of the environment is $\mathcal{Z}^2$ and the two components are the (true) environment state (z) and the observed environment state $(\tilde{z})$. Thus, the state space of the Markov chain modeling the complete system is $Q \triangleq (\mathcal{X}^2 \times \mathcal{U})^N \times \mathcal{Z}^2$. For each $i \in [N]$ the projection functions $x_i, \tilde{x}_i : Q \to \mathcal{X}$, $u_i : Q \to \mathcal{U}$ give the state, the observed state, and the control input of agent i at system state q. Similarly, $z, \tilde{z} : Q \to \mathcal{Z}$ give the environment state and the observed environment state at q. The space of observations for the Markov chain is $Y \triangleq \mathcal{X}^N \times \mathcal{Z}$. The transition probabilities from a state $q_{t-1} \in Q$ at time $t \in \mathbb{N}$ is defined by the following sequence of equations:

$$u_i(q_t) = g(x_i(q_{t-1}), \tilde{z}(q_{t-1}), t) \tag{1}$$
$$x_i(q_t) = f(x_i(q_{t-1}), u_i(q_t), z(q_{t-1})) \tag{2}$$
$$\tilde{x}_i(q_t) \sim r(x_i(q_t), \tilde{z}(q_{t-1}), t) \tag{3}$$
$$z(q_t) = h(x_1(q_t), \ldots, x_N(q_t)) \tag{4}$$
$$\tilde{z}(q_t) = h(\tilde{x}_1(q_t), \ldots, \tilde{x}_N(q_t)) \tag{5}$$

The first three equations define values of the control input (u_i), the agents state (x_i), and the environment state (z_i), for each $i \in [N]$, in the post state q_t as functions of q_{t-1}. The value of $\tilde{x}_i(q_t)$'s is chosen according to the probability distribution $r(x_i(q_t), t)$ and $\tilde{z}_i(q_t)$.

The observation function η of the Markov chain is defined as follows: for any state $q \in Q$, $\eta(q) \triangleq \langle \tilde{x}_1(q), \ldots, \tilde{x}_N(q), \tilde{z}(q) \rangle$. In other words, an observation is simply the projection of the state on the $\tilde{x}$ and $\tilde{z}$ components. We sometimes use $x_i(t)$ as $x_i(q_t)$ for short if the execution is clear in the context.

The initial state q_0 is specified by the global preference vector p and the aggregate function h. For each agent, the initial state is defined by the first point of its preference $x_i(q_0) = p_i(0)$. Then the aggregate $z(q_0) = h(x_1(q_0), \ldots, x_N(q_0))$. The initial control inputs (u_i) and the initial observed agent states $(\tilde{x}_i)$ and initial observed environment state $(\tilde{y})$ are set to 0.

Interpretation of Markov Transitions. From state q_t, the control input $u_i(q_t)$ for agent i is obtained by applying the possibly time-varying control function g to the agent's previous state $x_i(q_{t-1})$ and the observed aggregate $\tilde{z}(q_{t-1})$ at t. The new agent state $x_i(q_t)$ is obtained by applying the dynamics function f to the agent's state at step $t-1$, the state of the environment $z(q_{t-1})$ at $t-1$ and the newly computed control input $u_i(q_t)$. The agent i's observed state $\tilde{x}_i(q_t)$ is updated by choosing a value from the time-varying distribution $r(x_i(q_t), z(q_{t-1}), t)$. This distribution defines how noise is added to the actual value of the state $x_i(q_t)$. The actual environment state and the observed environment state are computed by applying the aggregation function h to the new agent states and the observed agent states. An *execution* of Markov chain $\mathcal{M}(p)$ of length T is a sequence $\alpha = q_0 \ldots q_{T-1}$, where $q_t \in (\mathcal{X}^2 \times \mathcal{U})^N \times \mathcal{Z}^2)$. We denote the set $\mathsf{Exec}_{p,T}$ to be the set of all executions up to time T of the distributed system parameterized by p. The observation up to time T $\eta(\alpha)$ is a sequence of $(\tilde{x}_1, \tilde{x}_2, \ldots, \tilde{x}_N, \tilde{z})$ that lives in the space $Y^T = (\mathcal{X}^N \times \mathcal{Z})^T$. The only sources of uncertainty in the behavior of a control system are (a) the preferences of the agents (p) and (b) the randomized observation map (r) which is used to disseminate noisy private information for the sake better performance. Thus, given a preference vector and an observation sequence, and the knowledge of the parameters f, g, h, it is possible to infer a unique execution of $\mathcal{M}(p)$. We formalize this notion in Proposition 1.

Figure 1: Close-loop distributed system

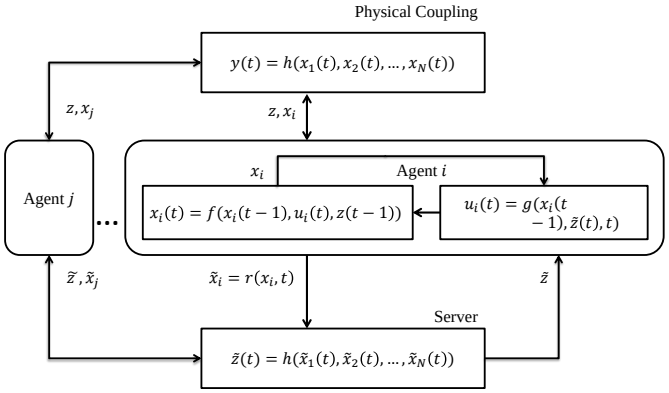

Proposition 1. *For any randomized control mechanism $\mathcal{M}$, given a preference vector p and an observation sequence β of length k, $\eta^{-1}_{\mathcal{M}(p)}(\beta)$ is a singleton set.*

Proof. The proof is by induction on the length of β. If β is of length one then $\eta^{-1}(\beta)$ is the single start state θ. As we mentioned previously, the agent i's state matches the first point of p_i, that is $x_i(\theta) = p_i(0)$ which is specified by p. Also, $z(\theta) = h(x_1(\theta), \ldots, x_N(\theta))$. And other variables are initialized as 0. Thus, the start state θ is fixed.

Suppose $\beta = \beta' y$ be an observation of length $k+1$, where $\eta^{-1}(\beta')$ is the unique execution α' ending with last state $q_k = \langle u, x, \tilde{x}, z, \tilde{z} \rangle$. It suffices to show that for the given state q_k and the observation y, there is a unique state q_{k+1} which makes $\alpha' q_{k+1} = \eta^{-1}(\beta)$. From Equation (1), it follows that for each $i \in [N]$, $u_i(q_{k+1})$ is uniquely defined as $g(x_i(q_k), \tilde{z}(q_k), t+1)$. This and Equation (2) implies that $x_i(q_{k+1})$ is uniquely defined. Similarly, the u_i's and the x_i's together with Equation (4) imply that $z(q_{k+1})$ is also uniquely defined. Finally, $\tilde{x}_i(q_{k+1}) = \tilde{x}_i(y)$ and $\tilde{z}(q_{k+1}) = \tilde{z}(y)$ by the definition of the observation function η. □

Thus, given an observation sequence β, there is a one-to-one correspondence between the global preference p and the execution α. Fix an observation sequence, If the adversary has a high confident guess of the preference vector p, then the agents' whole trajectory corresponding to such a p is also of high validity. Otherwise, if the preference vector is protected, the evolution of the agents is hidden. So it suffices to consider privacy of the preference vector p.

4.1 Privacy and Cost in Distributed Control

For formulating privacy of a distributed control system, we first define comparable and adjacent preference vectors. For a pair of preference vectors p and p', the corresponding Markov chains $\mathcal{M}(p)$ and $\mathcal{M}(p')$ are *comparable* if the observable spaces are identical, that is, $Y_{\mathcal{M}(p)} = Y_{\mathcal{M}(p')}$.

Definition 1. *A pair of preference vectors p and p' in $(\mathcal{X}^\omega)^N$ are* adjacent up to time T, *written as T-adj(p, p') in short, if there exists a $k \in [N]$, such that for all $t \leq T$, such that (i) $|p_k(t) - p'_k(t)| \leq 1$, and (ii) for all $i \neq k$ $p_i(t) = p'_i(t)$.*

In other words, two preference vectors are T-adjacent if they differ only in the preferences of a single agent upto time T, and the difference in terms of the L^1 norm is at most unity at each time. We adapt the standard definition of differential privacy to this framework of control mechanisms, where the protection of the individual agent's preferences have to be balanced with the benefits of information sharing for control in a shared environment.

Definition 2. *The randomized control mechanism $\mathcal{M}$ is* ϵ-differentially private *upto time T, if for any two T-adjacent preference vectors p and p' and any set of finite observation sequences Obs, $\mathcal{M}(p)$ and $\mathcal{M}(p')$ are comparable and*

$$\mathbb{P}_{\mathcal{M}(p)}[\eta^{-1}_{\mathcal{M}(p)}(Obs)] \leq e^{\epsilon} \mathbb{P}_{\mathcal{M}(p')}[\eta^{-1}_{\mathcal{M}(p')}(Obs)]. \quad (6)$$

This definition of differential privacy is similar to the one appears in [4] with two technical differences. First, we restrict the preferences to be adjacent upto a time bound. Secondly, owing our choice of the definition of $\mathcal{M}(p)$ which allows all the components of $\mathcal{M}(p)$ to possibly depend on p, for privacy of individual agent's preferences with respect to observation sequences produced from to Markov chains,

it is required that the output alphabets of the corresponding chains are same. This requirement is incorporated by making the chains comparable.

Performance of a distributed control system is measured by a cost function. It is standard to consider the following quadratic cost in optimal control theory. Given an execution of length $T+1$, $\alpha = q_0, q_1, \ldots, q_T$, the cost of control for an individual agent i upto T time is the sum of squared distance between the agent's state and its preferred state. That is, $cost_{\mathcal{M}(p),i}(\alpha) \triangleq \sum_{t=1}^{T}(x_i(q_t) - p_i(t))^\top (x_i(q_t) - p_i(t))$. The summation starts with $t = 1$ because by definition $x_i(q_0) = p_i(0)$ and no cost is paid at time $t = 0$. The cost function of agent i is the expectation of the function $cost_{\mathcal{M}(p),i}(\alpha)$ over the space of executions of length T,

$$cost_{\mathcal{M}(p),i}(T) \triangleq \mathbb{E}[\sum_{t=1}^{T}(x_i(q_t) - p_i(t))^\top (x_i(q_t) - p_i(t))]. \tag{7}$$

We introduce an example of distributed control problem.

Example 1 This example captures the routing of N agents on a 2-D plane whose motion is affected by the center of gravity of all the agents. The agent i's state $x_i \in \mathcal{X} \subseteq \mathbb{R}^2$ has two components, which are the x and y coordinates of agent i. Each agent has a preference which is a path $p_i \in (\mathbb{R}^2)^\omega$. The individual agent's state at time t is affected by three factors: the previous state $x_i(t-1)$, the aggregate state– which is the center of the gravity of the herd–$z(t-1)$ and the individual's control input $u_i(t)$. The update law of the ith agent's state at time $t+1$ follows

$$x_i(t) = 1.5x_i(t-1) + cz(t-1) + u_i(t). \tag{8}$$

The aggregate state $z \in \mathcal{Z} \subset \mathbb{R}^2$ is the center of gravity of the herd.

$$z(t) = \frac{1}{N}\sum_{i\in[N]} x_i(t).$$

Designing a controller for the ith agent, which cancels out the influence of the aggregate state on individual agent and drives it towards the goal $p_i(t+1)$, needs the actual states of all other agents. With the precise information of others, agent i can achieves its desirable individual cost by using some optimal control technique. For example, the following controller may be used:

$$u_i(t) = -cz(t-1) - 1.3x_i(t-1) + 0.8p_i(t). \tag{9}$$

Combined Equations (9) and (8), we get the update rule for the whole close-loop system

$$x_i(t) = 0.2x_i(t-1) + 0.8p_i(t). \tag{10}$$

The current state $x_i(t)$ is a linear combination of the previous state $x_i(t-1)$ and the current preference $p_i(t)$. If the the sequence of preference p_i is fixed for a few round, the state x_i converges to it geometrically. Otherwise if the sequence of preference is changing, the state x_i keeps tracking it. The cost of individual agent is defined by the some squared distance between its state x_i and p_i, that is,

$$cost(p, T) = \sum_{t=1}^{T} |x_i(t) - p_i(t)|_2^2.$$

In Section 6, we introduce a mechanism that guarantees differential privacy of individuals for this example.

□

4.2 Cost of Privacy

We define the cost of a randomized control mechanism $\mathcal{M}$ as the difference in the cost of two nearly identical Markov chains with observations $\mathcal{M}(p)$ and $\mathcal{M}'(p)$, where $\mathcal{M}(p)$ is the differentially private Markov chain and $\mathcal{M}'(p)$ is identical except that its observation map discloses perfect information about the agents' states. Formally, given $\mathcal{M}(p)$ defined by the parameters f, g, h, and r the perfectly observable version $\mathcal{M}'(p)$ is defined by the parameters f, g, h and r' where

$$r'(x_i(q_t), \tilde{z}(q_t), t) \triangleq \delta_{x_i(q_t)},$$

where δ_a is the Dirac delta distribution at a.

Definition 3. *For any $\epsilon > 0$ and time bound $T \in \mathbb{N}$, and an ϵ-differentially private randomized control mechanism $\mathcal{M}$, the* Cost of Privacy (CoP) *upto time T, is defined by the supremum of the difference between any individual's cost in $\mathcal{M}(p)$ and the corresponding perfectly observable chain $\mathcal{M}'(p)$ over all preference vector p:*

$$CoP(\epsilon, \mathcal{M}, T) \triangleq \sup_{p,i} \left(cost_{\mathcal{M}'(p),i}(T) - cost_{\mathcal{M}(p),i}(T)\right).$$

We will discuss the cost of privacy of Example 1 in Section 6.2

5. LAPLACE OBSERVATIONS OF DIFFERENTIAL PRIVACY

In this section, we introduce a strategy for creating observation maps that guarantees differential privacy of the agents' preferences. For the remainder of this paper let n be the length of local state x_i. In this design, at time t, each agent report $\tilde{x}_i(t)$ by adding a noise $\omega_i(t)$ on its actual state $x_i(t)$, that is

$$r(x_i(q_t)) \triangleq x_i(q_t) + \omega_i(t), \tag{11}$$

where $\omega_i(t)$ is a vector consists of n independent random noises drawn from Laplace distribution $Lap(M_t)$.

Before propose an actual design of M_t, we first define sensitivity of the system. Fix an observation sequence $\beta \in Y^T$ up to time T and a preference vectors p. As we mentioned in Proposition 1, $\eta^{-1}_{\mathcal{M}(p)}(\beta)$ is a singleton set. Then, $x(\eta^{-1}_{\mathcal{M}(p)}(\beta)(t))$ is the global state at time t corresponding to the execution $\eta^{-1}_{\mathcal{M}(p)}(\beta)$. To quantify the maximal difference of the system's global states at time t resulted from a pair of adjacent preference vectors p and p', we introduce the sensitivity of the system.

Definition 4. *For a mechanism $\mathcal{M}$, we define the sensitivity of $\mathcal{M}$ at time $t \in \mathbb{N}$ as*

$$\Delta(t) \triangleq \sup_{\beta} \sup_{adj(p,p')} |x(\eta^{-1}_{\mathcal{M}(p)}(\beta(t))) - x(\eta^{-1}_{\mathcal{M}(p')}(\beta(t)))|,$$

where the norm used is L^1-norm.

We assume that $\Delta(t)$ is bounded for any $t \in \mathbb{N}$ throughout the paper.

Theorem 2. *At each time $t \in [T]$, if each agent adds a noise vector $\omega_i(t)$ which consists of n independent Laplace noise $Lap(M_t)$ such that $\sum_{t=0}^{T} \frac{\Delta_D(t)}{M_t} \leq \epsilon$, then the distributed control system is ϵ-differentially private.*

Proof. Fix a pair of T-adjacent preference vectors $p, p' \in \mathcal{X}^{TN}$. and a set of observation sequence $Obs \subseteq Y^T$. we will denote the sets of executions $\eta^{-1}_{\mathcal{M}(p)}(Obs)$ and $\eta^{-1}_{\mathcal{M}(p')}(Obs)$ by A and A' respectively. First, we define a correspondence B between the sets A and A'. For $\alpha \in A$ and $\alpha' \in A'$, $B(\alpha) = \alpha'$ if and only if they are the observation sequence up to time T. That is $\eta(\alpha(t)) = \eta(\alpha'(t))$ for all $t \in [T]$. From Proposition 1, for any observation sequence $\beta \in Obs$ there is an unique execution $\alpha \in \mathsf{Execs}_p$ that can produce the observation. Similarly, α' is also unique in $\mathsf{Exec}_{p'}$. So B is indeed a bijection. we relate the probability measures of the sets of executions A and A'.

$$\frac{\mathbb{P}_{\mathcal{M}(p)}[\eta^{-1}_{\mathcal{M}(p)}(Obs)]}{\mathbb{P}_{\mathcal{M}(p')}[\eta^{-1}_{\mathcal{M}(p')}(Obs)]} = \frac{\int_{\alpha \in A} \mathbb{P}_{\mathcal{M}(p)}[\alpha] d\mu}{\int_{\alpha' \in A'} \mathbb{P}_{\mathcal{M}(p')}[\alpha'] d\mu'}. \tag{12}$$

Changing the variable using the bijection B we have,

$$\begin{aligned} \int_{\alpha' \in A'} \mathbb{P}_{\mathcal{M}(p')}[\alpha'] d\mu' &= \int_{B(\alpha) \in A'} \mathbb{P}_{\mathcal{M}(p)}[B(\alpha)] d\mu \\ &= \int_{\alpha \in A} \mathbb{P}_{\mathcal{M}(p)}[B(\alpha)] d\mu \end{aligned} \tag{13}$$

From Equations (1)-(5) the definition of r,

$$\int_{\alpha \in A} \mathbb{P}_{\mathcal{M}(p)}[\alpha] d\mu = \int_{\alpha \in A} \mathbb{P}_{\mathcal{M}(p)}[\tilde{x}(\alpha) | x(\alpha)] d\mu$$

where $x(t)$ is the vector of N agent states at t along execution α. Each $x_i(t)$ is a vector of length n. We denote the k state component of $x_i(t)$ by $x_i^{(k)}(t)$. As $\tilde{x}(t)$ is obtained by adding $n \times N$ independent noise values to $x(t)$, from the distribution $Lap(M_t)$, it follows that the probability density of an execution is reduced to

$$\mathbb{P}_{\mathcal{M}(p)}[\tilde{x}(\alpha) | x(\alpha)] = \prod_{\substack{i \in [N], k \in [n] \\ t \in [T]}} p_L(\tilde{x}_i^{(k)}(\alpha(t)) - x_i^{(k)}(\alpha(t)) | M_t), \tag{14}$$

where $p_L(x|b)$ is the probability density function at x with parameter b. Then, we relate the distance at time t between the state of α and $B(\alpha)$ with the sensitivity $\Delta(t)$. Let $\beta = \eta(\alpha)$ be the observation sequence corresponding to α.

By the Definition 4, we have

$$|x(\alpha(t)) - x(\alpha'(t))| \leq \Delta(t).$$

The norm in above equation is L^1-norm. The global state $x(t)$ consists of N local state $x_i(t)$, each of which has n component. So $|x(\alpha(t)) - x(\alpha'(t))|$ lives in space $\mathbb{R}^{nN}$. By definition of L^1-norm:

$$\begin{aligned} & \sum_{i=1}^{N} \sum_{k=1}^{n} |x_i^{(k)}(\alpha(t)) - x_i^{(k)}(\alpha'(t))| \\ = \; & |x_i(\alpha(t)) - x_i(\alpha'(t))| \leq \Delta(t). \end{aligned}$$

Recall that by definition of B, the observations of α and $B(\alpha)$ match, that is $\tilde{x}(\alpha(t)) = \tilde{x}(B(\alpha)(t))$. From the property of Laplace distribution,

$$\begin{aligned} & \prod_{i \in [N] k \in [n]} \frac{p_L(\tilde{x}_i^{(k)}(\alpha(t)) - x_i^{(k)}(\alpha(t)) | M_t)}{p_L(\tilde{x}_i^{(k)}(B(\alpha)(t)) - x_i^{(k)}(B(\alpha)(t)) | M_t)} \\ \leq & \prod_{i \in [N], k \in [n]} e^{\frac{|\tilde{x}(\alpha(t)) - x(\alpha(t)) - \tilde{x}(B(\alpha)(t)) + x(B(\alpha)(t))|}{M_t}} \\ = & \prod_{i \in [N], k \in [n]} e^{\frac{|x(\alpha(t)) - x(B(\alpha)(t))|}{M_t}} \\ = & e^{\sum_{i \in [N], k \in [n]} \frac{|x(\alpha(t)) - x(B(\alpha)(t))|}{M_t}} \\ \leq & e^{\frac{\Delta(t)}{M_t}}. \end{aligned} \tag{15}$$

Combining Equation (12), (13), (14) and (15), we derive

$$\begin{aligned} & \frac{\mathbb{P}_{\mathcal{M}(p)}[\eta^{-1}_{\mathcal{M}(p)}(Obs)]}{\mathbb{P}_{\mathcal{M}(p')}[\eta^{-1}_{\mathcal{M}(p')}(Obs)]} \\ \leq & \frac{\int_{\alpha \in A} \mathbb{P}_{\mathcal{M}(p)}[\tilde{x}(\alpha) | x(\alpha)] d\mu}{\int_{\alpha \in A} \mathbb{P}_{\mathcal{M}(p)}[\tilde{x}(B(\alpha)) | x(B(\alpha))] d\mu} \\ \leq & \prod_{t \in [T]} e^{\frac{\Delta(t)}{M_t}} \leq e^{\sum_{t \in [T]} \frac{\Delta(t)}{M_t}} \end{aligned}$$

If M_t satisfy $\sum_{t=0}^{T} \frac{\Delta_D(t)}{M_t} \leq \epsilon$, then $\prod_{t \in [T]} e^{\frac{\Delta(t)}{M_t}} \leq e^{\epsilon}$. Thus the lemma follows. □

We can also derive the following corollary from Theorem 2.

Corollary 3. *At each time $t \in [T]$ if each agent add an vector of independent Laplace noise $Lap(M_t)$, where $M_t = \frac{\Delta_D(t) T}{\epsilon}$ to its actual state, then the distributed control system is ϵ-differentially private.*

In this mechanism, the noise added is proportional to the sensitivity of the system and the time bound of the system T. Roughly, an adversary can examine a number of T observations of an individual agent. The parameter of the Laplace noises added is proportional to the length of the observation and the sensitivity of the system.

6. DIFFERENTIALLY PRIVATE LINEAR DISTRIBUTED CONTROL

In this section, we will specialize the general framework of Section 4 to linear control systems. Linear models for the physical dynamics and linear controller functions are the predominant models studied in control theory literature. In this setup, the optimal controller design problem can be formulated and solved effectively using convex optimization. We assume that agent i's state (x_i), its observed state ($\tilde{x}_i$), its control input (u_i), the environment state (z), and the observed environment state ($\tilde{z}$) are all points in $\mathbb{R}^n$, for some natural number n. Agent i's preference is an infinite (possibly repeated) sequence of points in $\mathbb{R}^n$. Next, we define the remaining four parameters of the control system. The linear dynamics function for the i^{th} agent is:

$$f(x_i, z, t) \triangleq Ax_i + cz + u_i,$$

where $A \in \mathbb{R}^{n \times n}$ is the dynamics matrix and $c \in \mathbb{R}$ is a coupling constant. The linear aggregation function h computes

the average of the agent states, which is defined as

$$h(x) \triangleq \frac{1}{N} \sum_{i \in [N]} x_i.$$

For this type of dynamics, a linear feedback controller suffices to drive the agent to any fixed preference point. We choose a general linear feedback control function of the form:

$$g(x_i, z, t) \triangleq (K - A)x_i + (I - K)p_i(t) - c\tilde{z},$$

where $K \in \mathbb{R}^{n \times n}$ is a *stable matrix* and I is the identity matrix. Finally, the form of the observation map is

$$r(x_i, t) \triangleq x_i + \omega_i(t),$$

where $\omega_i(t)$ is drawn from a time-dependent probability distribution to be defined below.

As in the general case (Section 4), given a preference vector p, the above parameters define the Markov chain $\mathcal{M}(p)$ which captures the evolution of the system. The system of equations defining the transitions of this Markov chain, corresponding to Equations (1)-(5), can be written as follows: At time $t \in \mathbb{N}$

$$u_i(t) = (K - A)x_i(t-1) + (I - K)p_i(t) - c\tilde{z}(t-1) \quad (16)$$

$$x_i(t) = Ax_i(t-1) + cz(t-1) + u_i(t) \quad (17)$$

$$\tilde{x}_i(t) = x_i(t) + \omega_i(t) \quad (18)$$

$$z(t) = \frac{1}{N} \sum_{i \in [N]} x_i(t) \quad (19)$$

$$\tilde{z}(t) = \frac{1}{N} \sum_{i \in [N]} \tilde{x}_i(t). \quad (20)$$

Combining the above equations, the closed-loop dynamics of agent i is:

$$x_i(t) = Kx_i(t-1) + (I - K)p_i(t) - \frac{c}{N} \sum_{i \in [N]} \omega_i(t-1). \quad (21)$$

Agent i's state at time t can be written as a function of its preference sequence $\{p_i(s)\}_{s \le t}$ and the sequence $\{\omega_i(s) : i \in [N], s \le t\}$ of noise vectors added in all previous rounds. By iteratively applying Equation (21), we obtain:

$$\begin{aligned} x_i(t) =& K^t p_i(0) + \sum_{s=1}^{t} K^{t-s}(I - K)p_i(s) \\ &- \frac{c}{N} \sum_{s=0}^{t-1} K^{t-s-1} \sum_{i \in [N]} \omega_i(s). \end{aligned} \quad (22)$$

Remark 1. By taking expectation on both side of Equation (21), we can write $x_(t) - p_(t) = K(x(t-1) - p(t))$. Given a stable the matrix K, for agent i, after update at time t, the new state gets closer to the preference $p_i(t)$. The more stable K is, the better tracking $x(t)$ performs towards $p(t)$.

For representing the dynamics of the complete system with N agents, we define two $nN \times nN$ matrices

$$\mathbf{K} \triangleq \begin{bmatrix} K & & \\ & \ddots & \\ & & K \end{bmatrix} \text{ and } \mathbf{C} \triangleq \frac{c}{N} \begin{bmatrix} I & \dots & I \\ \vdots & \ddots & \vdots \\ I & \dots & I \end{bmatrix},$$

where $\mathbf{K}$ is a block diagonal matrix with K matrices as its diagonal blocks and $\mathbf{C}$ is a block matrix with all the blocks set to $\frac{c}{N}$ times the identity matrix I. Combining the Equation (21) for all the N agents we obtain:

$$\begin{aligned} x(t) &= \mathbf{K}x(t-1) + (I - \mathbf{K})p(t) + \mathbf{C}x(t-1) - \tilde{z}(t-1) \\ &= (\mathbf{K} + \mathbf{C})x(t-1) + (I - \mathbf{K})p(t) - \tilde{z}(t-1). \end{aligned} \quad (23)$$

Given a preference vector p and an observation sequence β, by Proposition 1, we know that there is a unique execution $\eta^{-1}_{\mathcal{M}(p)}(\beta)$. The vector of agent states at time $t \ge 0$, along this execution is $x(\eta^{-1}_{\mathcal{M}(p)}(\beta)(t))$. Iteratively applying Equation (23) we obtain:

$$\begin{aligned} x(\eta^{-1}_{\mathcal{M}(p)}(\beta(t))) = (\mathbf{K} + \mathbf{C})^t p(0) - \sum_{s=0}^{t-1} (\mathbf{K} + \mathbf{C})^{t-s} \tilde{z}(\beta(t)) \\ + \sum_{s=1}^{t} (\mathbf{K} + \mathbf{C})^{t-s} (I - \mathbf{K}) p(s). \end{aligned}$$

6.1 Sensitivity of Linear Distributed Control

In this section, we state Theorem 4 which establishes bound on the sensitivity $\Delta(t)$. We refers the reader to the full version of the paper for the proof of this nontrivial result [9]. For proving this theorem, we fix two Markov chains of the system (Equations (16) -(20)) with adjacent preference vectors p and p' and compute the difference between two chains. Recall that p and p' are identical except the preference of one agent (i). Then, the difference between the two Markov chains has two components: (1) the change in agent i's state, and (2) the sum of changes in other agents' state. The sensitivity is then computed as a bound of the sum of above two components. With this bound on sensitivity, we introduce a Laplace mechanism defining the observation map (r) in Corollary 5 and then show that the mechanism achieves differential privacy of the linear distributed control system.

Theorem 4. *For the linear distributed control system, for all $t \in \mathbb{N}$ the sensitivity $\Delta(t)$ is upper bounded by $\kappa(t)$, where*

$$\kappa(t) \triangleq |G^t - K^t| + |K^t||H| \sum_{s=0}^{t-1} (|G^t - K^t| + |K^s|),$$

with $G \triangleq cI + K$ and $H \triangleq I - K$.

Remark 2. The upper bound on the sensitivity at time t, $\kappa(t)$ has two components:

(a) $|K^t| + |H| \sum_{s=1}^{t} |K^s|$ overapproximates the maximum change in agent i's state (x_i) if its own preference changes by at most unity at each time upto t, and

(b) $|G^s + K^s| + |H| \sum_{s=0}^{t-1} |G^s + K^s|$ overapproximates the sum of the changes in other agents' state given agent i's preference changes by at most unity upto t.

Remark 3. $\kappa(t)$ is independent to the number of agents (N). It only depends on matrix K, the coupling constant c and time t. K is specified by the individual's control function (Equation (16)), which assumes to be stable. The more stable matrix K is, the faster $|K^t|$ decays to 0. The coupling constant c quantifies the influence of the aggregate on each

individual agent. The matrix $G = cI + K$ captures the combined dynamics under the influence of the environment and the dynamics of the individual agents. The weaker physical coupling is, the smaller $|G^t|$ is. Therefore, we conclude that, as the individual agent dynamics becomes more stable or the physical coupling between agents becomes weaker, the sensitivity of the system decreases.

Remark 4. The dependence of $\kappa(t)$ on time t changes based on the stability of the K and G matrices. If G is stable, $\kappa(t)$ converges to a constant as $t \to \infty$. Otherwise if G is unstable, $\kappa(t)$ grows exponentially with t.

Theorems 2 and 4 immediately suggest an observation map (r) which guarantees differential privacy of the distributed linear control system.

Corollary 5. *For any time bound T and privacy parameter $\epsilon > 0$, for $M_t \triangleq \frac{T\kappa(t)}{\epsilon}$ and $\omega_i(t)$ chosen as noise vector of length n drawn independently from the distribution $Lap(M_t)$, the resulting observation map makes the linear distributed control system ϵ-differentially private up to time T.*

Example 2 Now we can apply the strategy explained above to Example 1, where $K = \frac{1}{5}I$ is a 2 by 2 matrix. $G = (c+\frac{1}{5})I$ in this case. By Theorem 4, the sensitivity is

$$\Delta(t) \leq \kappa(t) = \frac{4+20c}{20-25c} + \frac{16-45c}{20-25c}\left(c+\frac{1}{5}\right)^t$$

As stated in Remark 3, the sensitivity is independent of N. If G is stable, that is $|c+\frac{1}{5}| \leq 1$, the sensitivity $\Delta(t)$ is bounded and converges to a constant as $t \to \infty$. Otherwise, if $|c+\frac{1}{5}| > 1$, $\kappa(t)$ diverges. We choose the noise to be $M_t = \frac{\kappa(t)T}{\epsilon}$. By Corollary 5, the system guarantees ϵ-differential privacy upto time T . □

6.2 Cost of Privacy in Linear Distributed Control

The observation map of Corollary 5 adds independently drawn Laplace noise to the state of agent i observation at time t from the distribution $Lap(M_t)$. The noise parameter M_t depends on the individual's dynamics rather than the number of agents. In this section, we discuss the cost of privacy for this mechanism (see, Definition 3) compared to a perfectly observable system using the same controller.

Theorem 6. *The cost of privacy of the ϵ-differentially private mechanism $\mathcal{M}$ of Corollary 5 is inversely proportional to the number of agents N and the squared privacy parameter ϵ^2. In addition, if matrix G is stable, it is proportional to T^3. Otherwise if G is unstable, the cost of privacy grows exponentially with T.*

Proof. Given the ϵ-differentially private mechanism $\mathcal{M}$, the perfectly observable system $\mathcal{M}'$ is obtained by setting the noise values to be 0. We denote by $\bar{x}_i(t)$ the state of agent i for $\mathcal{M}'$ at time t. From Equation (22), by fixing $\omega_i(t) \equiv 0$, we get

$$\bar{x}_i(t) = K^t p_i(0) + \sum_{s=1}^{t} K^{t-s}(I-K)p_i(s).$$

We define a $n \times nN$ matrix $\mathbf{B} \triangleq \frac{c}{N}[I, \ldots, I]$. Let $x_i(t)$ be agent i's state corresponding to some execution of $\mathcal{M}(p)$. Again from Equation (22), the state of an individual agent i is

$$x_i(t) = \bar{x}_i(t) - \sum_{s=0}^{t-1} K^{t-s-1}\mathbf{B}\omega(s).$$

The cost of the mechanism $\mathcal{M}$ can be written as

$$\begin{aligned} & cost_{\mathcal{M},i}(T) = \mathbb{E}\left[\sum_{t=1}^{T} |x_i(t) - p_i(t)|_2^2\right] \\ = \;& \mathbb{E}\left[\sum_{t=1}^{T} |\bar{x}_i(t) - \sum_{s=0}^{t-1} K^{t-s-1}\mathbf{B}\omega(s) - p_i(t)|_2^2\right] \\ = \;& \sum_{t=1}^{T} \mathbb{E}[|\bar{x}_i(t) - p_i(t)|_2^2 + |\sum_{s=0}^{t-1} K^{t-s-1}\mathbf{B}\omega(s)|_2^2 \\ & -2(\bar{x}_i(t) - p_i(t))^\top \sum_{s=0}^{t-1} K^{t-s-1}\mathbf{B}\omega(s)] \end{aligned}$$

The first term on the right hand side is the cost of the system with perfect observations, that is, $cost_{\mathcal{M}',i}(T)$. The last term on the right hand side is the expectation of a linear combination of zero-mean noise terms, and therefore, equals 0. By Definition 3,

$$\begin{aligned} & CoP(\epsilon, \mathcal{M}, T) \\ = \;& \sup_{p,i}[cost_{\mathcal{M}(p),i}(T) - cost_{\mathcal{M}'(p),i}(T)] \\ = \;& \textstyle\sum_{t=1}^{T} \mathbb{E}\left[|\sum_{s=0}^{t-1} K^{t-s-1}\mathbf{B}\omega(s)|_2^2\right] \end{aligned} \tag{24}$$

In our Laplace mechanism, for different time steps s, τ, $\omega(s)$ and $\omega(\tau)$ are independent. Thus, $\mathbb{E}[\omega(s)^\top\omega(\tau)] = \mathbb{E}[\omega(s)]^\top\mathbb{E}[\omega(\tau)] = 0$. Then, the right hand side of Equation (24) reduces to

$$\sum_{t=1}^{T} \mathbb{E}\left[\sum_{s=0}^{t-1} \omega(s)^\top \mathbf{B}^\top (K^{t-s-1})^\top K^{t-s-1}\mathbf{B}\omega(s)\right].$$

Recall that each $\omega(s)$ consists of a noise vector $\omega_i(s)$ for each agent $i \in [N]$, and each of these vectors have n independent and identically distributed noise values drawn from $Lap(M_s)$. Each pair of vectors in $\omega(s)$ are independent. Denote $\omega^{(k)}(s)$, $k \in [nN]$, be the k^{th} element of the vector $\omega(s)$.

It follows that (a) for different indices $k, j \in [nN]$, $\mathbb{E}\left[\omega^{(k)}(s)\omega^{(j)}(s)\right]$ 0, and (b) for any $k \in [nN]$, $\mathbb{E}\left[\omega^{(k)}(s)\omega^{(k)}(s)\right] = Var\left[\omega^{(k)}(s)\right] = 2M_s^2$. Thus,the above expression is reduced to

$$\sum_{t=1}^{T}\sum_{s=0}^{t-1} M_s^2 \mathrm{Tr}(\mathbf{B}^\top (K^{t-s-1})^\top K^{t-s-1}\mathbf{B}), \tag{25}$$

where $\mathrm{Tr}(A)$ stands for the trace of matrix A. Recall that $\mathbf{B} \triangleq \frac{c}{N}[I, \ldots, I]$. It follows that

$$\begin{aligned} & \mathrm{Tr}(\mathbf{B}^\top (K^{t-s-1})^\top K^{t-s-1}\mathbf{B}) \\ =& \frac{c^2}{N}\mathrm{Tr}((K^{t-s-1})^\top K^{t-s-1}) = \frac{c^2}{N}|K^{t-s-1}|_2^2. \end{aligned}$$

Substituting the above equation into Equation (25) yields

$$CoP(\epsilon, \mathcal{M}, T) = \frac{c^2}{N}\sum_{t=1}^{T}\sum_{s=0}^{t-1} M_s^2 |K^{t-s-1}|_2^2$$

By interchanging the order of summation we get

$$\begin{aligned} CoP(\epsilon, \mathcal{M}, T) =& \frac{c^2}{N} \sum_{s=0}^{T-1} \sum_{t=s+1}^{T} M_s^2 |K^{t-s-1}|_2^2 \\ =& \frac{c^2}{N} \sum_{s=0}^{T-1} M_s^2 \sum_{t=0}^{T-s-1} |K^t|_2^2. \end{aligned} \tag{26}$$

Recall that in Corollary 5, $M_s = \frac{T\kappa(s)}{\epsilon}$. Combining this with Equation (26), we have

$$CoP(\epsilon, \mathcal{M}, T) = \frac{c^2 T^2}{N\epsilon^2} \sum_{s=0}^{T-1} \kappa(s)^2 \sum_{t=0}^{T-s-1} |K^t|_2^2.$$

From the above expression it is clear $CoP(\epsilon, \mathcal{M}, T)$ is inversely proportional to N and ϵ^2. As the matrix K is stable, $\sum_{t=0}^{T-s-1} |K^t|_2^2$ converges to some constant as $T \to \infty$. By Remark 4, if G is stable then $\kappa(s)$ converges to some constant as $s \to \infty$, $\sum_{s=0}^{T-1} \kappa(s)^2$ grows linearly with T and we have $CoP(\epsilon, \mathcal{M}, T) \sim O(T^3)$. Otherwise if G is unstable, $\kappa(s)$ grows exponentially with s and $CoP(\epsilon, \mathcal{M}, T)$ grows exponentially with T. □

Example 3 Continuing with the system described in Example 1, we now establish the cost of privacy associated with the communication strategy of Equation (26). In this example, $K = 0.2I$. We choose the coupling parameter c to be 0.4. Then, the close-loop system is stable. Therefore, the sensitivity is bounded by $\kappa(t) = 1.2 - 0.2 \times 0.6^t$. The cost of privacy of the system with N agents at time T follows $\frac{0.24T^3}{N\epsilon^2} + O(\frac{T^2}{N\epsilon^2})$.

We have explored an alternative communication strategy which also guarantees ϵ-differential privacy while minimizing the cost of privacy (see the Appendix of [9]). By this strategy, the cost of privacy of the system at time T follows $\frac{0.12T^3}{N\epsilon^2} + O(\frac{T^2}{N\epsilon^2})$. □

Example 4 We conclude with a simulation-based analysis of the traffic control Example 1. Consider a linear distributed control system in which each agent is a point on the plane moving towards a randomly chosen destination with dynamics described in Example 3 and control strategies given in Example 3. The cost of each agent is defined by the distance between its position to its destination. The coupling between agents is the repulsive force in the direction of the center of gravity (CM) of the population. Thus, if the control of an individual fights the force too strongly without the knowledge of the CM then a higher cost is incurred. We numerically simulated the system with different levels of privacy and different distributions of destinations and make the following observations.

Fig 2 shows the relative costs of control with (blue) no communication and (green) private communication, with respect to cost of control with complete (or brodcast) communication. First of all, if both the initial positions and the destinations are chosen with 0 mean, then the CM of the population hovers around the origin and in that case, the contribution of the coupling is small. As a result, there is not much to be gained through communication and we see (Figure 2) that the cost of the system with privacy is comparable to the cost of the system with no communication. When the destination comes from some distributions slightly biased from 0, we start to see that the cost of control with private communication starts to become smaller compared to those of systems with no communications.

Figure 3 shows that for the same distribution of initial positions and destinations the cost of privacy changes as predicted by Theorem 6. First of all, higher level of privacy comes with higher cost (Figure 3a). As ϵ changes from 0.2 to 2, the CoP changes from 10 to 0.1. Secondly, larger number of agents (N) gives lower cost of privacy (Figure 3b). As N changes from 10 to 100, the CoP decreases from 4 to 0.4. And finally a longer time horizon (T) translates to higher costs (Figure 3c). The simulation results suggest that the cost of privacy roughly has the order of $O(\frac{T^3}{N\epsilon^2})$. □

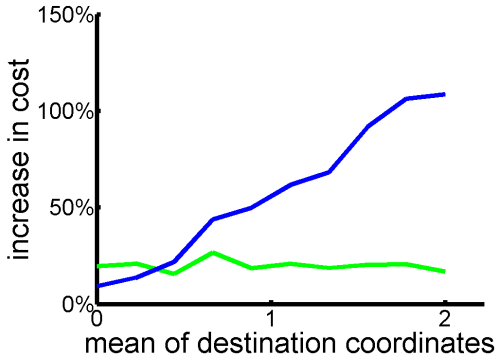

Figure 2: Increase in cost with biased sampled destinations. The blue and green lines capture the relative cost of control with no communication and private communication with respect to the cost of control with broadcast preferences respectively.

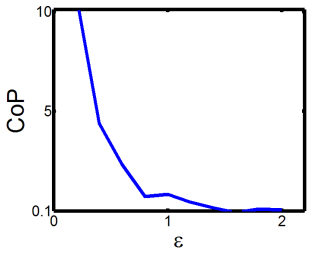

(a) CoP v.s. privacy level ϵ.

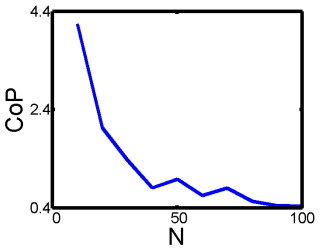

(b) CoP v.s. number of agents N.

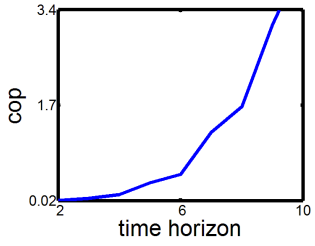

(c) CoP v.s. time horizon T.

Figure 3: Cost of Privacy for different privacy level, number of agents and time horizon.

7. CONCLUSIONS

We presented a general framework for studying cost of differential privacy for distributed control systems. We proposed a communication strategy by which individual agents

can share noisy information about their state which preserves ϵ-differential privacy while aiding the estimation of the aggregate environment and therefore improving control performance. Specializing to linear systems with quadratic costs, we showed that the sensitivity of η^{-1} and therefore the standard deviation of the required noise is independent of the number of participating agents. The sensitivity also decreases with the stability of the dynamics and with the weakening the environment's influence on an individual. For stable controllers, for preserving privacy over indefinite time horizons, the variance of the noise to be added is also independent of time. For unstable dynamics, on the other hand, the sensitivity can grow exponentially with time. The cost of ϵ-differential privacy for the proposed communication strategy up to time T for a system with N agents is at most $O(\frac{T^3}{N\epsilon^2})$ for stable systems. This suggests that the proposed communication strategy is best suited for distributed control systems with many short-lived participants.

The proposed framework should enable us to study more sophisticated communication strategies that incur smaller costs for more persistent agents. Another direction for future research will be to establish lower bounds on the best cost of privacy that can be achieved through any communication strategy, not just the form proposed here.

8. REFERENCES

[1] C. Dwork. Differential privacy. In *AUTOMATA, LANGUAGES AND PROGRAMMING*, volume 4052 of *Lecture Notes in Computer Science*, 2006.

[2] C. Dwork. Differential privacy: a survey of results. In *Proceedings of the 5th international conference on Theory and applications of models of computation*, TAMC'08, pages 1–19, Berlin, Heidelberg, 2008. Springer-Verlag.

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

[4] C. Dwork, M. Naor, G. Rothblum, and T. Pitassi. Differential privacy under continual observation. In *Proceedings of the 42nd ACM symposium on Theory of computing*, 2010.

[5] Q. Geng and P. Viswanath. Optimal noise-adding mechanism in differential privacy. *CoRR*, abs/1212.1186, 2012.

[6] M. Hardt and K. Talwar. On the geometry of differential privacy. In *Proceedings of the 42nd ACM symposium on Theory of computing*, STOC '10, pages 705–714, New York, NY, USA, 2010. ACM.

[7] J. Herrera, D. Work, R. Herring, X. Ban, Q. Jacobson, and A. Bayen. Evaluation of traffic data obtained via GPS-enabled mobile phones: The Mobile Century field experiment. *Transportation Research Part C*, 18(4):568–583, August 2010.

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

[9] Z. Huang, Y. Wang, S. Mitra, and G. Dullerud. On the cost of differential privacy in distributed control systems, 2013. Full version: http://users.crhc.illinois.edu/mitras/research/2013/cop_full.pdf.

[10] J. Le Ny and G. J. Pappas. Differentially Private Filtering. *ArXiv e-prints*, July 2012.

[11] J. Le Ny and G. J. Pappas. Differentially Private Kalman Filtering. *ArXiv e-prints*, July 2012.

[12] C. Li, M. Hay, V. Rastogi, G. Miklau, and A. McGregor. 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, pages 123–134, New York, NY, USA, 2010. ACM.

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

[14] S. Mitra. *A Verification Framework for Hybrid Systems*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA 02139, September 2007.

[15] J. Reed and B. C. Pierce. Distance makes the types grow stronger: a calculus for differential privacy. In *Proceedings of the 15th ACM SIGPLAN international conference on Functional programming*, ICFP '10, pages 157–168, New York, NY, USA, 2010. ACM.