2020 59th IEEE Conference on Decision and Control (CDC)
Jeju Island, Republic of Korea, December 14-18, 2020

Differentially Private Formation Control

Calvin Hawkins and Matthew Hale*

Abstract— **As multi-agent systems proliferate, there is increasing demand for coordination protocols that protect agents' sensitive information while allowing them to collaborate. To help address this need, this paper presents a differentially private formation control framework. Agents' state trajectories are protected using differential privacy, which is a statistical notion of privacy that protects data by adding noise to it. We provide a private formation control implementation and analyze the impact of privacy upon the system. Specifically, we quantify tradeoffs between privacy level, system performance, and connectedness of the network's communication topology. These tradeoffs are used to develop guidelines for calibrating privacy in terms of control theoretic quantities, such as steady-state error, without requiring in-depth knowledge of differential privacy. Additional guidelines are also developed for treating privacy levels and network topologies as design parameters to tune the network's performance. Simulation results illustrate these tradeoffs and show that strict privacy is inherently compatible with strong system performance.**

I. Introduction

Multi-agent systems, such as robotic swarms and social networks, require agents to share information to collaborate. In some cases, the information shared between agents may be sensitive. For example, self-driving cars share location data to be routed to a destination. Geo-location data and other data streams can be quite revealing about users and sensitive data should be protected. However, this data must still be useful for multi-agent coordination. Thus, privacy in multi-agent control must simultaneously protect agents' sensitive data while guaranteeing that privatized data enables the network to achieve a common task.

This type of privacy has recently been achieved using differential privacy. Differential privacy stems from the computer science literature, where it was originally used to protect sensitive data when databases are queried [1], [2]. Differential privacy is appealing because it is immune to post-processing and robust to side information [1]. These properties mean that privacy guarantees are not compromised by performing operations on differentially private data, and that they are not weakened by much by an adversary with additional information about data-producing agents [3].

Recently, differential privacy has been applied to dynamic systems [4]–[12]. One form of differential privacy in dynamic systems protects sensitive trajectory-valued data, and this is the notion of differential privacy used in this paper. Privacy of this form ensures that an adversary is unlikely to learn much about the state trajectory of a system by observing its outputs. In multi-agent control, this lets an agent share its outputs with other agents while protecting its state trajectory from those agents and eavesdroppers [4]–[7].

In this paper, we develop a framework for private multi-agent formation control using differential privacy. Formation control is a well-studied network control problem and can be robots physically assembling into geometric shapes or non-physical agents maintaining relative state offsets. For differential privacy, agents add privacy noise to their states before sharing them with other agents. The other agents use privatized states in their update laws, and then this process repeats at every time step. The private formation control protocol can be implemented in a completely distributed manner, and, contrary to some other privacy approaches, it does not require a central coordinator.

Beyond the privacy implementation, we develop guidelines for calibrating privacy in formation control. Specifically, we bound the quality of formation, or performance of the system, in terms of agents' privacy parameters and connectedness of the network. We develop guidelines by using these bounds to trade off degraded performance for stricter privacy requirements and a less connected communication topology. This ultimately allows us to formulate privacy guidelines based on control-theoretic properties without requiring users to have an in-depth understanding of differential privacy. Furthermore, we develop necessary and sufficient conditions for when private formation control networks achieve a desired performance level.

The rest of the paper is organized as follows. Section II gives graph theory and differential privacy background. Section III states the differentially private formation control problem and Section IV solves it. Section V provides guidelines for calibrating privacy based on performance requirements for specific communication topologies. In Section VI, we analyze the sensitivity of system performance to changes in privacy and communication topology. Next, Section VII provides simulations, and Section VIII concludes the paper.

II. Background and Preliminaries

In this section we briefly review the required background on graph theory and differential privacy.

A. Graph Theory Background

A graph $\mathcal{G} = (V, E)$ is defined over a set of nodes V and edges are contained in the set E. For N nodes, V is indexed over $\{1, ..., N\}$. The edge set of $\mathcal{G}$ is a subset $E \subseteq V \times V$, where the pair $(i, j) \in E$ if nodes i and j share a connection and $(i, j) \notin E$ if they do not. This paper considers

This work was supported by the AFOSR Center of Excellence on Assured Autonomy in Contested Environments and by NSF CAREER grant #1943275.

*The authors are with the Department of Mechanical and Aerospace Engineering, Herbert Wertheim College of Engineering, University of Florida. Emails: {calvin.hawkins,matthewhale}@ufl.edu.

undirected, weighted, simple graphs. Undirectedness means that an edge $(i,j) \in E$ is not distinguished from $(j,i) \in E$. Simplicity means that $(i,i) \notin E$ for all $i \in V$. Weightedness means that the edge $(i,j) \in E$ has a weight $w_{ij} = w_{ji} > 0$. Of particular interest are connected graphs.

Definition 1 (Connected Graph): A graph $\mathcal{G}$ is connected if, for all $i,j \in \{1,...,N\}$, $i \neq j$, there is a sequence of edges one can traverse from node i to node j. $\triangle$

This paper uses the weighted graph Laplacian, which is defined with weighted adjacency and weighted degree matrices. The weighted adjacency matrix $A(\mathcal{G}) \in \mathbb{R}^{N\times N}$ of $\mathcal{G}$ is defined element-wise as

$$A(\mathcal{G})_{ij} = \begin{cases} w_{ij} & (i,j) \in E \\ 0 & \text{otherwise} \end{cases}.$$

Because we only consider undirected graphs, $A(\mathcal{G})$ is symmetric. The weighted degree of node $i \in V$ is defined as $d_i = \sum_{j|(i,j)\in E} w_{ij}$. The maximum degree is $d_{max} = \max_i d_i$. The degree matrix $D(\mathcal{G}) \in \mathbb{R}^{N\times N}$ is the diagonal matrix $D(\mathcal{G}) = \text{diag}(d_1,...,d_N)$. The weighted Laplacian of $\mathcal{G}$ is then defined as $L(\mathcal{G}) = D(\mathcal{G}) - A(\mathcal{G})$.

Let $\lambda_k(\cdot)$ be the k^{th} smallest eigenvalue of a matrix. By definition, $\lambda_1(L(\mathcal{G})) = 0$ for all graph Laplacians and

$$0 = \lambda_1(L(\mathcal{G})) \leq \lambda_2(L(\mathcal{G})) \leq \cdots \leq \lambda_N(L(\mathcal{G})).$$

The value of $\lambda_2(\mathcal{G}))$ plays a key role in this paper.

Definition 2 (Algebraic Connectivity [13]): The algebraic connectivity of a graph $\mathcal{G}$ is the second smallest eigenvalue of its Laplacian and $\mathcal{G}$ is connected if and only if $\lambda_2(L(\mathcal{G})) > 0$. $\triangle$

Agent i's neighborhood set $N(i)$ is the set of all agents agent i can communicate with, $N(i) = \{j \mid (i,j) \in E\}$.

B. Differential Privacy Background

This section provides a brief description of the differential privacy background needed for the remainder of the paper. More complete expositions can be found in [4], [14]. Overall, the goal of differential privacy is to make similar pieces of data appear approximately indistinguishable from one another. Differential privacy is appealing because its privacy guarantees are immune to post-processing [14]. For example, private data can be filtered without threatening its privacy guarantees [4], [15]. More generally, arbitrary post-hoc computations on private data do not harm differential privacy. In addition, after differential privacy is implemented, an adversary with complete knowledge of the mechanism used to implement privacy has no advantage over another adversary without mechanism knowledge [1], [2].

In this paper we use differential privacy to privatize state trajectories of mobile autonomous agents. We consider vector-valued trajectories of the form $Z = (Z(1), Z(2), ..., Z(k), ...)$, where $Z(k) \in \mathbb{R}^d$ for all k. The ℓ_p norm of Z is defined as $\|Z\|_{\ell_p} = \left(\sum_{k=1}^{\infty} \|Z(k)\|_p^p\right)^{\frac{1}{p}}$, where $\|\cdot\|_p$ is the ordinary p-norm on $\mathbb{R}^d$. Define the set

$$\ell_p^d := \{Z \mid Z(k) \in \mathbb{R}^d, \|Z\|_{\ell_p} < \infty\}.$$

The set ℓ_p^d only contains trajectories that converge to the origin. However, we want to privatize arbitrary trajectories, including those that do not converge at all. To do so, we consider a larger set of trajectories. Let the truncation operator P_T be defined as

$$P_T[y] = \begin{cases} y(k) & k \leq T \\ 0 & k > T \end{cases}.$$

Then we define the set

$$\tilde{\ell}_p^d = \{Z \mid Z(k) \in \mathbb{R}^d, P_T[Z] \in \ell_p^d \text{ for all } T \in \mathbb{N}\},$$

and we will privatize state trajectories in this set.

Consider a network of N agents, where agent i's state trajectory is denoted by y_i. The k^{th} element of agent i's trajectory is $y_i(k) \in \mathbb{R}^n$ for $n \in \mathbb{N}$. Agent i's state trajectory belongs to $\tilde{\ell}_2^n$.

Differential privacy is defined with respect to an adjacency relation. We provide privacy to single agents' state trajectories (rather than collections of trajectories as in some other works), and our choice of adjacency relation is defined for single agents. In the case of dynamic systems, the adjacency relation gives a notion of how similar trajectories are and specifies which trajectories must be made approximately indistinguishable from each other.

Definition 3 (Adjacency [5]): Fix an adjacency parameter $b_i > 0$ for agent i. $\text{Adj}_{b_i} : \tilde{\ell}_2^n \times \tilde{\ell}_2^n \to \{0,1\}$ is defined as

$$\text{Adj}_{b_i}(v_i, w_i) = \begin{cases} 1 & \|v_i - w_i\|_{\ell_2} \leq b_i \\ 0 & \text{otherwise.} \end{cases} \quad \triangle$$

In words, two state trajectories of agent i are adjacent if and only if the ℓ_2-norm of their difference is upper bounded by b_i. This means that every state trajectory within distance b_i from agent i's state trajectory must be made approximately indistinguishable from it to enforce differential privacy.

To calibrate differential privacy's protections, agent i selects privacy parameters ϵ_i and δ_i. These parameters determine the level of privacy afforded to x_i. Typically, $\epsilon_i \in [0.1, \ln 3]$ and $\delta_i \leq 0.01$ for all i [5]. The value of δ_i can be regarded as the probability that differential privacy fails for agent i, while ϵ_i can be regarded as the information leakage about agent i.

The implementation of differential privacy in this work provides differential privacy for each agent individually. This will be accomplished by adding noise to sensitive data directly, an approach called "input perturbation" privacy in the literature [16]. The noise is added by a privacy mechanism, which is a randomized map. We now provide a formal definition of differential privacy, which states the guarantees a mechanism must provide. First, fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. We are considering outputs in $\tilde{\ell}_2^n$ and use a σ-algebra over $\tilde{\ell}_2^n$, denoted Σ_2^n [17].

Definition 4 (Differential Privacy): Let $\epsilon_i > 0$ and $\delta_i \in [0, \frac{1}{2})$ be given. A mechanism $M : \tilde{\ell}_2^n \times \Omega \to \tilde{\ell}_2^n$ is (ϵ_i, δ_i)-differentially private if, for all adjacent $y_i, y_i' \in \tilde{\ell}_2^n$, we have

$$\mathbb{P}[M(y_i) \in S] \leq e^{\epsilon_i}\mathbb{P}[M(y_i') \in S] + \delta_i \text{ for all } S \in \Sigma_2^n. \ \triangle$$

The Gaussian mechanism will be used to implement differential privacy. The Gaussian mechanism adds zero-mean i.i.d. noise drawn from a Gaussian distribution pointwise in time. Stating the required distribution uses the Q-function, defined as $Q(y) = \frac{1}{\sqrt{2\pi}} \int_y^\infty e^{-\frac{z^2}{2}} dz$.

Lemma 1 (Gaussian Mechanism [4]): Let $b_i > 0$, $\epsilon_i > 0$, and $\delta_i \in (0, \frac{1}{2})$ be given, and fix the adjacency relation Adj_{b_i}. Let $y_i \in \tilde{\ell}_2^n$. The Gaussian mechanism for (ϵ_i,δ_i)-differential privacy takes the form $\tilde{y}_i(k) = y_i(k) + w_i(k)$, where w_i is a stochastic process with $w_i(k) \sim \mathcal{N}(0, \sigma_i^2 I_n)$ and $\sigma_i \geq \frac{b_i}{2\epsilon_i}(K_{\delta_i} + \sqrt{K_{\delta_i}^2 + 2\epsilon_i})$ where $K_{\delta_i} = Q^{-1}(\delta_i)$. This mechanism provides (ϵ_i,δ_i)-differential privacy to y_i. ■

For convenience, let $\kappa(\delta_i, \epsilon_i) = \frac{1}{2\epsilon_i}(K_{\delta_i} + \sqrt{K_{\delta_i}^2 + 2\epsilon_i})$.

III. Problem Formulation

In this section we state and analyze the differentially private formation control problem.

Problem 1: Consider a network of N agents with communication topology modeled by the undirected, simple, connected, and weighted graph $\mathcal{G}$. Let $y_i(k)$ be agent i's state at time k, $N(i)$ be agent i's neighborhood set, $\gamma > 0$, and w_{ij} be a positive weight on the edge $(i, j) \in E$. We define $\Delta_{ij} \in \mathbb{R}^n$ for all $(i, j) \in E$ as the desired relative distance between agents i and j.

i. Implement the formation control protocol

$$y_i(k+1) = y_i(k) + \gamma \sum_{j \in N(i)} w_{ij}(y_j(k) - y_i(k) - \Delta_{ij}), \tag{1}$$

in a differentially private manner.

ii. Analyze the relationship between network performance, privacy, and the underlying graph topology. △

We will solve Problem 1 by bounding the performance of the network in terms of the privacy parameters of each agent and the algebraic connectivity of the underlying graph. This will allow us to analyze the relationship between performance, privacy, and topology.

Remark 1: We consider formation control in $\mathbb{R}^n$, which is equivalent to running n independent scalar-valued formation controllers. Therefore, for simplicity we analyze the scalar case.

Before solving Problem 1, we give the necessary definitions for formation control. First, we define agent- and network-level dynamics. Then, we detail how each agent will enforce differential privacy. Lastly, we explain how differentially private communications affect the performance of a formation control protocol and how to quantify quality of a formation.

A. Multi-agent Formation control

The goal of formation control is for agents in a network to assemble into some geometric shape or set of relative states. Multi-agent formation control is a well researched problem and there are several mathematical formulations one can use to achieve similar results [18]–[24]. We will define relative distances between agents that communicate and the control objective is for all agents to maintain the relative distances to each of their neighbors. This approach is similar to that of [20] and the translationally invariant formations in [24].

For the formation to be feasible, $\Delta_{ij} = -\Delta_{ji}$ for all $(i, j) \in E$. The network control objective is driving $\lim_{k\to\infty}(y_j(k) - y_i(k)) = \Delta_{ij}$ for all $(i, j) \in E$. It is important to note that there is an infinite set of points that can be in formation; the formation can be centered around any point in $\mathbb{R}^n$ and meet the control requirement, i.e., we allow formations to be translationally invariant [24].

Now we define the agents' update law. Let $\{p_1, ..., p_N\}$ be any collection of points in formation such that $p_j - p_i = \Delta_{ij}$ for all $(i, j) \in E$ and let $p = (p_1^T, \ldots, p_N^T)^T \in \mathbb{R}^{nN}$ be the network-level formation specification. We consider the formation control protocol in Equation (1). As noted in Remark 1, we analyze convergence of Equation (1) at the component level. Thus, while $y_i \in \mathbb{R}^n$, we select an arbitrary $l \in \{1, \ldots, n\}$ and provide analysis for

$$x(k) = (y_{1,l}(k) \ldots y_{N,l}(k))^T \in \mathbb{R}^N,$$

i.e., each agents l^{th} component, which proceeds identically for each $l \in \{1, \ldots, n\}$. Below, we also use the vector of l^{th} components of p, denoted

$$q = (p_{1,l} \ldots p_{N,l})^T.$$

Let $\bar{x}(k) = x(k) - q$. Then we analyze

$$\bar{x}(k+1) = (I - \gamma L(\mathcal{G}))\bar{x}(k). \tag{2}$$

Letting $P = I - \gamma L(\mathcal{G})$, we may write $\bar{x}(k+1) = P\bar{x}(k)$. In this form, we have the following convergence result.

Lemma 2 ([23], Theorem 2): If $\mathcal{G}$ is connected, P is doubly stochastic, and $\gamma \in (0, \frac{1}{d_{max}})$, then the protocol in Equation (2) reaches consensus asymptotically and $\bar{x}(k) \to \mathbb{1}^T \frac{1}{n} \bar{x}(0) \mathbb{1}$. ■

Because the protocol in Equation (2) reaches consensus over $\bar{x}$, it solves the translationally invariant formation control problem [24]. Using δ_{ij} to denote the state offset between agents j and i in the appropriate dimension, the node-level protocol in Equation (1) can be rewritten for a single component as

$$x_i(k+1) = x_i(k) + \gamma \sum_{j \in N(i)} w_{ij}(x_j(k) - x_i(k) - \delta_{ij}),$$

which we use below.

B. Private Communications

When agent j transmits $\bar{x}_j(k)$ to the agents in $N(j)$, it is potentially exposing its state trajectory, x_j, to them and adversaries or eavesdroppers. Agent j therefore sends a differentially private version of $\bar{x}_j(k)$ to its neighborhood.

Agent j starts by selecting privacy parameters $\epsilon_j > 0$, $\delta_j \in (0, \frac{1}{2})$, and adjacency relation Adj_{b_j} with $b_j > 0$. Agent j then privatizes its state trajectory x_j with the Gaussian mechanism. Let $\tilde{x}_j$ denote the differentially private version of x_j, where, pointwise in time, $\tilde{x}_j(k) = x_j(k) + v_j(k)$, with $v_j(k) \sim \mathcal{N}(0, \sigma_j^2)$ and $\sigma_j \geq \kappa(\delta_j, \epsilon_j) b_j$. Thus agent j keeps the trajectory x_j differentially private. Agent j then shares $\tilde{\bar{x}}_j(k) = \tilde{x}_j(k) - q_j$, which is also differentially private because subtracting q_j is merely post-processing [14].

C. Private Formation Control

When each agent is sharing differentially private information, the node-level formation control protocol becomes

$$\bar{x}_i(k+1) = \bar{x}_i(k) + \gamma \sum_{j \in N(i)} w_{ij}(\tilde{\bar{x}}_j(k) - \bar{x}_i(k)), \quad (3)$$

where agent i uses $\bar{x}_i$ rather than $\tilde{\bar{x}}_i$ because it always has access to its own unprivatized state. The stochastic nature of this protocol implies that agents no longer exactly reach a formation, and, in particular, the states will never exactly converge to a steady-state value.

To analyze performance, let $\beta(k) := \frac{1}{N}\mathbb{1}^T x(k)\mathbb{1} + q - \frac{1}{N}\mathbb{1}^T q\mathbb{1}$, which is the state vector the protocol in Equation (2) would converge to with initial state $x(k)$ and without privacy. Also let $e(k) = x(k) - \beta(k)$, which is the distance of the current state to the state the protocol would converge to without differential privacy. To quantity the effects of privacy on the network as a whole, let $e_{\text{agg}}(k) := \frac{1}{n}\sum_{i=1}^{n} E[e_i^2(k)]$ be the aggregate error of the network, and let

$$e_{ss} := \limsup_{k \to \infty} e_{\text{agg}}(k)$$

be the steady-state error of the network.

Problem 1 requires us to quantify the relationship between privacy, encoded by (ϵ_i, δ_i); performance, encoded by e_{ss}; and topology, encoded by λ_2. These quantitative tradeoffs are the subject of the next section.

IV. Differentially Private Formation Control

In this section we solve Problem 1. First, we show how the private formation control protocol can be modeled as a Markov chain. Then, we solve Problem 1 by deriving performance bounds that are functions of the underlying graph topology and each agent's privacy parameters.

A. Formation Control as a Markov chain

Problem 1 takes the form of a consensus protocol with Gaussian i.i.d. noise perturbing each agent's state, which has been previously studied in [18]. We begin by expanding $\tilde{\bar{x}}_j(k)$ in Equation (3), which yields

$$\bar{x}_i(k+1) = \bar{x}_i(k) + \gamma \sum_{j \in N(i)} w_{ij}(\bar{x}_j(k) + v_j(k) - \bar{x}_i(k)). \quad (4)$$

For the purposes of analysis, we will consider equivalent network-level dynamics given as follows.

Lemma 3: Let agents use the communication graph $\mathcal{G}$ with weighted Laplacian $L(\mathcal{G})$. Then Equation (4) can be represented at the network level as $\bar{x}(k+1) = P\bar{x}(k) + z(k)$, where $P = I - \gamma L(\mathcal{G})$ and $z(k) \sim \mathcal{N}(0, Z)$ where $Z = \text{diag}(s_1^2, ..., s_N^2)$, with $s_i^2 = \gamma^2 \sum_{j \in N(i)} w_{ij}^2 \sigma_j^2$.

Proof: See Lemma 3 in [25]. ∎

For analysis, we use the network-level update law $\bar{x}(k+1) = P\bar{x}(k) + z(k)$. The main result of this paper uses the fact that a stochastic matrix P can serve as the transition matrix of a Markov chain and the properties of the Markov chain can be used to analyze the network dynamics.

B. Solving Problem 1

Now we state the first of our main results: a bound on performance in terms of agents' level of privacy and underlying graph topology.

Theorem 1: Consider the network-level private formation control protocol $\bar{x}(k+1) = (I - \gamma L(\mathcal{G}))\bar{x}(k) + z(k)$. If $\gamma \sum_{j \in N(i)} w_{ij} < 1$, $\gamma \in \left(0, \frac{1}{d_{max}}\right)$, $\mathcal{G}$ is connected and undirected, and $\sigma_i \geq \kappa(\delta_i, \epsilon_i) b_i$ for all i, then e_{ss} is upper-bounded by

$$e_{ss} \leq \frac{\gamma(N-1)^2 \max_i \kappa(\delta_i, \epsilon_i)^2 b_i^2}{N\lambda_2(L(\mathcal{G}))(2 - \gamma\lambda_2(L(\mathcal{G})))}.$$

Proof: See [25, Theorem 1]. ∎

We can simplify Theorem 1 when each agent has the same privacy parameters. Next, and from this point on, we consider the case where $\sigma = \kappa(\delta, \epsilon)b$ so that each agent adds the minimum amount of noise needed to attain (ϵ, δ)-differential privacy.

Corollary 1 (Homogeneous Privacy Parameters): Let each agent in the network have the privacy parameters ϵ and δ and the adjacency parameter b. Then $e_{ss} \leq \frac{\gamma\kappa(\delta,\epsilon)^2 b^2 (N-1)^2}{N\lambda_2(L(\mathcal{G}))(2-\gamma\lambda_2(L(\mathcal{G})))}$.

The rest of the paper focuses on the homogeneous case presented in Corollary 1, though all forthcoming results are easily adapted to the heterogeneous case by considering minima and maxima over all agents where appropriate.

V. Network Design Guidelines

In this section we give guidelines for designing a differentially private formation control network. The goal is to design the network so that e_{ss} does not exceed a given limit e_R. The question of interest is: Given a specific communication topology, how much privacy is each agent allowed to have for $e_{ss} \leq e_R$? As noted in Remark 1, we do this for each dimension of formation control individually. A smaller value of ϵ corresponds to being more private. Therefore an upper bound on e_{ss}, which is the measure of system performance, implies a lower bound on ϵ, each agent's privacy parameter.

We derive an impossibility result and sufficient conditions in terms of ϵ for $e_{ss} \leq e_R$ for specific networks. We consider connected graphs $\mathcal{G}$ with uniform weights, where $w_{ij} = w$ for all $(i,j) \in E$. By construction, the graphs we consider in this paper are weight-balanced, which implies that for any weights, the protocol in Equation (2) will converge to the unweighted average as seen in Lemma 2. Throughout this section we fix δ to be some small number and let ϵ vary to tune the level of privacy, which is common in differential privacy implementations [26].

Theorem 2 (Impossibility Result): Given a network of N agents with specified ϵ, δ, b, and e_R, compute $\lambda_2(\mathcal{G})$. Then $e_{ss} \leq e_R$ cannot be assured if $\epsilon < \frac{2bz_1}{Ne_R\lambda_2(\mathcal{G})}\left(b + \frac{e_R K_\delta \lambda_2(\mathcal{G}) N}{\sqrt{e_R z_1 \lambda_2(\mathcal{G}) N}}\right)$, where $z_1 = \frac{\gamma(N-1)^2}{2-\gamma\lambda_2(\mathcal{G})}$.

Proof: See Theorem 2 in [25]. ∎

We now derive necessary and sufficient conditions for assuring $e_{ss} \leq e_R$ for common graphs: the complete graph, line graph, cycle graph, and star graph. These conditions

$\mathcal{G}$ \ N	10	100	1,000	10,000
Complete	0.0074	0.0081	0.0084	0.0116
Cycle	0.0380	1.4514	199.35	159591
Line	0.7533	3.2127	714.70	635752
Star	0.0235	0.0820	0.2661	0.8849

TABLE I: Comparison of the lower bounds on ϵ for various communication topologies and numbers of agents. This table illustrates that more-connected graphs accommodate privacy better when the network size grows, because they allow ϵ to be smaller, which gives stronger privacy protections.

can easily be checked a priori and give a network designer a simple means of determining whether a specific network will meet performance requirements. Proofs of Corollaries 3-5 are similar to that of Corollary 2 and are omitted.

Corollary 2 (Complete graph): The complete graph has algebraic connectivity $\lambda_2(L(\mathcal{G})) = wN$ [27]. Consider a network of N agents with specified ϵ, δ, γ, b, w, and e_R, and communication topology modeled by the complete graph. The network can be shown to satisfy $e_{ss} \leq e_R$ if and only if $\epsilon \geq \frac{2b\gamma(N-1)^2}{N^2 e_R w(2-\gamma w N)}\left(b + \frac{e_R K_\delta w N\sqrt{2-\gamma w N)}}{(N-1)\sqrt{e_R\gamma w}}\right)$.

Proof: See Corollary 2 in [25]. ■

Corollary 3 (Cycle Graph): The cycle graph has algebraic connectivity $\lambda_2(L(\mathcal{G})) = 2w\left(1-\cos\left(\frac{2\pi}{N}\right)\right)$ [27]. Consider a network of N agents with specified ϵ, δ, γ, b, w, and e_R, and communication topology modeled by the cycle graph. The network is assured to satisfy $e_{ss} \leq e_R$ if and only if $\epsilon \geq \frac{bz_2}{Ne_R 2w\left(1-\cos\left(\frac{2\pi}{N}\right)\right)} + \frac{K_\delta}{\sqrt{z_2 e_R w\left(1-\cos\left(\frac{2\pi}{N}\right)\right)N}}$, where $z_2 = \frac{(N-1)^2\gamma}{1-\gamma w\left(1-\cos\left(\frac{2\pi}{N}\right)\right)}$.

Corollary 4 (Line Graph): The line graph has algebraic connectivity $\lambda_2(L(\mathcal{G})) = 2w\left(1-\cos\left(\frac{\pi}{N}\right)\right)$ [27]. Consider a network of N agents with specified ϵ, δ, γ, b, w, and communication topology modeled by the line graph. The network is assured to satisfy $e_{ss} \leq e_R$ if and only if $\epsilon \geq \frac{bz_3}{Ne_R 2w\left(1-\cos\left(\frac{\pi}{N}\right)\right)} + \frac{K_\delta}{\sqrt{z_3 e_R w\left(1-\cos\left(\frac{\pi}{N}\right)\right)N}}$, where $z_3 = \frac{(N-1)^2\gamma}{1-\gamma w\left(1-\cos\left(\frac{\pi}{N}\right)\right)}$.

Corollary 5 (Star Graph): The star graph has algebraic connectivity $\lambda_2(L(\mathcal{G})) = w$ [27]. Consider a network of N agents with specified ϵ, δ, γ, b, w, and e_R, and communication topology modeled by the star graph. The network is assured to satisfy $e_{ss} \leq e_R$ if and only if $\epsilon \geq \frac{2b\gamma(N-1)^2}{Ne_R w(2-\gamma w)}\left(b + \frac{e_R K_\delta w N\sqrt{2-\gamma w)}}{(N-1)\sqrt{e_R\gamma w N}}\right)$.

Remark 2: Fix $\delta = 0.01$, $b = 5$,$w = 1$, $\gamma = 10^{-4}$, and $e_R = 100$. The lower bounds on ϵ found in Corollaries 2-5 were calculated numerically for networks with a varying number of agents, the results of which are in Table I.

VI. Sensitivity Results

Theorem 1 and Corollaries 2-5 show that performance of a network is a function of the network topology, each agent's privacy parameters, adjacency relationship, step size, and the number of agents. Some of these parameters are global, in that the parameter depends on the entire network, and some are local, in that the parameter can change at the agent level. For example, the network communication topology is a global parameter while each agent's privacy parameter, ϵ, is a local parameter.

Consider the following example: Given a network that is not performing as desired, one option is to change the network's topology and allow more agents to communicate, while another option is loosening the agents' privacy requirements. Depending on design constraints, it may be more effective to allow more agents to communicate or to relax privacy requirements. It is useful to understand when changing ϵ is more effective than changing the network topology and vice versa. In this section we therefore analyze how sensitive network performance is to local changes in privacy and global changes in topology.

Theorem 3: Let $\eta_1 = \frac{2\epsilon\gamma+\gamma K_\delta^2+2}{2\gamma} + \frac{1}{2}\sqrt{2\epsilon K_\delta^2 + K_\delta^4}$, $\eta_2 = \frac{2\epsilon\gamma+\gamma K_\delta^2+2}{2\gamma} - \frac{1}{2}\sqrt{2\epsilon K_\delta^2 + K_\delta^4}$, $\alpha = \epsilon^2 + \frac{3\epsilon K_\delta^2}{2} + \frac{1}{\gamma^2} + \frac{K_\delta^4}{2}$, and $\mu = \frac{K_\delta^2(4\epsilon^2+4\epsilon K_\delta^2+K_\delta^4)}{2\sqrt{2\epsilon K_\delta^2+K_\delta^4}}$. Then e_{ss} is more sensitive to λ_2 than ϵ when $\lambda_2 > \eta_1 - \sqrt{\alpha+\mu}$ or when $\lambda_2 < \eta_2 - \sqrt{\alpha-\mu}$.

Proof: See Theorem 3 in [25] ■

Remark 3: These results can be formulated in such a way that there is some cost associated with changing $\lambda_2(L(\mathcal{G}))$ and a cost associated with changing ϵ. Making an optimal change to achieve performance criteria will largely depend on application. This will be explored in a future publication.

The results presented in Theorem 3 can be instantiated for specific graphs. For example, consider the following.

Corollary 6: Let $\delta = 0.00135$, such that $K_\delta = 3$, and let $\epsilon = 0.01$. Let $\gamma = \frac{1}{10}$. The network's performance is more sensitive to the network topology than ϵ when $\lambda_2 > 5.55134$.

To illustrate these results, consider the following. The star graph over $N = 10$ nodes has $\lambda_2 = 1$, which implies that the network's performance is more sensitive to changes in the privacy parameter ϵ. The complete graph over $N = 10$ nodes has $\lambda_2 = 10$, which implies the network's performance is more sensitive to changes in the network topology.

VII. Simulation Results

In this section, we present private formation control simulation results. Consider a network of $N = 5$ agents running a differentially private formation controller. Agents i's state at time k is $y_i(k) \in \mathbb{R}^2$, and every agent's state trajectory is in $\tilde{\ell}_2^2$. The agents' communication topology is modeled by the star graph over 5 nodes with weights $w_{ij} = 1$ for all $(i,j) \in E$. The network's algebraic connectivity is $\lambda_2 = 1$. The formation specification is

$$p = \begin{bmatrix} 0 & -20 & 20 & 20 & -20 \\ 0 & 20 & 20 & -20 & -20 \end{bmatrix}^T,$$

where row i denotes agent i's desired location in the formation. Thus p specifies a formation where agents 2-5 will form a square with agent 1 at the center.

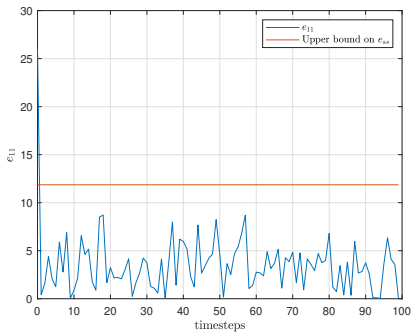

Fig. 1: Agent 1's error in the first element of its state and the upper bound on e_{ss}. The upper bound is on the steady state value of e_{ss}, however it holds point-wise in time for e_{11} and components of other agents' states.

We consider the homogeneous case where each agent has identical privacy parameters, $(\epsilon_i, \delta_i) = (\ln 3, 0.00135)$ for all i and every agent also has an identical adjacency parameter $b_i = 2$ for all i. Let $\gamma = \frac{1}{5}$. Let e_{11} denote the error of the first element of agent 1's state. The protocol in Equation (3) was run for 100 time steps

Figure 1 shows e_{11} at every time step as well as the upper bound found in Theorem 1, where we see that e_{11} never converges to 0 due to the stochastic nature of the protocol, but remains in some neighborhood of 0. The bound on e_{ss} presented in Theorem 1 is on the expected steady state value of square aggregate error, though we see that this bound also holds point-wise in time for e_{11} in this simulation. These results were typical throughout numerous simulation runs.

VIII. Conclusions

In this paper, we have studied the problem of differentially private formation control. This work enables agents to assemble formations while only sharing differentially private output data with a bounded steady state error. We developed guidelines for calibrating privacy under different control-theoretic requirements. The tunable parameters in this work are the privacy parameters and the topology itself, balancing the corresponding trade offs is a subject of future work.

References

[1] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.

[2] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of cryptography conference*. Springer, 2006, pp. 265–284.

[3] S. P. Kasiviswanathan and A. Smith, "On the'semantics' of differential privacy: A bayesian formulation," *Journal of Privacy and Confidentiality*, vol. 6, no. 1, 2014.

[7] J. Le Ny and M. Mohammady, "Differentially private mimo filtering for event streams," *IEEE Transactions on Automatic Control*, vol. 63, no. 1, pp. 145–157, 2017.

[4] J. Le Ny and G. J. Pappas, "Differentially private filtering," *IEEE Transactions on Automatic Control*, vol. 59, no. 2, pp. 341–354, 2013.

[5] K. Yazdani, A. Jones, K. Leahy, and M. Hale, "Differentially private lq control," *arXiv preprint arXiv:1807.05082*, 2018.

[6] M. T. Hale and M. Egerstedt, "Cloud-enabled differentially private multiagent optimization with constraints," *IEEE Transactions on Control of Network Systems*, vol. 5, no. 4, pp. 1693–1706, 2017.

[8] A. Jones, K. Leahy, and M. Hale, "Towards differential privacy for symbolic systems," in *2019 American Control Conference (ACC)*. IEEE, 2019, pp. 372–377.

[9] 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*, ser. WPES '12. New York, NY, USA: Association for Computing Machinery, 2012, p. 81–90. [Online]. Available: https://doi.org/10.1145/2381966.2381978

[10] Y. Wang, Z. Huang, S. Mitra, and G. E. Dullerud, "Differential privacy in linear distributed control systems: Entropy minimizing mechanisms and performance tradeoffs," *IEEE Transactions on Control of Network Systems*, vol. 4, no. 1, pp. 118–130, 2017.

[11] Z. Xu, K. Yazdani, M. T. Hale, and U. Topcu, "Differentially private controller synthesis with metric temporal logic specifications," in *2020 American Control Conference (ACC)*. IEEE, 2020, pp. 4745–4750.

[12] Y. Wang, M. Hale, M. Egerstedt, and G. E. Dullerud, "Differentially private objective functions in distributed cloud-based optimization," in *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 3688–3694.

[13] M. Fiedler, "Algebraic connectivity of graphs," *Czechoslovak mathematical journal*, vol. 23, no. 2, pp. 298–305, 1973.

[14] C. Dwork, "Differential privacy," *Automata, languages and programming*, pp. 1–12, 2006.

[15] K. Yazdani and M. Hale, "Error bounds and guidelines for privacy calibration in differentially private kalman filtering," in *2020 American Control Conference (ACC)*, 2020, pp. 4423–4428.

[16] J. Le Ny, *Differential Privacy for Dynamic Data*. Springer, 2020.

[17] B. Hajek, *Random processes for engineers*. Cambridge university press, 2015.

[18] A. Jadbabaie and A. Olshevsky, "Scaling laws for consensus protocols subject to noise," 2015.

[19] L. Krick, M. E. Broucke, and B. A. Francis, "Stabilisation of infinitesimally rigid formations of multi-robot networks," *International Journal of control*, vol. 82, no. 3, pp. 423–439, 2009.

[20] W. Ren, R. W. Beard, and E. M. Atkins, "Information consensus in multivehicle cooperative control," *IEEE Control systems magazine*, vol. 27, no. 2, pp. 71–82, 2007.

[21] W. Ren, "Consensus strategies for cooperative control of vehicle formations," *IET Control Theory & Applications*, vol. 1, no. 2, pp. 505–512, 2007.

[22] J. A. Fax and R. M. Murray, "Information flow and cooperative control of vehicle formations," *IEEE transactions on automatic control*, vol. 49, no. 9, pp. 1465–1476, 2004.

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

[24] M. Mesbahi and M. Egerstedt, *Graph theoretic methods in multiagent networks*. Princeton University Press, 2010.

[25] C. Hawkins and M. Hale, "Differentially private formation control," *arXiv preprint arXiv:2004.02744*, 2020.

[26] J. Hsu, M. Gaboardi, A. Haeberlen, S. Khanna, A. Narayan, B. C. Pierce, and A. Roth, "Differential privacy: An economic method for choosing epsilon," in *2014 IEEE 27th Computer Security Foundations Symposium*. IEEE, 2014, pp. 398–410.

[27] N. M. M. De Abreu, "Old and new results on algebraic connectivity of graphs," *Linear algebra and its applications*, vol. 423, no. 1, pp. 53–73, 2007.