Differentially Private Formation Control

Calvin Hawkins & Matthew Hale
Department of Mechanical and Aerospace Engineering
University of Florida

Center of Excellence on Assured Autonomy in Contested Environments
Fall 2020 Review
October 29th, 2020

UF | Department of Mechanical & Aerospace Engineering
UNIVERSITY of FLORIDA

UF | UNIVERSITY of FLORIDA
Duke UNIVERSITY
TEXAS The University of Texas at Austin
UC SANTA CRUZ

Motivation: Work together but keep secrets

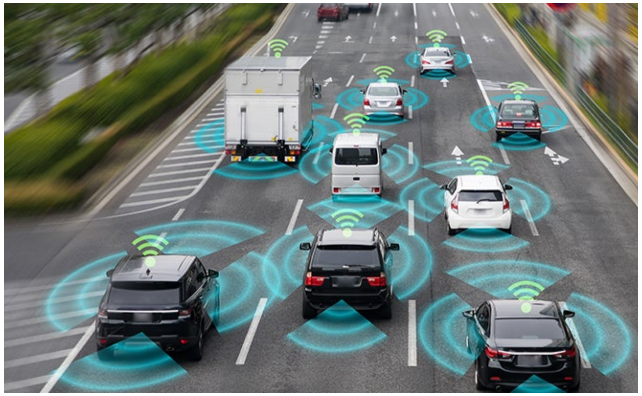

- Allow agents to collaborate while protecting their sensitive information.
- Examples:
  - Coalitions collaborating but maintaining secrecy
  - Autonomous vehicles sharing location data
  - Social Networks sharing personal information
  - Data-driven control sharing sensitive state information

UF UNIVERSITY of FLORIDA · Duke UNIVERSITY · TEXAS The University of Texas at Austin · UC SANTA CRUZ

Differential Privacy Can Help Us

- Statistical notion of privacy from computer science

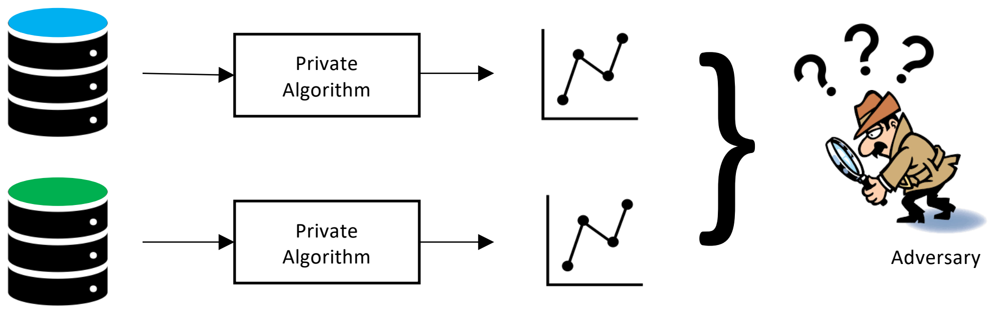

- Immune to post-processing and robust to side information.
- Used by Apple, Google, Uber, and the 2020 Census.
- In multi-agent control, agents can share state trajectory data while protecting itself from other agents and eavesdroppers.

Differential Privacy Definitions

- Goal of Differential Privacy: Make "similar" pieces of data appear "approximately indistinguishable"
- Adjacency defines when pieces of data are "similar:"

$$\mathrm{Adj}_{\mathrm{b}_i}(\mathrm{x}_i, \mathrm{x}'_i) = \begin{cases} 1 & \|\mathrm{x}_i - \mathrm{x}'_i\|_{\ell_p} \leq \mathrm{b}_i \\ 0 & \text{else} \end{cases}$$

Definition of Differential Privacy (Approximate indistinguishability)

Let $\epsilon_i > 0$ and $\delta_i \in \left[0, \frac{1}{2}\right)$. A randomized mechanism M is (ϵ_i, δ_i) −differentially private for agent i if, for all adjacent x_i, x_i', we have

$$P[M(x_i) \in S] \leq e^{\epsilon_i} P[M(x_i') \in S] + \delta_i$$

Sensitive Data → Privacy Mechanism tuned by ϵ_i, δ_i → Private Data

We will privatize formation control

- Consider a network of N agents where agent i has state $x_i(k) \in \mathbb{R}^n$ at time k
- The network communication topology is modeled by a weighted, undirected graph $\mathcal{G}$
- If agents i and j communicate, they maintain a distance of $\Delta_{ij} \in \mathbb{R}^n$

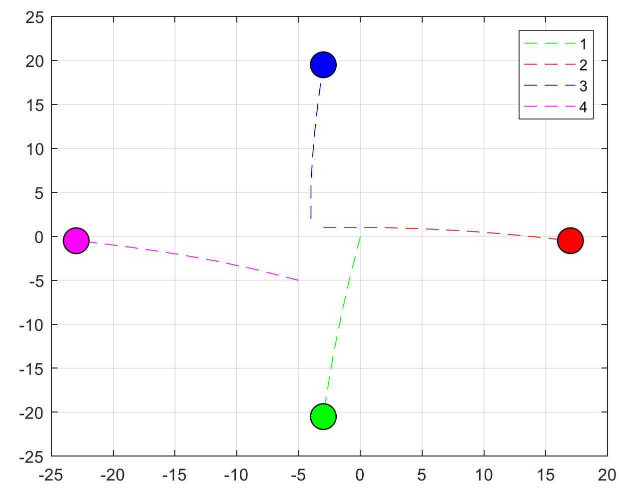

- Without privacy, this is achieved by the formation control protocol

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

Problem Statement: Private Formation Control

Problem Statement:

(i) Implement the formation control protocol

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

in a differentially private manner

(ii) Quantify tradeoffs between network performance, privacy, and graph topology

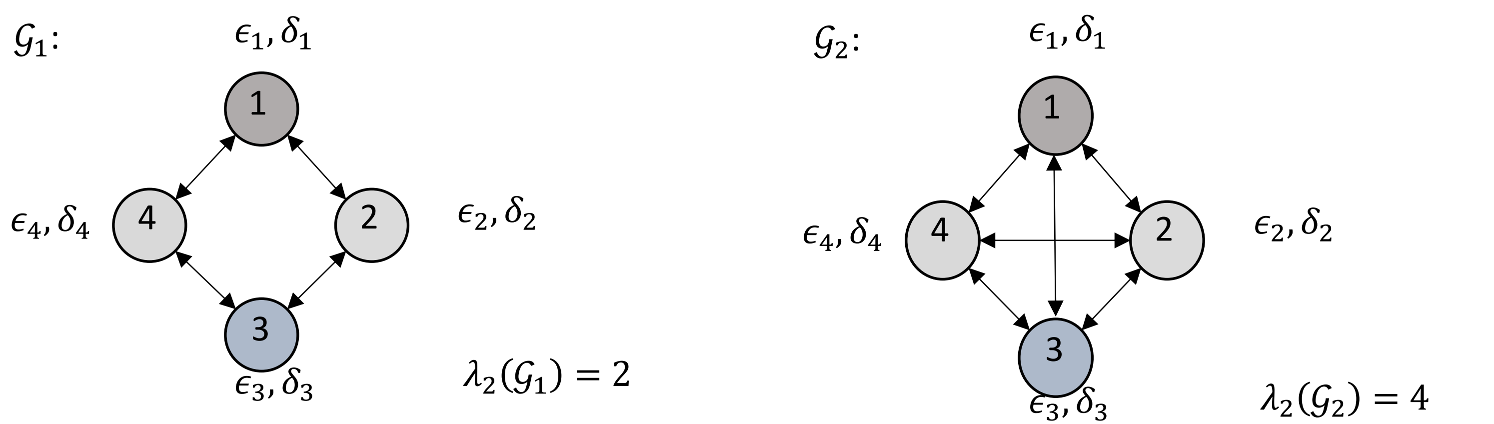

Privacy is enforced when communicating

- Agent i must send its state to its neighbors in $N(i)$ at every timestep k
- Agent i will send a private version of its state, denoted $\tilde{x}_i(k)$
- Differential privacy is achieved at the trajectory level with the Gaussian Mechanism:

$$\tilde{x}_i(k) = x_i(k) + v_i(k)$$
$$v_i(k) \sim \mathcal{N}(0, \sigma_i^2 I_n)$$

Lemma: The Gaussian mechanism is (ϵ_i, δ_i) −differentially private for agent i if $\sigma_i \geq \kappa(\epsilon_i, \delta_i) b_i$, where $\kappa(\delta_i, \epsilon_i) = \frac{1}{2\epsilon_i}\left(K_{\delta_i} + \sqrt{K_{\delta_i}^2 + 2\epsilon_i}\right)$, and $K_{\delta_i} = Q^{-1}(\delta_i)$.

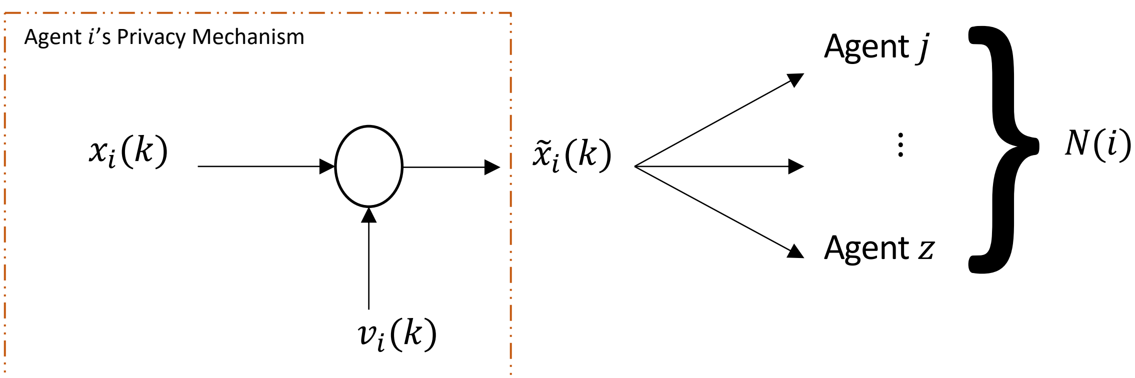

We have private formation control

- With privacy, the formation control protocol is

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

- Privacy induces uncertainty ⇒ formations are imperfect

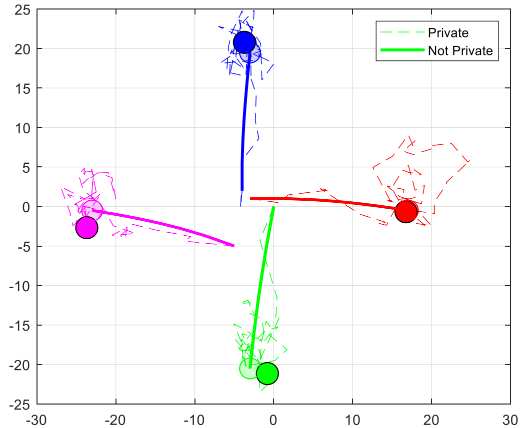

- Let $e_i(k) = x_i(k) - \beta_i(k)$, where $\beta(k)$ is the state the non-private protocol converges to with initial condition $x(k)$.

- To quantify performance at the network level, let

$$e_{ss} = \limsup_{k\to\infty} \sum_{i=1}^{N} E[e_i(k)^2]$$

Formation Error Bounds

Theorem 1: Bounds on Steady-State Error

A network running the formation control protocol

$$x_i(k+1) = x_i(k) + \gamma \sum_{j \in N(i)} w_{ij}(\color{red}{\tilde{x}_j(k)} - x_i(k) - \Delta_{ij})$$

over a connected, undirected, weighted graph $\mathcal{G}$, is differentially private and has e_{ss} upper bounded by

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

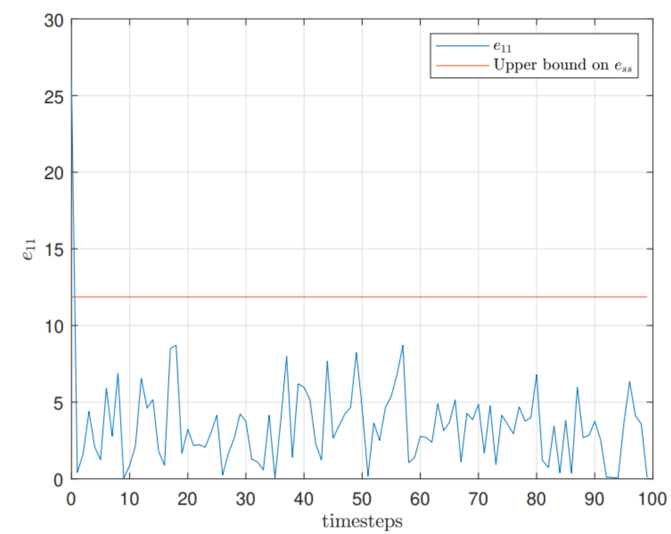

Example: The effect of ϵ on performance

- Fix $\delta_i = 0.05$ for all i. Fix the communication topology $\mathcal{G}$.
- Recall: Smaller $\epsilon_i \implies$ stronger privacy for agent i.

$\mathcal{G}$:

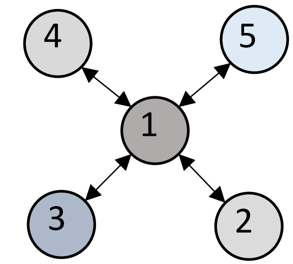

<u>Case 1:</u>
$\epsilon_i = 0.8$ for all i.
$e_{ss} \leq 6.25$

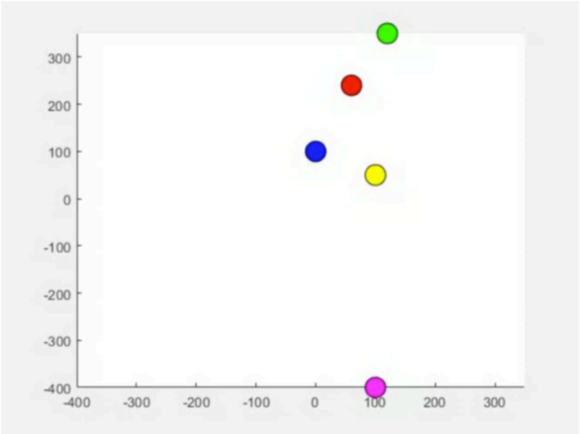

<u>Case 2:</u>
$\epsilon_i = 0.1$ for all i.
$e_{ss} \leq 303.79$

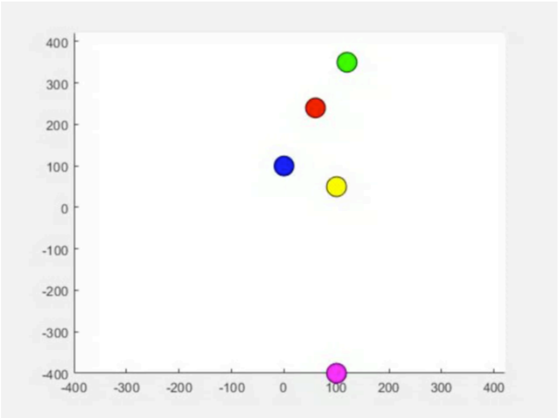

There are fundamental limits to privacy

- Suppose we must design a private formation control network: we are given that the steady state error of the system must not exceed e_R
- Given a graph G and homogeneous privacy parameter ε, will it work?

Theorem 3: Impossibility Result

It is impossible to construct a differentially private formation controller that meets the performance requirement e_R if

$$\epsilon < \frac{2bz}{N e_R \lambda_2(\mathcal{G})}\left(b + \frac{e_R K_\delta \lambda_2(G) N}{\sqrt{e_R z \lambda_2(\mathcal{G}) N}}\right),$$

where $z = \frac{\gamma(N-1)^2}{2-\gamma\lambda_2(\mathcal{G})}$.

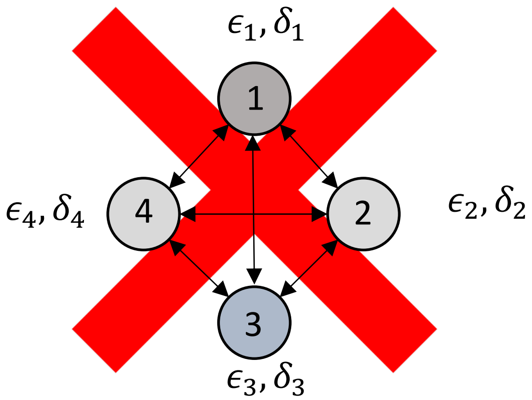

Next steps: Privacy & Network Codesign

- Agents want to be as private possible but also want to maximize performance
- Seen impact of changing ε. What about changing the topology of G?

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

- What is the optimal network design? Who communicates with whom?
  - Constraints: Formation error, edge budget, user preferences

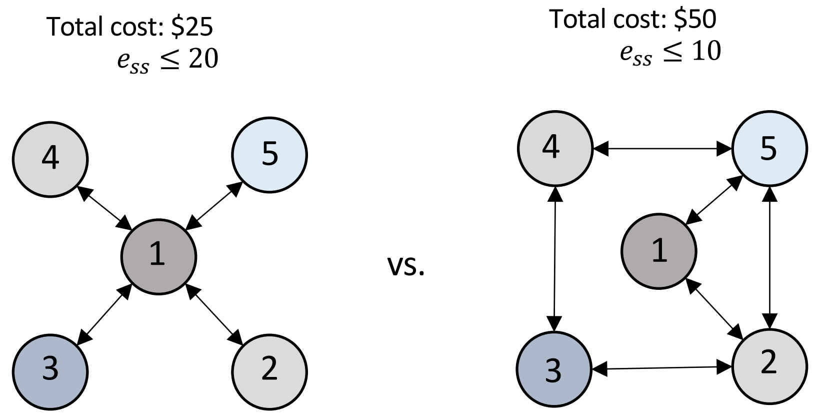

- Preliminary results: problem is quasiconvex, numerically difficult

Thank you