Deception in Linear Quadratic Control

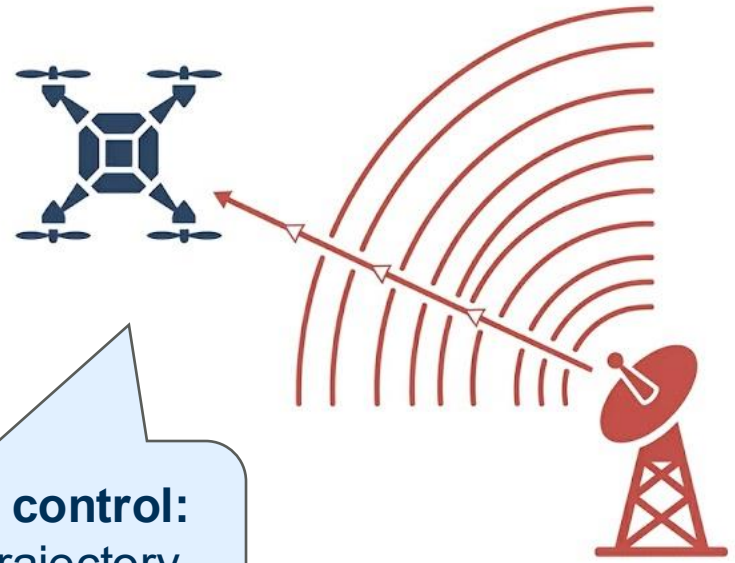
Yerin Kim[1], Haosheng Zhou[2], Alexander Benvenuti[1], Ruimeng Hu[2], Matthew Hale[1]

[1]Georgia Institute of Technology, [2]University of California Santa Barbara

Georgia Tech

Optimal LQ control reveals information

- **Linear-Quadratic (LQ) control** is a standard tool in control engineering.
  - Smart grid
  - UAV
  - Satellites
- The optimal LQ control **reveals sensitive information** such as:
  - Cost function[1]
  - Target states[2]
- Revealing such information allows informed attacks, resulting in system vulnerability.



Standard LQ control: Optimize its trajectory under given objective

Adversarial observer: Infer information from observed trajectory

[1] H. Zhang, Y. Li, and X. Hu, "Inverse optimal control for finite-horizon discrete-time linear quadratic regulator under noisy output," in 2019 IEEE 58th conference on decision and control (CDC). IEEE, 2019, pp. 6663–6668.
[2] Y. Li, C. Yu, H. Fang, and J. Chen, "Inverse optimal control for linear quadratic tracking with unknown target states," Automatica, vol. 185, p. 112819, 2026.

Georgia Tech

We will perturb observations to hide objectives

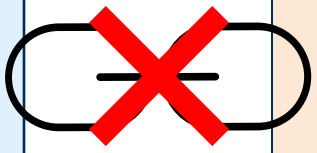
- To hide information, we break the link between **what we do** and **what they see**.

Keep



what we do

- Primary objective of the system
- Achieved by minimizing cost function



Perturb



what they see

- Observed state trajectories
- Naturally reveals information about system objective

Question 1: Can we introduce deceptive perturbations to trajectories while pursuing primary objective?

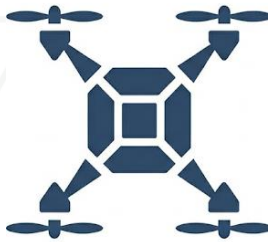
Question 2: Even when the attacker knows the presence of deception, can we secretly perturb trajectories without being detected?

Georgia Tech

We consider red team vs. blue team setting

- We consider red team vs. blue team setting, where blue team uses LQ control.

Blue Team



- Blue team's state dynamics with Gaussian noise

$$\begin{cases} \boldsymbol{v}_{t+1} = \alpha_t + \boldsymbol{v}_t + \boldsymbol{z}_t & \text{Acceleration control} \quad \text{(Used for primary task)} \\ \boldsymbol{p}_{t+1} = (\boldsymbol{v}_t + \beta_t) + \boldsymbol{p}_t + \boldsymbol{y}_t & \text{Instantaneous velocity control} \quad \text{(Used for deception)} \end{cases}$$

  - Gaussian noise follows $\boldsymbol{z}_t \sim \mathcal{N}(0, \Sigma_z^2), \boldsymbol{y}_t \sim \mathcal{N}(0, \Sigma_y^2)$
  - This system can also be expressed in linear form: $\boldsymbol{x}_{t+1} = A\boldsymbol{x}_t + B\boldsymbol{u}_t + \boldsymbol{w}_t$

- Blue team's primary objective

Note: Non-deceptive optimal control gives $\beta_t^* \equiv 0$.

  - Follow reference trajectory $\{\overline{\boldsymbol{v}}_t\}_{t=0}^T$ with minimum control effort.

Red Team



- Red team observes blue team's trajectories and try to find out if blue team is using deception.

If blue team is **NOT using deception**: Trust observed trajectory and make incorrect downstream inference

If blue team is **introducing deception**: Do not trust observed trajectory and potentially take counter-actions

Georgia Tech

We model blue team as LQR and red team using SPRT

- **Blue team model:** Standard LQ control with reference trajectory (No deception)

Running cost

Terminal cost

$$\underset{\boldsymbol{u}}{\text{minimize}} \quad \mathbb{E}\left[\sum_{t=0}^{T-1}\left(\frac{1}{2}(\boldsymbol{x}_t-\bar{\boldsymbol{x}}_t)^T Q_t(\boldsymbol{x}_t-\bar{\boldsymbol{x}}_t)+\frac{1}{2}\boldsymbol{u}_t^T R\boldsymbol{u}_t\right)+\frac{1}{2}(\boldsymbol{x}_T-\bar{\boldsymbol{x}}_T)^T Q_T(\boldsymbol{x}_T-\bar{\boldsymbol{x}}_T)\right] = J^{primary}(\boldsymbol{u})$$

$$\text{subject to.} \quad \boldsymbol{x}_{t+1}=A\boldsymbol{x}_t+B\boldsymbol{u}_t+\boldsymbol{w}_t \quad \leftarrow \text{Linear state dynamics}$$

A control law $\beta_t = F_t^b \boldsymbol{v}_t + F_t^c \boldsymbol{p}_t + \boldsymbol{f}_t^d$ is given, which encodes a direction of perturbation.

- **Red team model:** Sequential Probability Ratio Test (SPRT)
  - Null hypothesis (H_0): Blue team is not using deception ($\beta_t^* = 0$)
  - Alternative hypothesis (H_1): Blue team is introducing deception ($\beta_t^* = F_t^b \boldsymbol{v} + F_t^c \boldsymbol{x} + \boldsymbol{f}_t^d$)

$$\log L_T = \log \frac{\mathbb{P}_{H_1}(\{\boldsymbol{x}_t\}_{t=0}^T)}{\mathbb{P}_{H_0}(\{\boldsymbol{x}_t\}_{t=0}^T)}$$

Red team accepts H_1 **(Deception detected)**

Red team accepts H_0 **(Deception undetected)**

Georgia Tech

The blue team introduce deception by augmenting cost

- We use SPRT statistics to augment the cost function.

$$\begin{aligned} \underset{\boldsymbol{u}}{\text{minimize}} \quad & J^{primary}(\boldsymbol{u}) - \lambda \mathbb{E} \log L_T \;\; = J^{blue}(\boldsymbol{u}) \\ \text{subject to.} \quad & \boldsymbol{x}_{t+1} = A\boldsymbol{x}_t + B\boldsymbol{u}_t + \boldsymbol{w}_t \end{aligned}$$

Note:

$\mathbb{E} \log L_T$

β_t gets close to $\beta_t^* = F_t^b \boldsymbol{v}_t + F_t^c \boldsymbol{x}_t + \boldsymbol{f}_t^d$

β_t gets close to $\beta_t^* = 0$

- Maximizing $\mathbb{E} \log L_T$ results in more deceptive control.
- $\mathbb{E} \log L_T$ can be nicely integrated into quadratic running cost.

$$J^{blue}(\boldsymbol{u}) = \mathbb{E}\left[\sum_{t=0}^{T-1} \left(\frac{1}{2}\boldsymbol{x}_t^T Q_t \boldsymbol{x}_t + \frac{1}{2}\boldsymbol{u}_t^T R \boldsymbol{u}_t + \boldsymbol{x}_t^T N_t \boldsymbol{u}_t + \boldsymbol{q}_t^T \boldsymbol{x}_t + \boldsymbol{r}_t^T \boldsymbol{u}_t + d_t\right) + \frac{1}{2}(\boldsymbol{x}_T - \bar{\boldsymbol{x}}_T)^T Q_T (\boldsymbol{x}_T - \bar{\boldsymbol{x}}_T)\right]$$

$$Q_t = \begin{bmatrix} R_v + \lambda F_t^{bT}(\Sigma_y^2)^{-1} F_t^b & \lambda F_t^{bT}(\Sigma_y^2)^{-1} F_t^c \\ \lambda F_t^{cT}(\Sigma_y^2)^{-1} F_t^b & \lambda F_t^{cT}(\Sigma_y^2)^{-1} F_t^c \end{bmatrix}, R = \begin{bmatrix} R_\alpha & 0 \\ 0 & R_\beta \end{bmatrix}, N_t = \begin{bmatrix} 0 & -\lambda F_t^{bT}(\Sigma_y^2)^{-1} \\ 0 & -\lambda F_t^{cT}(\Sigma_y^2)^{-1} \end{bmatrix},$$

$$q_t = \begin{bmatrix} -R_v \bar{\boldsymbol{v}}_t + \lambda F_t^{bT}(\Sigma_y^2)^{-1} \boldsymbol{f}_t^d \\ \lambda F_t^{cT}(\Sigma_y^2)^{-1} \boldsymbol{f}_t^d \end{bmatrix}, r_t = \begin{bmatrix} 0 \\ -\lambda(\Sigma_y^2)^{-1} \boldsymbol{f}_t^d \end{bmatrix}, d_t = \frac{1}{2}\bar{\boldsymbol{v}}_t^T R_v \bar{\boldsymbol{v}}_t + \frac{\lambda}{2}\boldsymbol{f}_t^{dT}(\Sigma_y^2)^{-1}\boldsymbol{f}_t^d$$

← Preserves quadratic structure!

Takeaway: As augmented problem preserves LQ structure, optimal control can be solved efficiently!

$$\boldsymbol{u}_t^* = -H_t^{-1}(G_t^T \boldsymbol{x}_t + \boldsymbol{r}_t + B^T \boldsymbol{s}_{t+1})$$

Georgia Tech

Solutions to augmented LQ are well behaved

- Augmented problem has unique solution under certain condition.

Gain in cost for minimizing β

Variance of Gaussian noise y

$$I_n - \lambda\left(\Sigma_y^2\right)^{-\frac{1}{2}} R_\beta^{-1}\left(\Sigma_y^2\right)^{-\frac{1}{2}} \succcurlyeq 0$$

Under this condition, there exists unique solution to the optimal control.

- In 1D case, the condition reduces to $\lambda \in [0, R_\beta \Sigma_y^2]$.



$\lambda = 0$

$\beta_t = 0$

(No deception)

$\lambda = R_\beta \Sigma_y^2$

$\beta_t = F_t^b \boldsymbol{v}_t + F_t^c \boldsymbol{x}_t + \boldsymbol{f}_t^d$

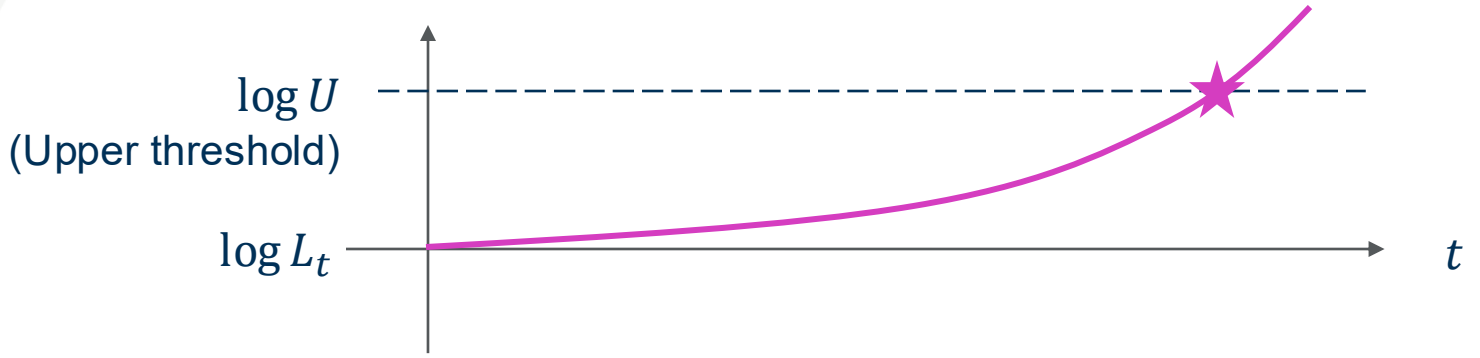
(Maximal deception)

Takeaway: We can introduce deceptive perturbation by augmenting quadratic cost with $\mathbb{E} \log L_T$.

Georgia Tech

The blue team requires stealthy deception

- Red team detects the deception when the $\log L_t$ gets larger.



Red team accepts H_1
(Deception detected)

- At time t, the blue team's deception is detected if $\log L_t \geq \log U$.

- Therefore, blue team wants to ensure low probability of detection by ensuring

$$\underbrace{\mathbb{P}(\log L_t \geq \log U)}_{\text{Probability of detection}} \leq \underbrace{\epsilon}_{\text{Stealthiness requirement}}, \qquad \forall t \in \{0, 1, \ldots, T\}.$$

Georgia Tech

The probability of detection is provably bounded

- A provable upper bound exists for $\mathbb{P}(\log L_t \geq \log U)$.

[Theorem 1] Probabilistic bound of detection

For any $\epsilon_1, \epsilon_2 > 0$ such that $\epsilon_1 + \epsilon_2 = \log U - \mathbb{E}\log L_t$,

$$\mathbb{P}(\log L_t \geq \log U)$$

$$\leq \exp\left\{-\frac{\epsilon_1^2}{2\left\|(\Sigma_t^w)^{\frac{1}{2}}\left(\mathcal{L}_t^1 + 2\mathcal{L}_t^{2,s} m_t^w\right)\right\|_2^2}\right\} + \exp\left\{-\frac{1}{8}\min\left(\frac{\epsilon_2^2}{\left\|(\Sigma_t^w)^{\frac{1}{2}}\mathcal{L}_t^{2,s}(\Sigma_t^w)^{\frac{1}{2}}\right\|_F^2}, \frac{\epsilon_2}{\left\|(\Sigma_t^w)^{\frac{1}{2}}\mathcal{L}_t^{2,s}(\Sigma_t^w)^{\frac{1}{2}}\right\|_2}\right)\right\}.$$

  - This bound is conservative and may not provide a wide range of admissible λ.

- Practically, $\mathbb{P}(\log L_t \geq \log U)$ can be approximated using Monte-Carlo method.

Based on this bound and approximation, the blue team tunes its stealthiness level with parameter λ.

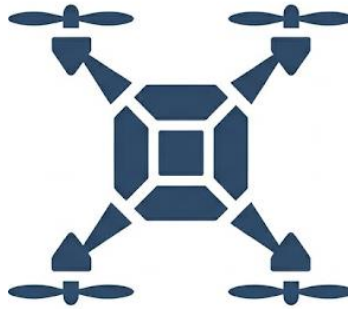
Georgia Tech

The blue team tunes stealthiness with λ

- Blue team tunes stealthiness with λ based on following procedure.

$$\mathbb{P}(\log L_t \geq \log U) \leq \epsilon, \qquad \forall t \in \{0, 1, \ldots, T\}$$

(Probability of detection) (Stealthiness requirement)



1 Select required level of stealthiness ϵ.

2 Search for satisfactory λ :

2-1. Select λ that satisfies well-posedness condition.

2-2. Estimate $\mathbb{P}(\log L_t \geq \log U)$ under selected λ.

2-3. If $\mathbb{P}(\log L_t \geq \log U) \geq \epsilon$ for any t, go back to 2-1.

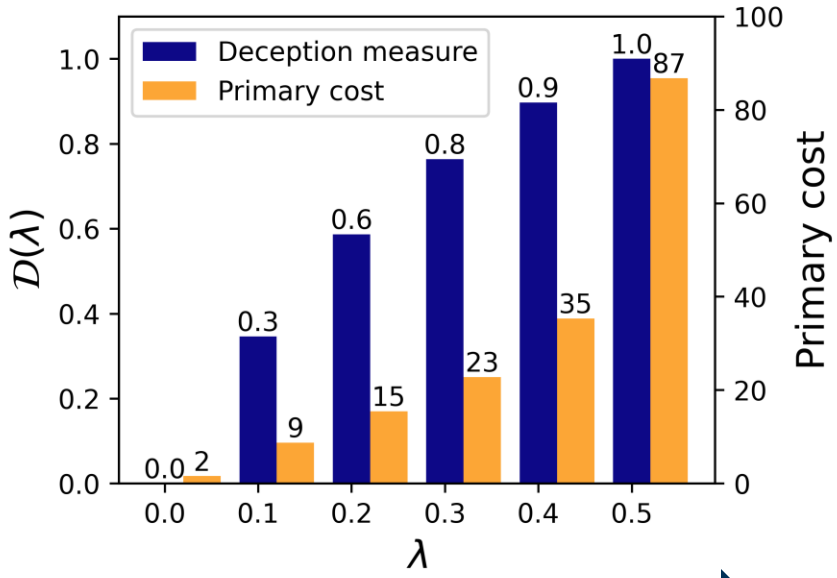
3 Accept λ and execute deceptive LQ control under λ.

Georgia Tech

Larger λ introduces more deception

- Larger λ introduces more deception, but lower performance and higher risk of detection.

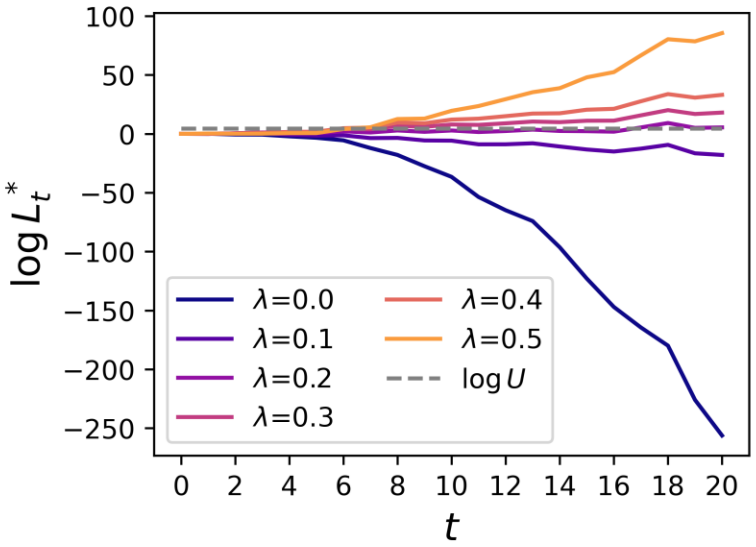
Blue team's control results



More deception introduced

Worse primary performance

Red team's SPRT results



Deception detected!

Question 1: Can we introduce deceptive perturbation to trajectories while pursuing primary objective?

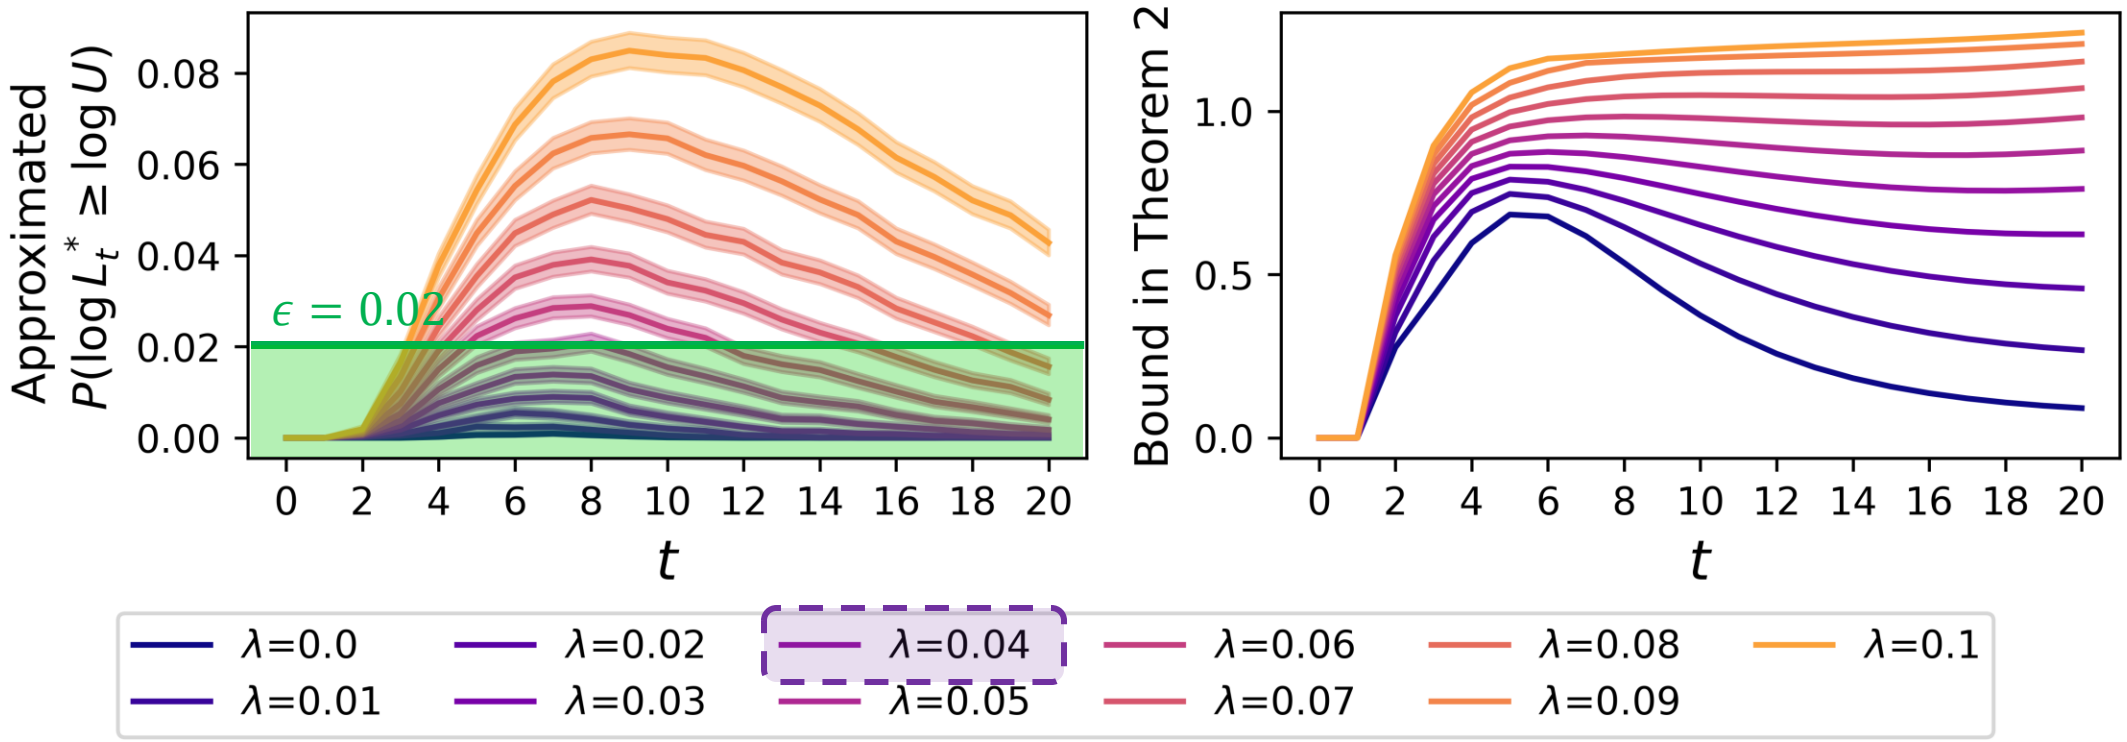
Yes!

- To avoid this, blue team needs to introduce deception up to a certain level.

Georgia Tech

The blue team selects λ based on detection probability

- Blue team selects λ based on detection probability.
  - For example, when $\epsilon = 0.02$ is required, blue team can select $\lambda = 0.04$ based on the approximated probability.

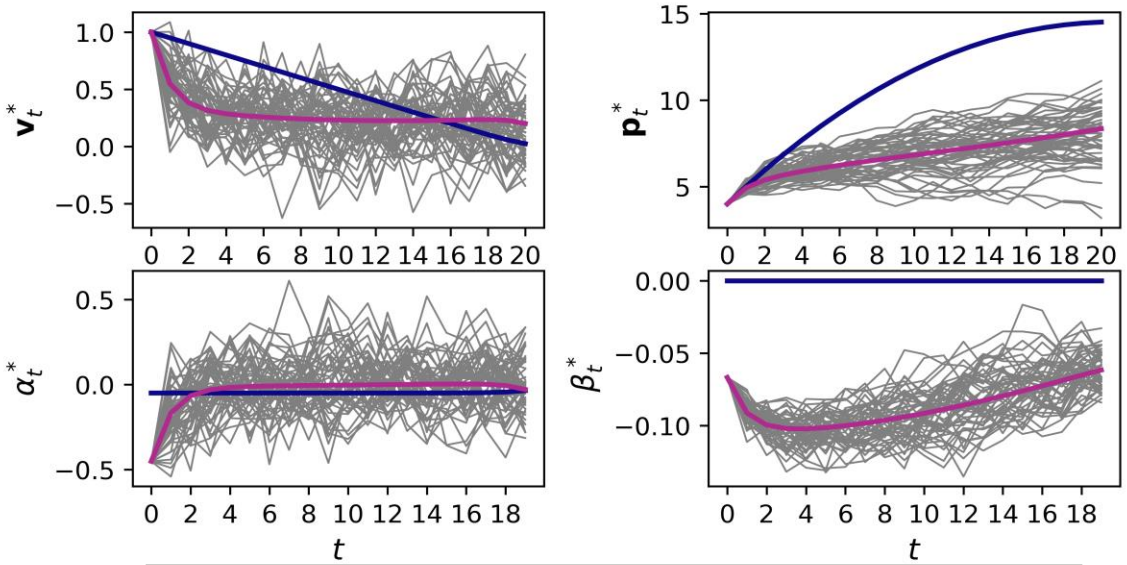


Georgia Tech

Carefully selected λ results in undetected deception

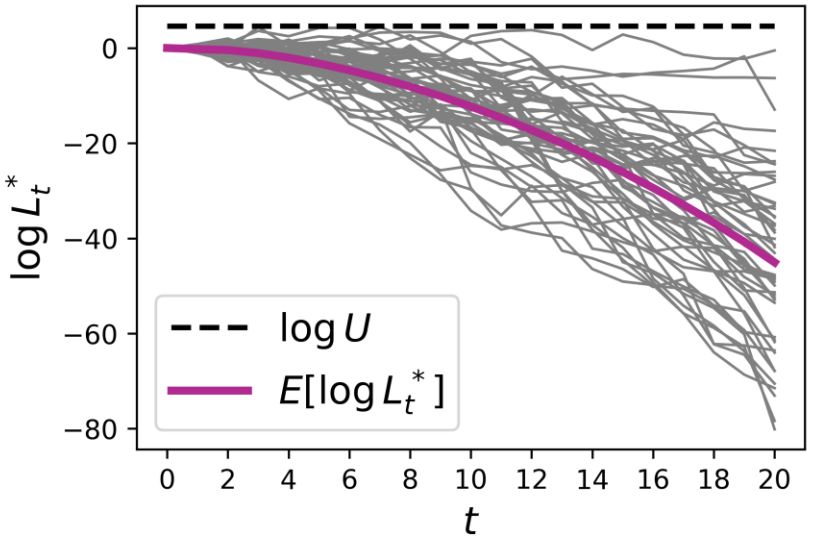
- Carefully selected λ results in undetected deception.
  - $\lambda = 0.04$ is selected to avoid detection with 98% probability.
  - Blue team's deception is conducted with 50 random scenarios.

Blue team's state trajectory (Observation)



The trajectories are perturbed in average.

Red team's SPRT result



Deception is mostly not detected.

Question 2: Even when the attacker knows the presence of deception, can we secretly perturb trajectories without being detected?

Yes!

Georgia Tech

Thank you.

Email: yerinkim@gatech.edu

Georgia Tech