Minimizing Information Leakage of Abrupt Changes in Stochastic Systems

Alessio Russo and Alexandre Proutiere

Control Decision Conference (CDC), 2021

KTH, Royal Institute of Technology, Stockholm

Problem Motivation and Background

Introduction

This work is motivated by current trends in privacy:

- More and more data is being published online.
- Most of the sensors are connected to the internet, perhaps using unencrypted connections.
- Even the window size of a browser can be used to identify someone.

Problem motivation

We study the scenario where an eavesdropper tries to detect a change in a controlled system $\mathcal{S}$.

- Eavesdropping leads to a loss of privacy.
- This privacy loss may reveal private information.
- Eavesdropping is more likely to happen if the system has many sensors.
- **Goal: how can we make the job of the eavesdropper as hard as possible?**

Problem formulation

We consider a Markov Decision Process (MDP) M that undergoes a change at some point ν.

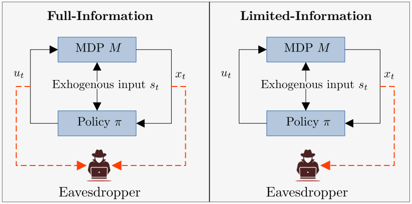

M is described by a tuple $(\mathcal{X}, \mathcal{U}, P, r)$, where $\mathcal{X}$ and $\mathcal{U}$ are the state and action spaces, P is the transition density and r is the reward function.

Problem formulation

We consider a Markov Decision Process (MDP) M that undergoes a change at some point ν.

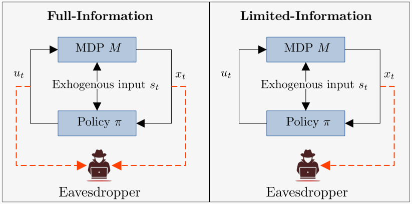

M is described by a tuple $(\mathcal{X}, \mathcal{U}, P, r)$, where $\mathcal{X}$ and $\mathcal{U}$ are the state and action spaces, P is the transition density and r is the reward function.

We focus on single-change problems. We model this change as an exogenous binary input $s_t = \mathbb{1}_{\{t \geq \nu\}}$, so that the transition model is

$$P(x'|x,u,s) = \begin{cases} P_0(x'|x,u) & \text{if } s = 0, \\ P_1(x'|x,u) & \text{if } s = 1 \end{cases}$$

Problem formulation

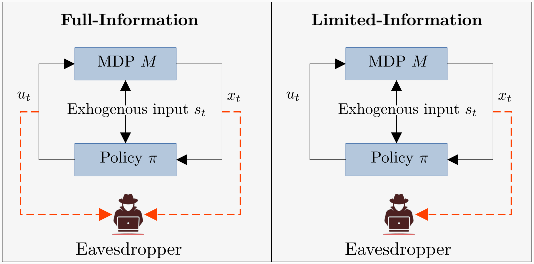

Assumption

- *The victim can observe s_t.*
- *The eavesdropper wishes to infer the change point ν by observing the system's dynamics.*
    - **Full-information***: the eavesdropper can measure (x_t, u_t).*
    - **Limited-Information***: the eavesdropper only measures (x_t).*
- **The goal of the victim is to make the inference of the change point ν as hard as possible**.

Modeling the inference problem

We use minimax Quickest Change Detection theory [3,4] to model the eavesdropper's problem.

There are two fundamental ingredients:

1. **A measure of performance for a detection rule** T [1,2]:

$$\overline{\mathbb{E}}_1(T) := \sup_{\nu \geq 1} \operatorname{ess\,sup} \mathbb{E}_\nu[(T-\nu)^+ | \mathcal{F}_{\nu-1}] \text{ s.t. } \mathbb{E}_\infty[T] \geq \bar{T}$$

Worst case scenario

Detection delay
$(x)^+ = \max(x, 0)$

Expected time to false alarm

Modeling the inference problem

We use minimax Quickest Change Detection theory [3,4] to model the eavesdropper's problem.

There are two fundamental ingredients:

1. **A measure of performance for a detection rule** T [1,2]:

$$\overline{\mathbb{E}}_1(T) := \sup_{\nu \geq 1} \operatorname{ess\,sup} \mathbb{E}_\nu[(T-\nu)^+ | \mathcal{F}_{\nu-1}] \text{ s.t. } \mathbb{E}_\infty[T] \geq \bar{T}$$

Worst case scenario

Detection delay $(x)^+ = \max(x, 0)$

Expected time to false alarm

2. **A lower bound** [2-4]:

$$\liminf_{\bar{T} \to \infty} \inf_{T \in D(\bar{T})} \frac{\overline{\mathbb{E}}_1(T)}{\ln \bar{T}} \geq I^{-1}.$$

Time to false alarm goes to infinity

Set of detection rules

I is the information rate

The idea

The idea is to exploit the lower bound [2]:

$$\liminf_{\bar{T}\to\infty} \inf_{T\in D(\bar{T})} \frac{\overline{\mathbb{E}}_1(T)}{\ln \bar{T}} \geq I^{-1}$$

where $I = \lim_{n\to\infty} n^{-1} \sum_{t=\nu}^{\nu+n} Z_t$, with $Z_i = \ln \frac{f_1(Y_i|Y_1,\ldots,Y_{i-1})}{f_0(Y_i|Y_1,\ldots,Y_{i-1})}$ and Y_i is the i-th observation of the eavesdropper. f_0 indicates the density function before the change (f_1 after the change).

The idea: make the inference problem as hard as possible by minimizing the information rate I.

We also define the privacy level to be $\mathcal{I} = I^{-1}$.

The idea

The idea is to exploit the lower bound [2]:

$$\liminf_{\bar{T}\to\infty} \inf_{T\in D(\bar{T})} \frac{\overline{\mathbb{E}}_1(T)}{\ln \bar{T}} \geq I^{-1}$$

where $I = \lim_{n\to\infty} n^{-1} \sum_{t=\nu}^{\nu+n} Z_t$, with $Z_i = \ln \frac{f_1(Y_i|Y_1,\dots,Y_{i-1})}{f_0(Y_i|Y_1,\dots,Y_{i-1})}$ and Y_i is the i-th observation of the eavesdropper. f_0 indicates the density function before the change (f_1 after the change).

The idea: make the inference problem as hard as possible by minimizing the information rate I.

Differential Privacy: what is the connection with differential privacy?

- We are not interested in minimizing the statistical difference between two trajectories (τ, τ'), but the difference in any trajectory before and after the abrupt change.
- Minimizing I is equivalent to minimizing the on-avg. KL-Privacy [5]

The idea

The idea is to exploit the lower bound [2]:

$$\liminf_{\bar{T}\to\infty} \inf_{T\in D(\bar{T})} \frac{\overline{\mathbb{E}}_1(T)}{\ln \bar{T}} \geq I^{-1}$$

where $I = \lim_{n\to\infty} n^{-1} \sum_{t=\nu}^{\nu+n} Z_t$, with $Z_i = \ln \frac{f_1(Y_i|Y_1,\ldots,Y_{i-1})}{f_0(Y_i|Y_1,\ldots,Y_{i-1})}$ and Y_i is the i-th observation of the eavesdropper. f_0 indicates the density function before the change (f_1 after the change).

Problem: how can we balance the impact on performance?

Use two policies: π_0 used before the change, and π_1 used after the change. Solve the following performance-privacy optimization problem

$$\sup_{\pi_0,\pi_1} \rho V_0^{\pi_0} + (1-\rho)V_1^{\pi_1} - \lambda I(\pi_0,\pi_1),$$

(ρ,λ) tune the performance-privacy trade-off, and $I(\pi_0,\pi_1)$ measures the information rate.

$V_0^{\pi_0}$ **is the average reward of the system controlled by** π_0 (sim. $V_1^{\pi_1}$)

Full-information scenario

Information rate in the full-information case

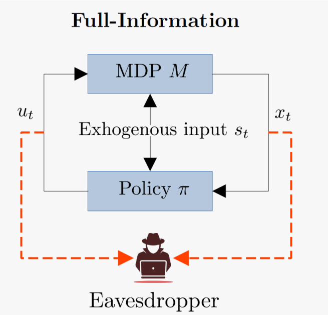

Theorem

In the full-information case (i.e., the eavesdropper measures $Y_t = (X_t, U_t)$), under suitable assumptions of ergodicity we have

$$I = \mathbb{E}_{x\sim\mu_1^{\pi_1}, u\sim\pi_1(x)} \left[D(P_1(x,u), P_0(x,u))\right] + \mathbb{E}_{x\sim\mu_1^{\pi_1}} \left[D(\pi_1(x), \pi_0(x))\right].$$

- $\mu_1^{\pi_1}$ is the stationary measure of the MDP controlled by π_1 after the change
- $D(P,Q)$ is the KL-divergence between P and Q.

Performance-privacy trade-off

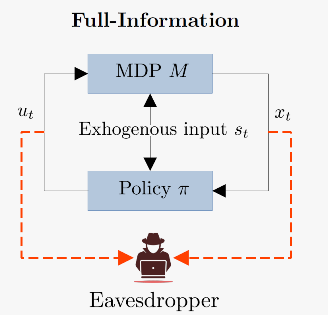

Theorem

In finite state-action spaces solving $\sup_{\pi_0,\pi_1} \rho V_{M_0}^{\pi_0} + (1-\rho)V_{M_1}^{\pi_1} - \lambda I(\pi_0,\pi_1)$ **amounts to solving a concave problem**.

- It can be solved using methods from DC programming (Difference of Convex functions).
- Convex problem if $\pi_1 = \pi_0$ (equivalent to having $\rho = 1$).

Limited-information scenario

Information rate in the Limited-information case

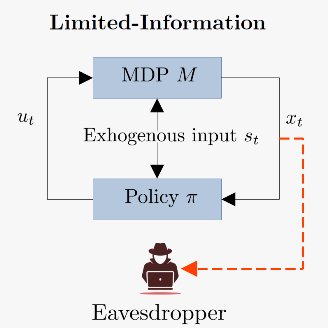

Theorem

In the limited-information case (i.e., the eavesdropper measures $Y_t = (X_t)$), under suitable assumptions of ergodicity we have

$$I = \mathbb{E}_{x\sim\mu_1^{\pi_1}}\left[D\left(P_1^{\pi_1}(x), P_0^{\pi_0}(x)\right)\right].$$

where $P_i^{\pi_i}(x'|x) = \mathbb{E}_{a\sim\pi_i(\cdot|x)}[P_i(x'|x,a)]$.

- I **is smaller compared to the full-information case** (it is an application of the log-sum inequality).
- However, **computing policies that attain the best level of achievable privacy is more challenging** (even computing the minimum value of I is a concave program).
- **Solving** $\sup_{\pi_0,\pi_1} \rho V_{M_0}^{\pi_0} + (1-\rho)V_{M_1}^{\pi_1} - \lambda I(\pi_0,\pi_1)$ **in finite state-action spaces is still a concave problem.**

Examples and numerical results

Linear systems: information rate

Consider a linear system:

$$x_{t+1} = \underbrace{Ax_t + Bu_t}_{\text{Nominal dynamics}} + \underbrace{F\theta s_t}_{\text{Abrupt change}} + \underbrace{w_t}_{\text{White noise}} ,$$

where F and θ are constant terms, $s_t = \mathbb{1}_{\{t \geq \nu\}}$ and $w_t \sim \mathcal{N}(0, Q)$.

Linear systems: information rate

Consider a linear system:

$$x_{t+1} = \underbrace{Ax_t + Bu_t}_{\text{Nominal dynamics}} + \underbrace{F\theta s_t}_{\text{Abrupt change}} + \underbrace{w_t}_{\text{White noise}},$$

where F and θ are constant terms, $s_t = \mathbb{1}_{\{t\geq\nu\}}$ and $w_t \sim \mathcal{N}(0, Q)$.

Proposition

Consider the following policy $u_t = \pi_0(x_t)s_t + \pi_1(x_t)(1 - s_t)$. The lowest possible value of the information rate in the two scenarios is

- **Full information case**

$$\inf_{\pi_i} I(\pi_0, \pi_1) = \frac{1}{2}\theta^\top F^\top Q^{-1} F\theta \Rightarrow \textit{The more noise the better}$$

- **Limited information case**

$$\inf_{\pi_0,\pi_1} I(\pi_0, \pi_1) = \frac{1}{2}\theta^\top F^\top G^\top Q^{-1} GF\theta \Rightarrow \textit{Depends on the inv. of } B$$

where $G = I - B(B^\top QB)^{-1}B^\top Q$.

Linear systems: trade-off - numerical example

Consider $x_{t+1} = \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix} x_t + \begin{bmatrix} 0.01 \\ 1 \end{bmatrix} u_t + \begin{bmatrix} 0.5 \\ 0.7 \end{bmatrix} s_t + w_t$, with $Q = I$.

We study the solution to the performance-privacy problem

$$\sup_{\pi_0,\pi_1} \rho V_0^{\pi_0} + (1-\rho) V_1^{\pi_1} - \lambda I(\pi_0, \pi_1),$$

where $V_i^{\pi_i}$ **is the avg. reward, with reward** $r(x,u) = \|x\|_2^2$. (*we omit the closed form solution for brevity*).

Linear systems: trade-off - numerical example

Consider $x_{t+1} = \begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix} x_t + \begin{bmatrix} 0.01 \\ 1 \end{bmatrix} u_t + \begin{bmatrix} 0.5 \\ 0.7 \end{bmatrix} s_t + w_t$, with $Q = I$.

We study the solution to the performance-privacy problem

$$\sup_{\pi_0, \pi_1} \rho V_0^{\pi_0} + (1-\rho) V_1^{\pi_1} - \lambda I(\pi_0, \pi_1),$$

where $V_i^{\pi_i}$ **is the avg. reward, with reward** $r(x,u) = \|x\|_2^2$. (*we omit the closed form solution for brevity*).

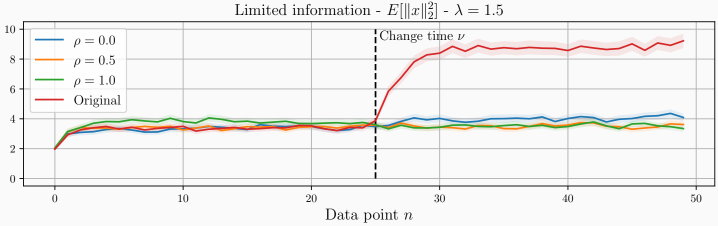

Figure 1: Value of $\mathbb{E}[\|x\|_2^2]$ in the limited information case for $\lambda = 1.5$ and different values of ρ. Shadow area indicates 95% confidence interval.

Linear systems: trade-off - numerical example

$$\sup_{\pi_0,\pi_1} \underbrace{\rho V_0^{\pi_0} + (1-\rho)V_1^{\pi_1}}_{V} - \lambda I(\pi_0,\pi_1),$$

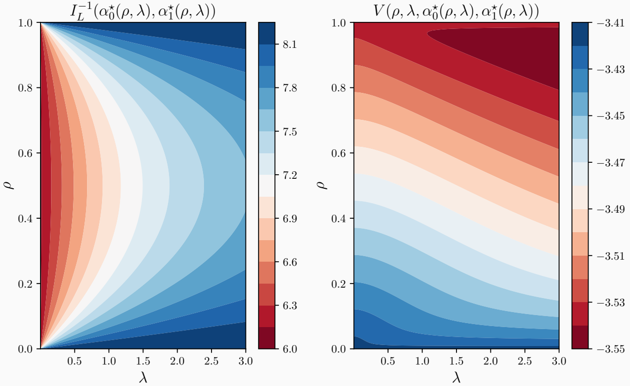

Figure 2: Privacy level I^{-1} (*left*) and Average reward $\rho V_0^{\pi_0} + (1-\rho)V_1^{\pi_1}$ (*right*) as function of ρ and λ.

3-states MDP

Consider an MDP with 3 states and 2 actions. **We analyse the minimum information rate between P_0 and P_θ**, where

$$P_\theta(x'|x,u) = \theta P_0(x'|x,u) + (1-\theta)P_b(x'|x,u), \quad \theta \in [0,1]$$

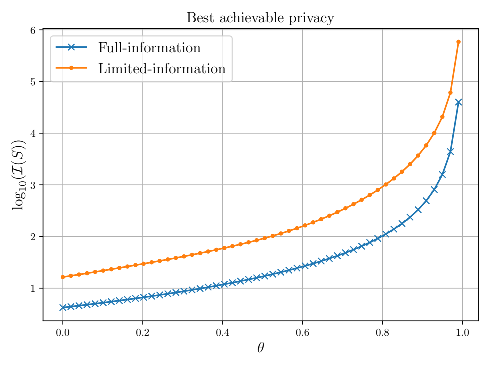

Figure 3: Logarithmic value of I^{-1} as a function of θ

Conclusion

Conclusion & Future work

Conclusion[1]:

- We analysed the problem of making the inference of an abrupt change as hard as possible using the tools from QCD
- Our approach is equivalent to minimizing the On-average KL-Privacy
- For general MDPs the problem is hard to solve, but for linear systems we get nice results
- Future work: consider the learning problem

Thank you for listening!

[1] Code available here `https://github.com/rssalessio/PrivacyStochasticSystems`

References

1. Lorden, Gary. "Procedures for reacting to a change in distribution." The Annals of Mathematical Statistics (1971): 1897-1908.
2. Lai, Tze Leung. "Information bounds and quick detection of parameter changes in stochastic systems." IEEE Transactions on Information Theory 44.7 (1998): 2917-2929.
3. V. V. Veeravalli and T. Banerjee, "Quickest change detection," in Academic Press Library in Signal Processing. Elsevier, 2014, vol. 3, pp. 209–255.
4. A. Tartakovsky, I. Nikiforov, and M. Basseville, Sequential analysis: Hypothesis testing and changepoint detection. CRC Press, 2014.
5. Wang, Yu-Xiang, Jing Lei, and Stephen E. Fienberg. "On-average kl-privacy and its equivalence to generalization for max-entropy mechanisms." International Conference on Privacy in Statistical Databases. Springer, Cham, 2016.