Leveraging Unique CPS Properties to Design Better Privacy-Enhancing Algorithms

Jairo Giraldo
Alvaro Cardenas
Murat Kantarcioglu

Erik Jonsson School of Engineering
and Computer Science

Cyber-Physical Systems

Integration of a physical process with embedded computation and communication networks that can make the system safer, more efficient, and smarter.

Cyber-Physical Systems

Integration of a physical process with embedded computation and communication networks that can make the system safer, more efficient, and smarter.

Integration of a physical process with embedded computation and communication networks that can make the system safer, more efficient, and smarter.

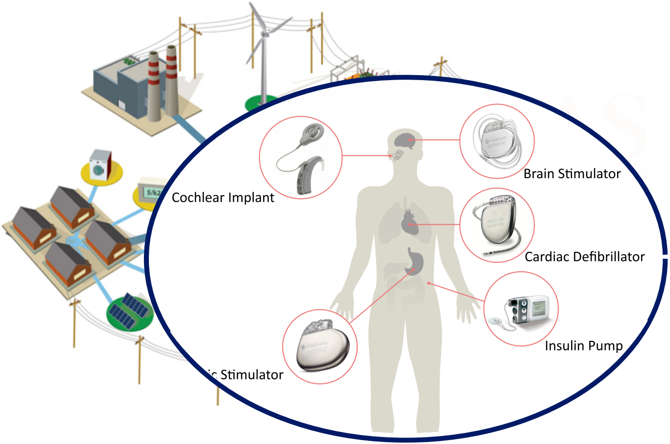

Cyber-Physical Systems

There is an incentive for most CPS to gather sensitive information. Unfortunately, that information can be used by adversaries. For example...

[1] G.W. Hart, Nonintrusive Applicance Load Monitoring, *Proceedings of the IEEE*, 80 (12):1870-1891, 1992.

Cyber-Physical Systems

There is an incentive for most CPS to gather sensitive information. Unfortunately, that information can be used by adversaries. For example...
It is possible to infer the behavior of electricity users by analyzing their consumption patterns [1].

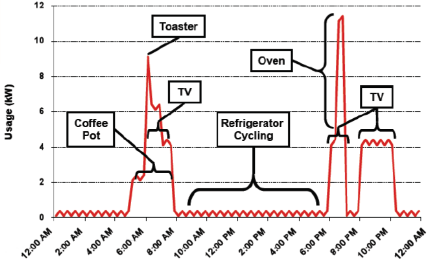

[1] G.W. Hart, Nonintrusive Applicance Load Monitoring, *Proceedings of the IEEE*, 80 (12):1870-1891, 1992.

Cyber-Physical Systems

There is an incentive for most CPS to gather sensitive information. Unfortunately, that information can be used by adversaries. For example...

It is possible to infer the behavior of electricity users by analyzing their consumption patterns [1].

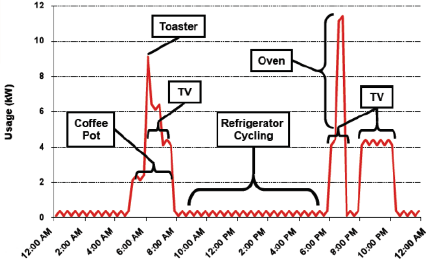

Privacy algorithms are relevant in CPS!!!

[1] G.W. Hart, Nonintrusive Applicance Load Monitoring, *Proceedings of the IEEE*, 80 (12):1870-1891, 1992.

CPS Posses Unique Properties

- CPS are noisy (e.g., sensor noise, environmental disturbances).
- Feedback loops and Controls can attenuate/amplify noise.
- Some systems are very susceptible to noise (stability).

Our Contributions

Can be divided in two groups:

- **Differential Privacy Framework**
  - Using tools from stochastic control theory, we characterize the inherent noise.
  - We define ***Inherent Differential Privacy***
  - Find the minimum external noise that should be injected to ensure a desired level of privacy.
- **Data Minimization in Multi-agent Control Systems**
  - We propose event-based privacy.
  - We modify the sensor sampling period to hide relevant information.

Differential Privacy for Databases

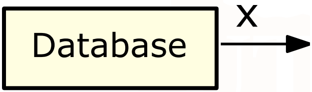

Name	Income
Penny	50
Leonard	30
Howard	25
Raj	40
Sheldon	35

Differential Privacy for Databases

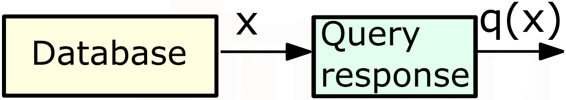

Name	Income
Penny	50
Leonard	30
Howard	25
Raj	40
Sheldon	35

Average 36

Differential Privacy for Databases

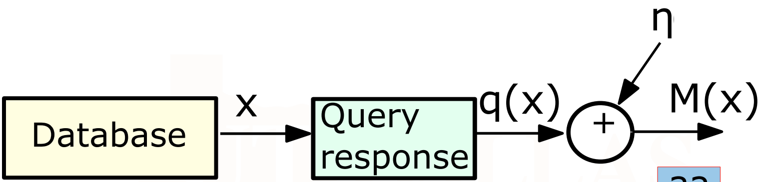

Name	Income
Penny	50
Leonard	30
Howard	25
Raj	40
Sheldon	35

Average 36

32
39
35
37

Differential Privacy for Databases

[2] Now consider two adjacent databases x, x' that differ **only in one element.** For all pairs x, x', $\epsilon \in (0, 1)$, $\delta > 0$, and $S \subseteq range(M)$

(ϵ, δ)-Differential Privacy

$$P(M(x) \in S) \leq e^{\epsilon} P(M(x') \in S) + \delta$$

[2] C. Dwork et al., The algorithmic foundations of differential privacy, *in Foundations and Trends in Theoretical Computer Science*, pp. 211-407, 2014

[2] Now consider two adjacent databases x, x' that differ **only in one element.** For all pairs x, x', $\epsilon \in (0,1)$, $\delta > 0$, and $S \subseteq range(M)$

(ϵ, δ)-Differential Privacy

$$P(M(x) \in S) \leq e^{\epsilon} P(M(x') \in S) + \delta$$

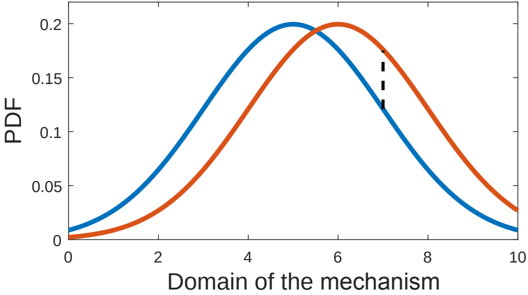

[2] C. Dwork et al., The algorithmic foundations of differential privacy, *in Foundations and Trends in Theoretical Computer Science*, pp. 211-407, 2014

How much noise to add?

Depends on the maximum change of the query response when only one element of the database is modified.

Sensitivity

$$\Delta_{q,p} = \max_{x,x'} \|q(x) - q(x')\|_p$$

How much noise to add?

Depends on the maximum change of the query response when only one element of the database is modified.

Sensitivity

$\Delta_{q,p} = \max_{x,x'} \|q(x) - q(x')\|_p$

Gaussian Mechanism

Noise from a Gaussian distribution, $\eta \sim N(0, \sigma^2)$. If

$$\sigma \geq \frac{\sqrt{2\ln(1.25/\delta)}\Delta_{q,2}}{\epsilon},$$

(ϵ, δ)-differential privacy is guaranteed.

How much noise to add?

Depends on the maximum change of the query response when only one element of the database is modified.

Sensitivity

$\Delta_{q,p} = \max_{x,x'} \|q(x) - q(x')\|_p$

Gaussian Mechanism

Noise from a Gaussian distribution, $\eta \sim N(0, \sigma^2)$. If

$$\boxed{\sigma \geq \frac{\sqrt{2\ln(1.25/\delta)}\Delta_{q,2}}{\epsilon}},$$

(ϵ, δ)-differential privacy is guaranteed.

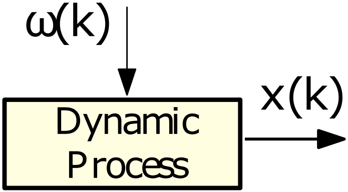

$$x(k+1)=Ax(k)+Bu(k)+\omega(k)$$

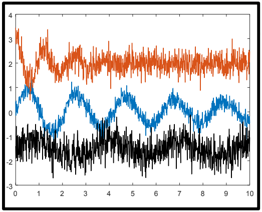

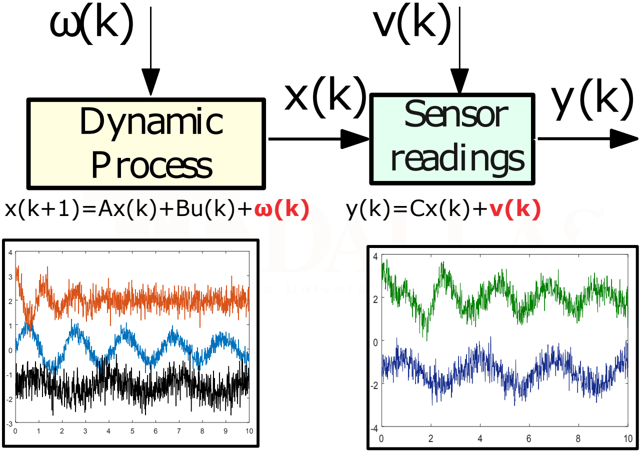
ω(k)
v(k)
x(k)
y(k)
Dynamic Process
Sensor readings
x(k+1)=Ax(k)+Bu(k)+ω(k)
y(k)=Cx(k)+v(k)

General Model of CPS

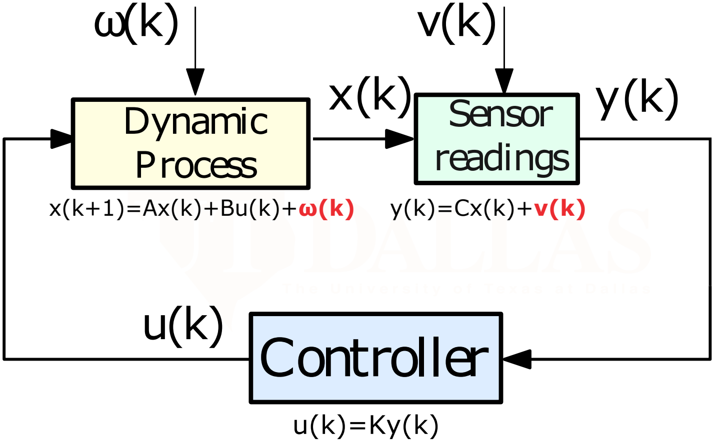

Properties of CPS with feedback

- There are inherent sources of uncertainties:
  - $\omega(k)$ represents environmental disturbances or random changes in the process
  - $v(k)$ describes the sensor noise
- There is an incentive to share $y(k)$, and keep it private
- **The output $y(k)$ is already noisy, and its variance evolves over time**

Properties of CPS with feedback

- There are inherent sources of uncertainties:
  - $\omega(k)$ represents environmental disturbances or random changes in the process
  - $v(k)$ describes the sensor noise
- There is an incentive to share $y(k)$, and keep it private
- **The output $y(k)$ is already noisy, and its variance evolves over time**

Knowing how noisy $y(k)$ is, we can characterize the level of privacy!!

How to characterize the variance of $y(k)$?

CPS Model

$$
\begin{aligned}
x(k+1) &= Ax(k) + Bu(k) + \omega(k) \\
y(k) &= Cx(k) + v(k) \\
u(k) &= Ky(k)
\end{aligned}
$$

$\omega_i \sim N(0, \sigma^2_{\omega,i})$ with $R_\omega = \mathrm{diag}(\sigma^2_{\omega,1}, \dots, \sigma^2_{\omega,n})$. Similarly, $v_i(k) \sim N(0, \sigma^2_{v,i})$ and R_v.
Defining $\bar{A} = A + BKC$,

$$\mathbf{x}(k+1) = \bar{A}\mathbf{x}(k) + \underbrace{BK\mathbf{v}(k) + \omega(k)}_{\varphi(k)}$$

$$R_\varphi = E[\varphi(k)\varphi(k)^\top] = BKR_vK^\top B^\top + R_\omega.$$

How to characterize the variance of $y(k)$?

Let $m(k) = E[x(k)]$. From stochastic control theory [3], the covariance matrix of the states is defined by $Q(k) = E[(x(k) - m(k))(x(k) - m(k))^\top]$ and it evolves according to

$$Q(k+1) = \bar{A}Q(k)\bar{A}^\top + R_\varphi.$$

[3] G. Chen, et al., Linear stochastic control systems, *CRC Press*, 1995

How to characterize the variance of $y(k)$?

Let $m(k) = E[x(k)]$. From stochastic control theory [3], the covariance matrix of the states is defined by
$Q(k) = E[(x(k) - m(k))(x(k) - m(k))^\top]$ and it evolves according to

$$Q(k+1) = \bar{A}Q(k)\bar{A}^\top + R_\varphi.$$

The variance of the output vector $y(k)$ at each instant k is

$$\boxed{Q_y(k) = CQ(k)C^\top + R_v}$$

and depends on the system and control parameters.

The variance of each output $y_i(k)$ is $\sigma^2_{y,i}(k)$ and correspond to the diagonal elements of $Q_y(k)$.

[3] G. Chen, et al., Linear stochastic control systems, *CRC Press*, 1995

How much privacy does $y(k)$ guarantees?

For a given δ, and sensitivity $\Delta_{y,2}$, the **inherent level of privacy** (or inherent privacy loss) is then

$$\epsilon_y(k) = \frac{\sqrt{2\ln(1.25/\delta)}\Delta_{y,2}}{\min_i \sigma_{y,i}(k)}.$$

How much privacy does $y(k)$ guarantees?

For a given δ, and sensitivity $\Delta_{y,2}$, the **inherent level of privacy** (or inherent privacy loss) is then

$$\epsilon_y(k) = \frac{\sqrt{2\ln(1.25/\delta)}\Delta_{y,2}}{\min_i \sigma_{y,i}(k)}.$$

$(\epsilon_y(k), \delta)$-differential privacy is guaranteed without adding any external mechanism.

How to ensure a desired level of (ϵ, δ)-differential privacy?

Recall that for a desired ϵ, δ, the standard deviation of the output noise should be

$$\sigma \geq \sqrt{2\ln(1.25/\delta)}\Delta_{y,2}/\epsilon\,.$$

How to ensure a desired level of (ϵ, δ)-differential privacy?

Recall that for a desired ϵ, δ, the standard deviation of the output noise should be

$$\sigma \geq \sqrt{2\ln(1.25/\delta)}\Delta_{y,2}/\epsilon .$$

If $\min_i \sigma_{y,i}(k) < \sigma$, extra noise $\eta(k)$ should be added.

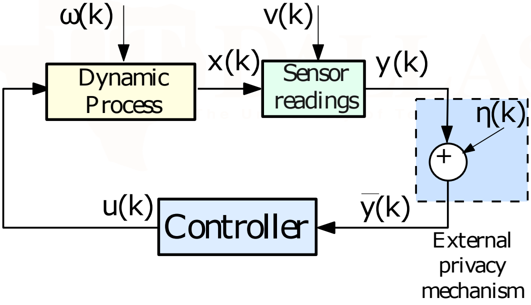

Injecting Minimum Noise

If we don't consider the inherent noise, the variance of $\eta(k)$ would be

$$R_\eta = \sigma^2 I_N$$

Injecting Minimum Noise

If we don't consider the inherent noise, the variance of $\eta(k)$ would be

$$R_\eta = \sigma^2 I_N$$

However, the **minimum noise** $\eta(k)$ has a variance that evolves over time and depends on the inherent noise,

$$\boxed{R_\eta(k) = \sigma^2 I_N - CQ(k)C^\top - R_v}$$

Clearly, since $CQ(k)C^\top + R_v > 0$, less noise is injected.

Case Study: Real-Time Pricing

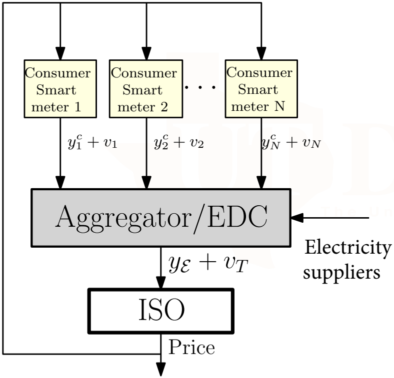

- Consumers
- Electricity suppliers
- EDC (Energy Data Center) gathers information.
- ISO (Independent system operator) takes the aggregated and set the price $\lambda(k)$

Case Study: Real-Time Pricing

- RTP can be modeled as a linear system of the form $x(k+1) = Ax(k) + B\lambda(k)$ [4]
- $y_{\varepsilon}(k) = y_T^s(k) - y_T^c(k)$ is the supply-demand mismatch received by the ISO
- The controller objective is to drive the supply-demand mismatch to zero

[4] R. Tan et al., Impact of integrity attacks on real-time pricing in smart grids, *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security, pp. 439-450, 2013*

Case Study: Real-Time Pricing

- RTP can be modeled as a linear system of the form $x(k+1) = Ax(k) + B\lambda(k)$ [4]
- $y_\varepsilon(k) = y_T^s(k) - y_T^c(k)$ is the supply-demand mismatch received by the ISO
- The controller objective is to drive the supply-demand mismatch to zero

Setting the Price

The control strategy that sets the price is

$$\lambda(k+1) = \lambda(k) + Ky_\varepsilon(k)$$

[4] R. Tan et al., Impact of integrity attacks on real-time pricing in smart grids, *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security, pp. 439-450, 2013*

Inherent Privacy with No external mechanism

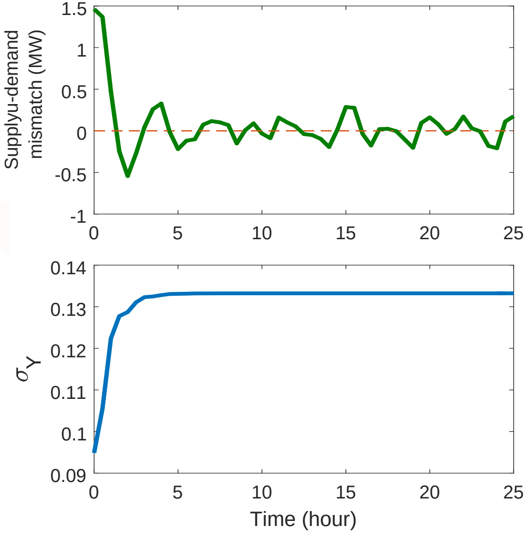

Inherent Privacy with No external mechanism

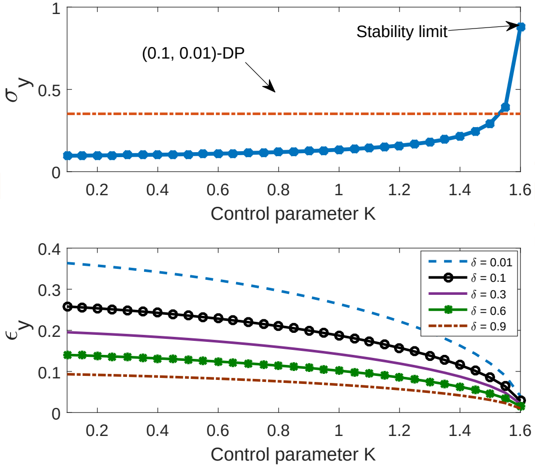

Adding the minimum noise

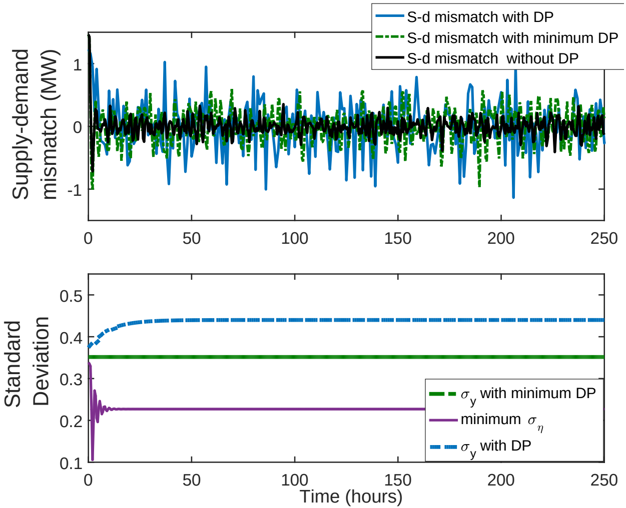

Some Systems Cannot Tolerate Noise

Distributed Frequency Control in the Smart Grid with a DP mechanism.

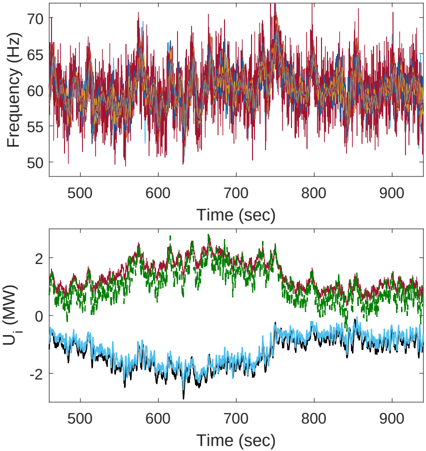

Event-based privacy aims to keep private specific events in the system.

For instance, changes in the power consumption.

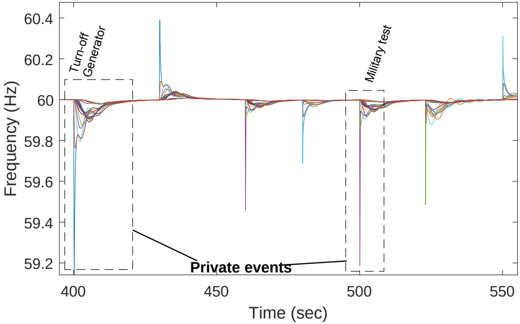

How to provide event-based privacy?

We propose two approaches:

We propose two approaches:

Periodic Sampling

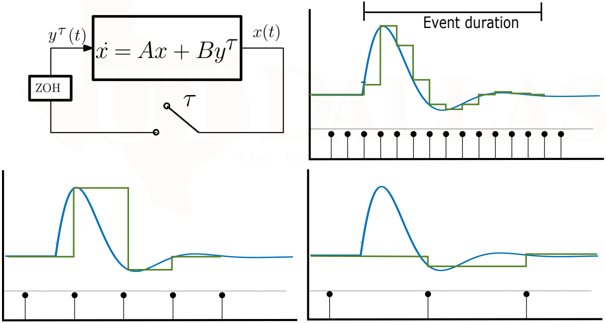

It does not require much knowledge about the events.

Discretionary Sampling

- We select when to sample and when to lie.
- We lie by sending old information ($y(k) = y(k-1)$) for some sampling periods.
- It requires prior knowledge of the events and their duration.
- This ensures complete privacy, but it increases the settling time.

Case Study: Control in the Smart Grid

Distributed frequency control for the IEEE 30 bus system benchmark with distributed generation [5]

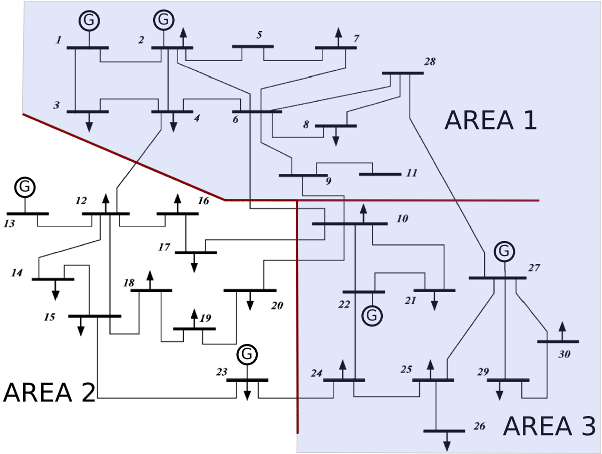

[5] W. El-Khattam et al., Investigating distributed generation systems performance using monte carlo simulation, *IEEE Transactions on Power Systems*, pp. 524–532, 2006.

Periodic sampling vs. Discretionary sampling

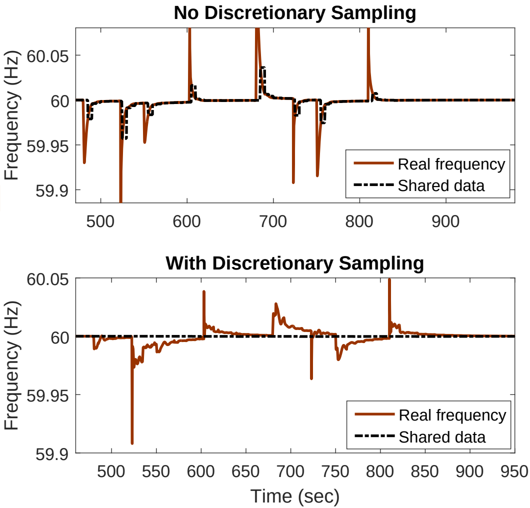

Conclusions

- It is possible to use tools from control theory to analyze privacy in CPS.
- Inherent uncertainties in CPS can be amplified/attenuated to provide certain levels of differential privacy
- Considering the inherent noise, we can minimize the amount of noise to be injected.
- It is possible to hide events by changing the amount of information transmitted, but it causes performance degradation.