Differential Privacy in Cyber-Physical Systems

Jerome Le Ny and George J. Pappas
Joseph Moore Professor
Departments of ESE and CIS
University of Pennsylvania
pappasg@seas.upenn.edu

The trillion sensor challenge

- 5 Billion people to be connected by 2015 (Source: NSN)
- 7 trillion wireless devices serving 7 billion people in 2017 (Source: WWRF)
  - 1000 wireless devices per person?

[Courtesy: Niko Kiukkonen, Nokia]

TerraSwarms: Swarm at the edge of the cloud

TerraSwarm

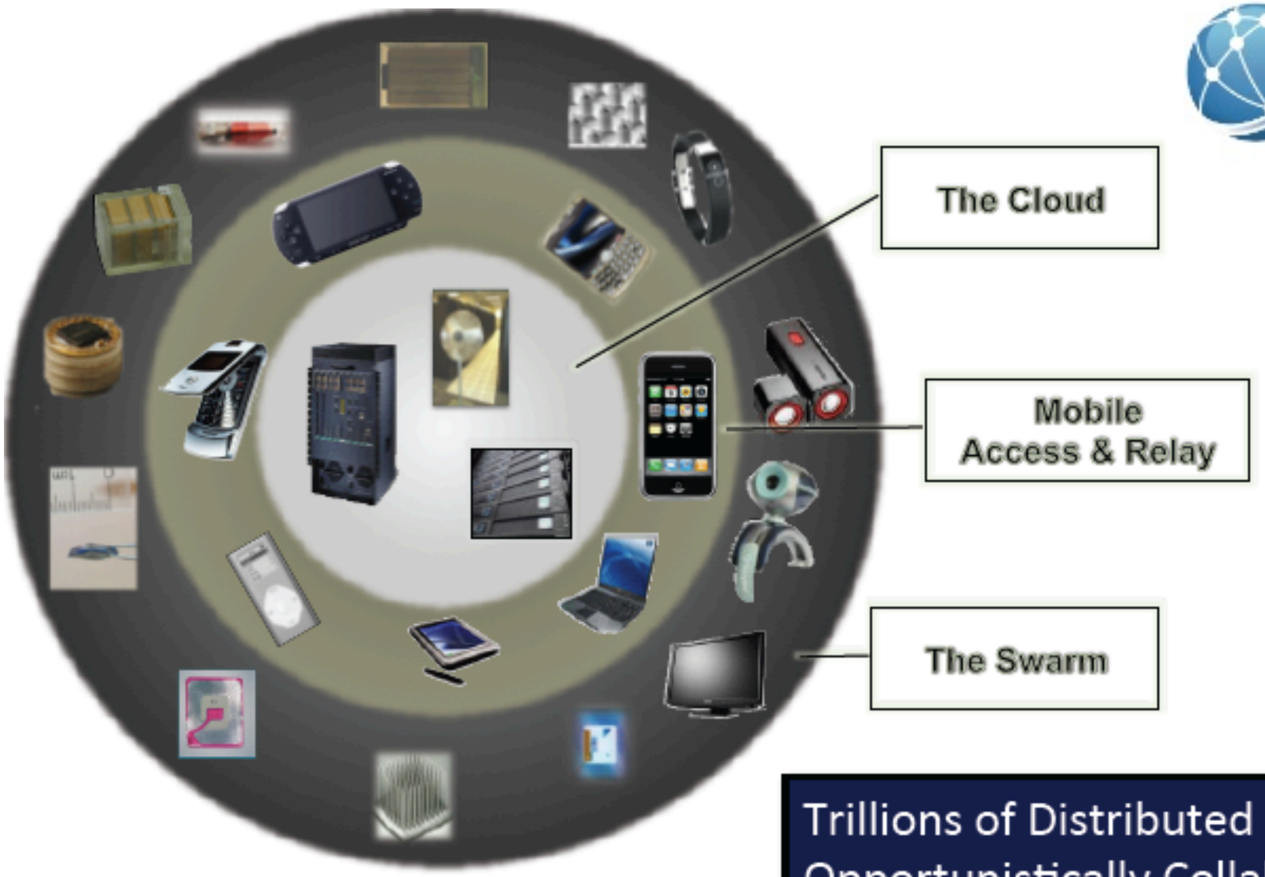

The Cloud

Mobile
Access & Relay

The Swarm

[J. Rabaey, ASPDAC'08]

Trillions of Distributed Connected Devices
Opportunistically Collaborating to Present
Unique Experiences or Fulfill Common Goals

TerraSwarm Challenges

Penn
University *of* Pennsylvania

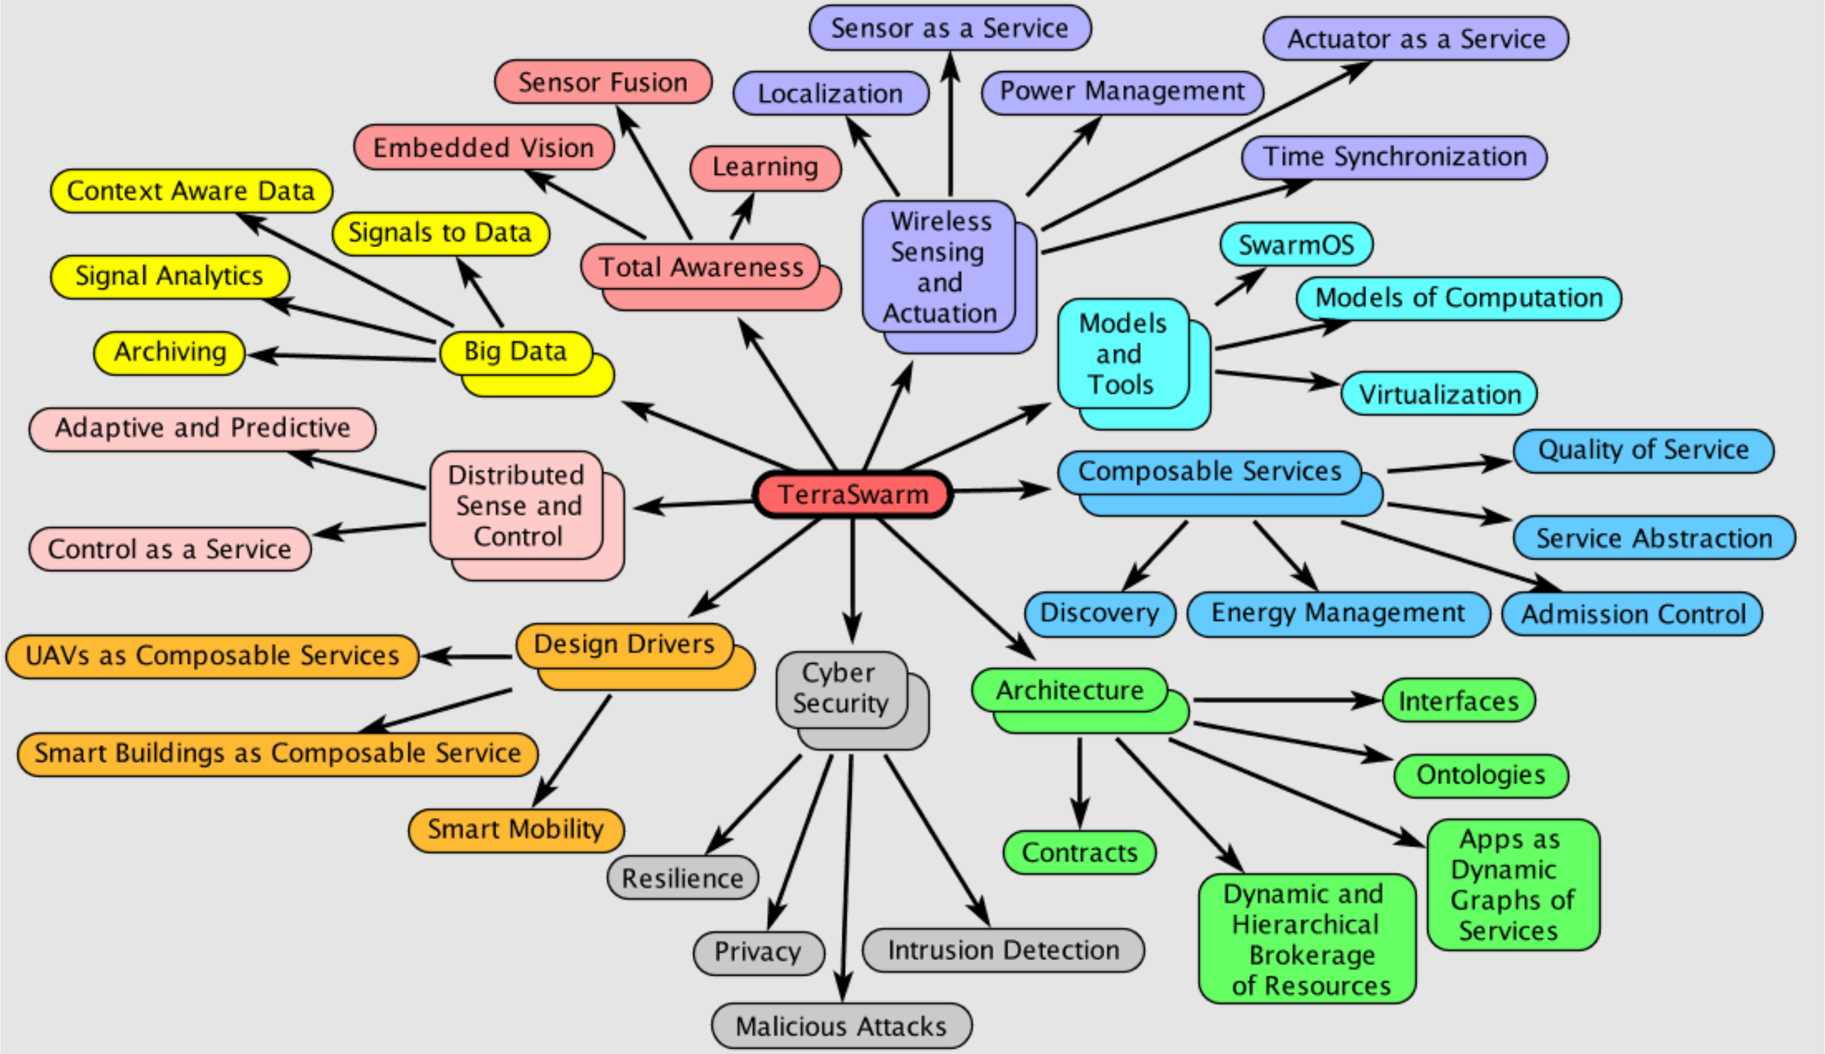

Privacy Concerns - Cyber-physical applications

[image: Inst. Mech. Eng.]

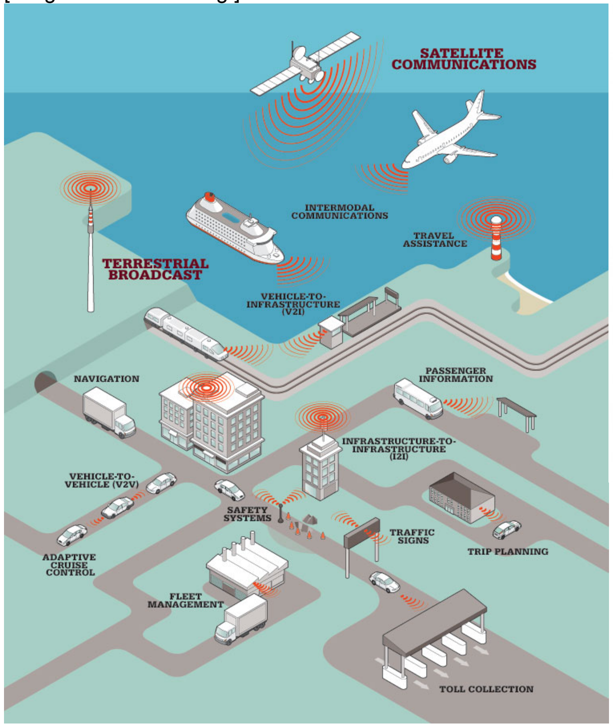

Intelligent Transportation Systems

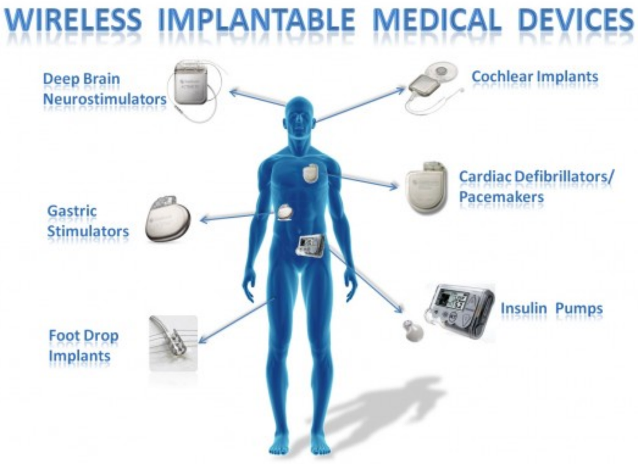
WIRELESS IMPLANTABLE MEDICAL DEVICES
Deep Brain Neurostimulators
Cochlear Implants
Gastric Stimulators
Cardiac Defibrillators/ Pacemakers
Foot Drop Implants
Insulin Pumps

[UCR]

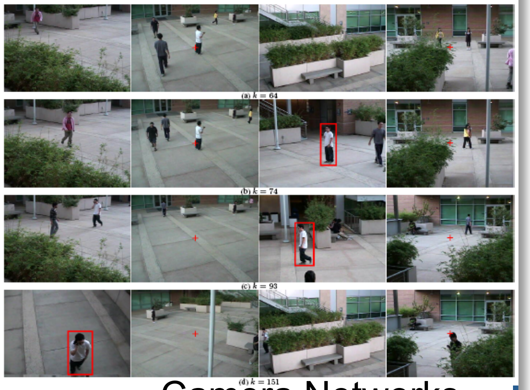

Camera Networks

Location based services

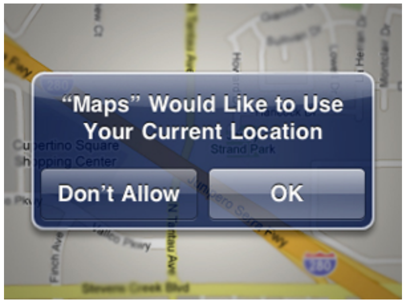

“Maps” Would Like to Use Your Current Location

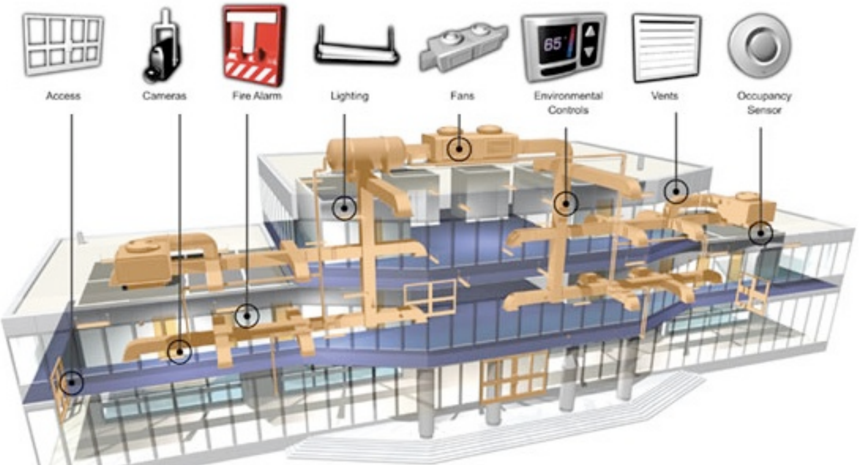

Building Automation

Privacy Concerns - SmartGrid

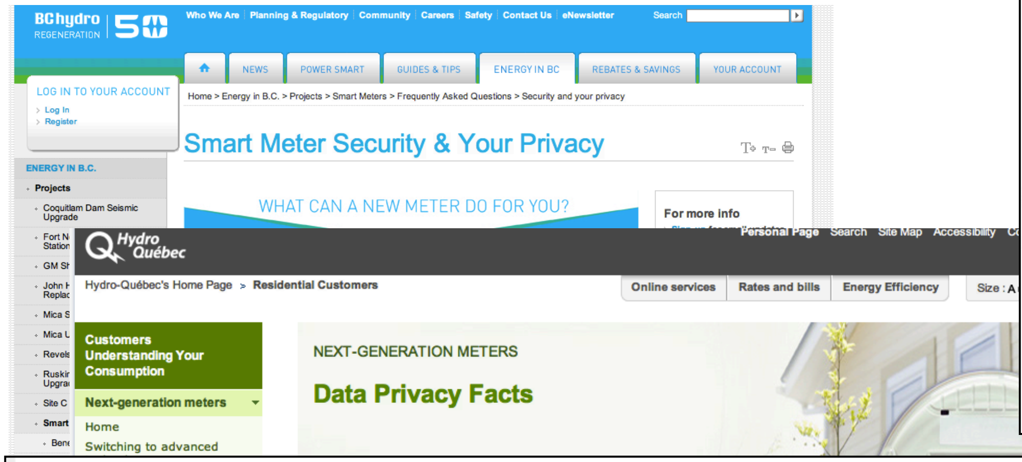

Home > Energy in B.C. > Projects > Smart Meters > Frequently Asked Questions > Security and your privacy

Smart Meter Security & Your Privacy

WHAT CAN A NEW METER DO FOR YOU?

For more info

Hydro-Québec's Home Page > Residential Customers

NEXT-GENERATION METERS

Data Privacy Facts

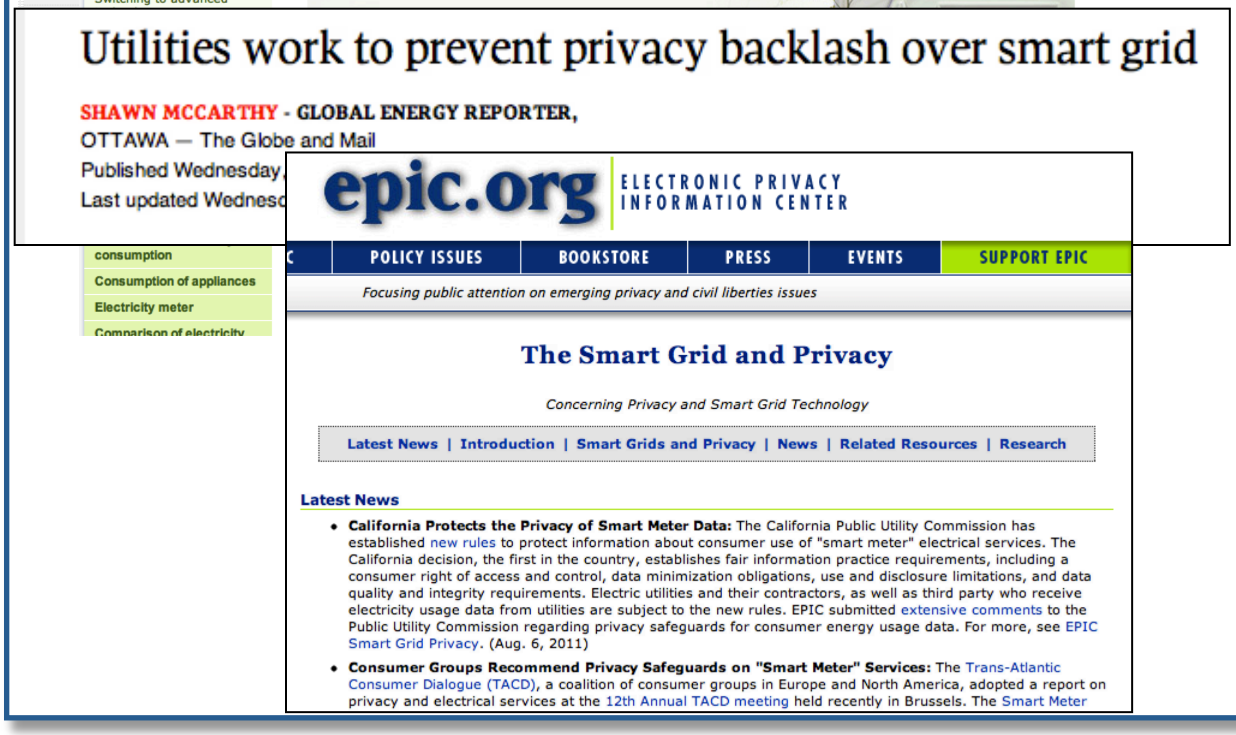

Utilities work to prevent privacy backlash over smart grid

SHAWN MCCARTHY - GLOBAL ENERGY REPORTER,
OTTAWA — The Globe and Mail
Published Wednesday,
Last updated Wednesd

epic.org ELECTRONIC PRIVACY INFORMATION CENTER

Focusing public attention on emerging privacy and civil liberties issues

The Smart Grid and Privacy

Concerning Privacy and Smart Grid Technology

Latest News | Introduction | Smart Grids and Privacy | News | Related Resources | Research

Latest News

- **California Protects the Privacy of Smart Meter Data:** The California Public Utility Commission has established new rules to protect information about consumer use of "smart meter" electrical services. The California decision, the first in the country, establishes fair information practice requirements, including a consumer right of access and control, data minimization obligations, use and disclosure limitations, and data quality and integrity requirements. Electric utilities and their contractors, as well as third party who receive electricity usage data from utilities are subject to the new rules. EPIC submitted extensive comments to the Public Utility Commission regarding privacy safeguards for consumer energy usage data. For more, see EPIC Smart Grid Privacy. (Aug. 6, 2011)
- **Consumer Groups Recommend Privacy Safeguards on "Smart Meter" Services:** The Trans-Atlantic Consumer Dialogue (TACD), a coalition of consumer groups in Europe and North America, adopted a report on privacy and electrical services at the 12th Annual TACD meeting held recently in Brussels. The Smart Meter

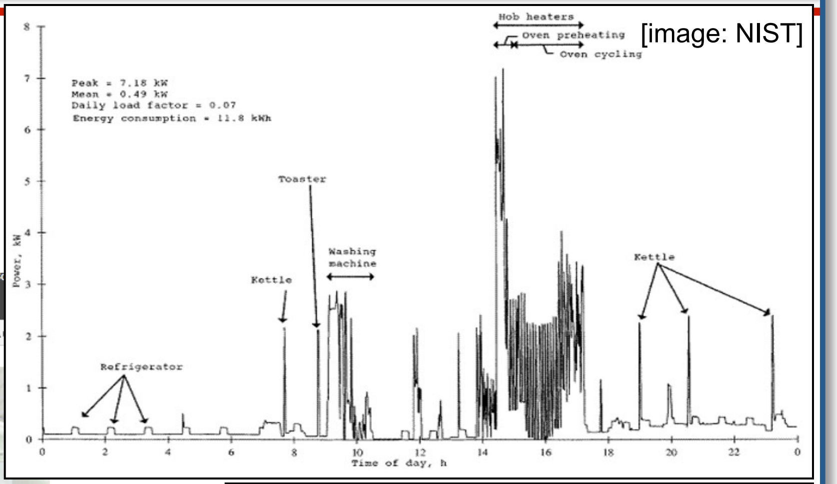

[image: NIST]

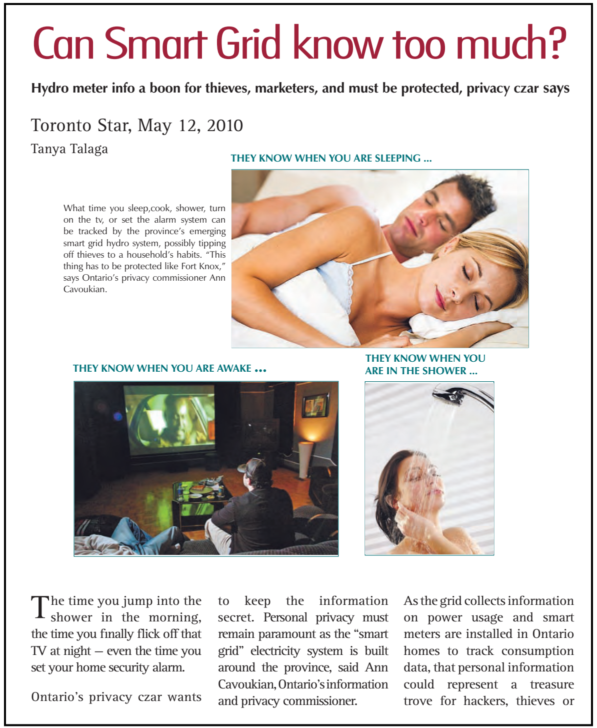

Can Smart Grid know too much?

Hydro meter info a boon for thieves, marketers, and must be protected, privacy czar says

Toronto Star, May 12, 2010

Tanya Talaga

THEY KNOW WHEN YOU ARE SLEEPING ...

What time you sleep,cook, shower, turn on the tv, or set the alarm system can be tracked by the province's emerging smart grid hydro system, possibly tipping off thieves to a household's habits. "This thing has to be protected like Fort Knox," says Ontario's privacy commissioner Ann Cavoukian.

THEY KNOW WHEN YOU ARE AWAKE ...

THEY KNOW WHEN YOU ARE IN THE SHOWER ...

The time you jump into the shower in the morning, the time you finally flick off that TV at night – even the time you set your home security alarm.

Ontario's privacy czar wants to keep the information secret. Personal privacy must remain paramount as the "smart grid" electricity system is built around the province, said Ann Cavoukian, Ontario's information and privacy commissioner.

As the grid collects information on power usage and smart meters are installed in Ontario homes to track consumption data, that personal information could represent a treasure trove for hackers, thieves or

Privacy Concerns – Mobile Millenium

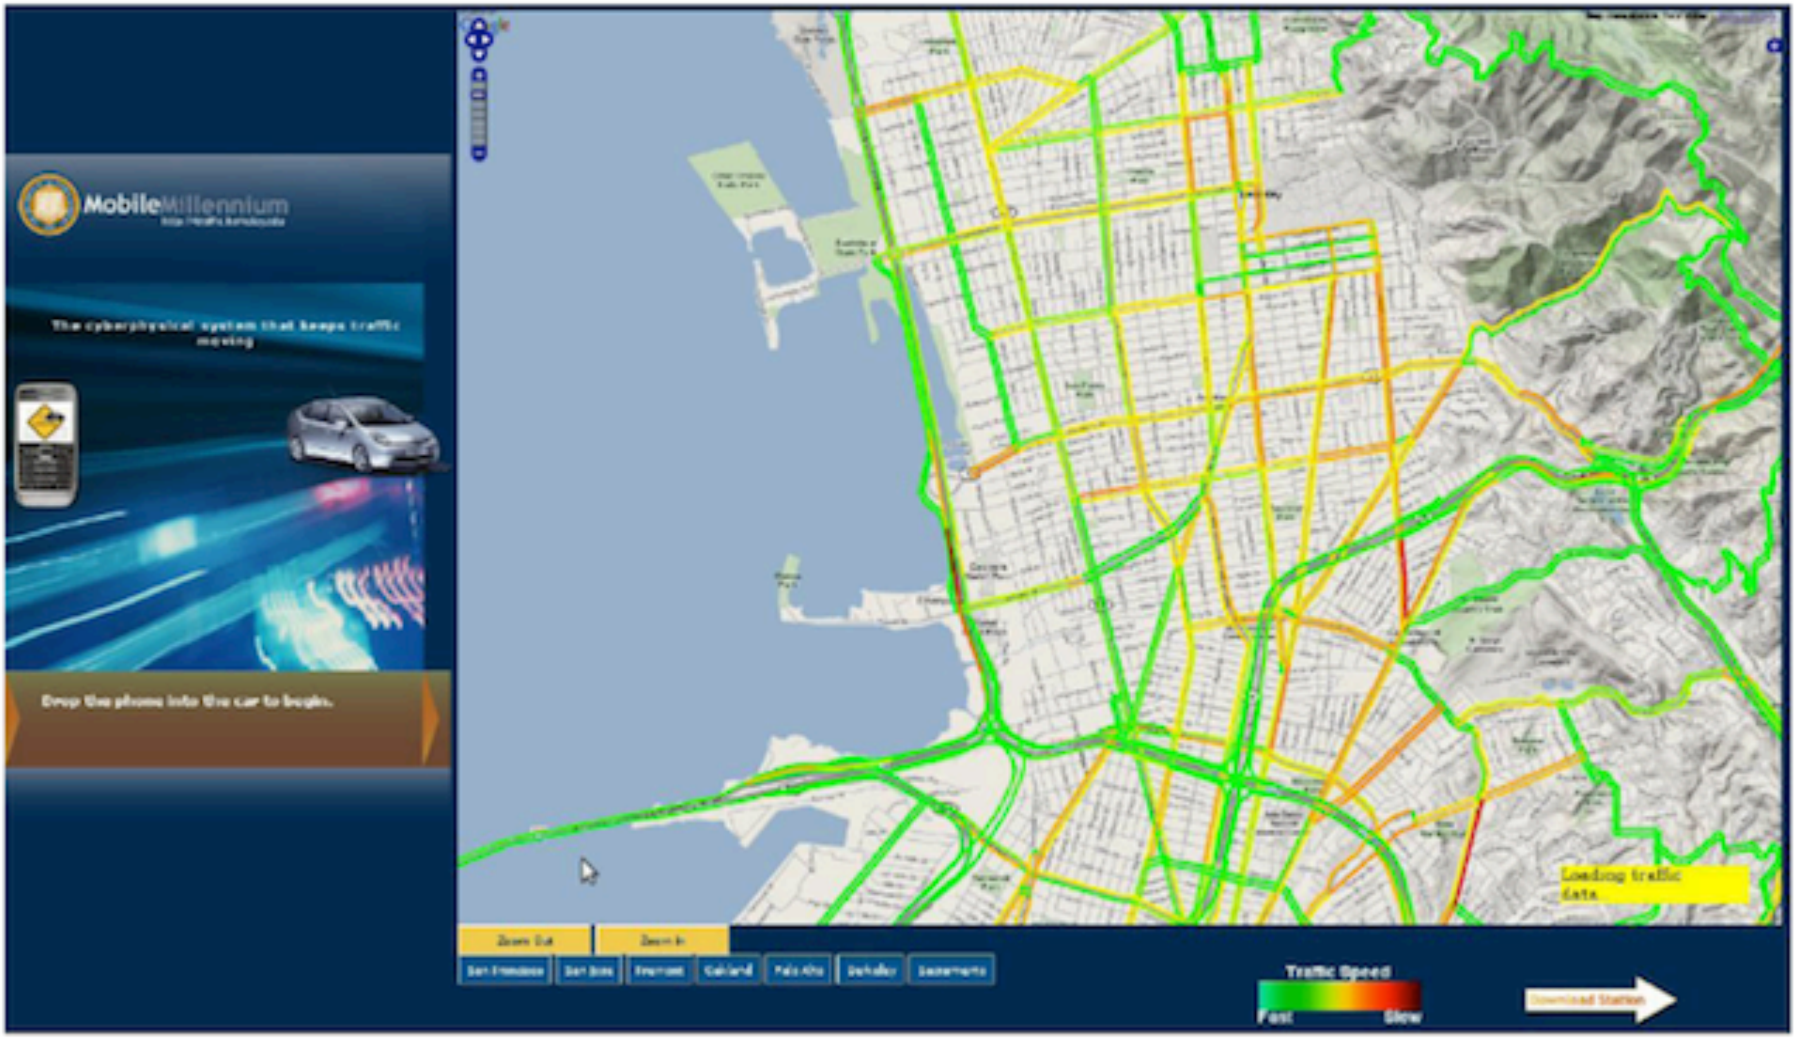

The cyberphysical system that keeps traffic moving

Drop the phone into the car to begin.

Loading traffic data

Privacy challenges

- What is privacy, formally?
- Is there a tradeoff between privacy and utility?
- Privacy-aware estimation and control
- Systems and control tools for privacy

Privacy is not anonymity

- Privacy breaches generally due to existence of side information
  - Mass. GIC medical db w/ voter registration db (Sweeney, 1997)
  - Netflix prize w/ IMDB (Narayana & Shmatikov, 2008)
  - Individual online transactions w/ changes in public recommendation systems (Calandrino et al., 2011)
  - Anonymity in location based services
- Can't know what the adversary knows, or might know in the future.

Differential Privacy

Cynthia Dwork. "Differential privacy." Automata, languages and programming. Springer Berlin Heidelberg, 2006. 1-12.

Differential Privacy

Cynthia Dwork

Microsoft Research
dwork@microsoft.com

Abstract. In 1977 Dalenius articulated a desideratum for statistical databases: nothing about an individual should be learnable from the database that cannot be learned without access to the database. We give a general impossibility result showing that a formalization of Dalenius' goal along the lines of semantic security cannot be achieved. Contrary to intuition, a variant of the result threatens the privacy even of someone not in the database. This state of affairs suggests a new measure, *differential privacy*, which, intuitively, captures the increased risk to one's privacy incurred by participating in a database. The techniques developed in a sequence of papers [8, 13, 3], culminating in those described in [12], can achieve any desired level of privacy under this measure. In many cases, extremely accurate information about the database can be provided while simultaneously ensuring very high levels of privacy.

1 Introduction

A statistic is a quantity computed from a sample. If a database is a representative sample of an underlying population, the goal of a privacy-preserving statistical database is to enable the user to learn properties of the population as a whole, while protecting the privacy of the individuals in the sample. The work discussed herein was originally motivated by exactly this problem: how to reveal useful information about the underlying population, as represented by the database, while preserving the privacy of individuals. Fortuitously, the techniques developed in [8, 13, 3] and particularly in [12] are so powerful as to broaden the scope of private data analysis beyond this orignal "representatitive" motivation, permitting privacy-preserving analysis of an object that is itself of intrinsic interest. For instance, the database may describe a concrete interconnection network – not a sample subnetwork – and we wish to reveal certain properties of the network without releasing information about individual edges or nodes. We therefore treat the more general problem of *privacy-preserving analysis of data*.

A rigorous treatment of privacy requires definitions: What constitutes a failure to preserve privacy? What is the power of the adversary whose goal it is to compromise privacy? What auxiliary information is available to the adversary (newspapers, medical studies, labor statistics) even without access to the database in question? Of course, utility also requires formal treatment, as releasing no information or only random noise clearly does not compromise privacy; we

M. Bugliesi et al. (Eds.): ICALP 2006, Part II, LNCS 4052, pp. 1–12, 2006.

Differential Privacy, informally

- Set-up:

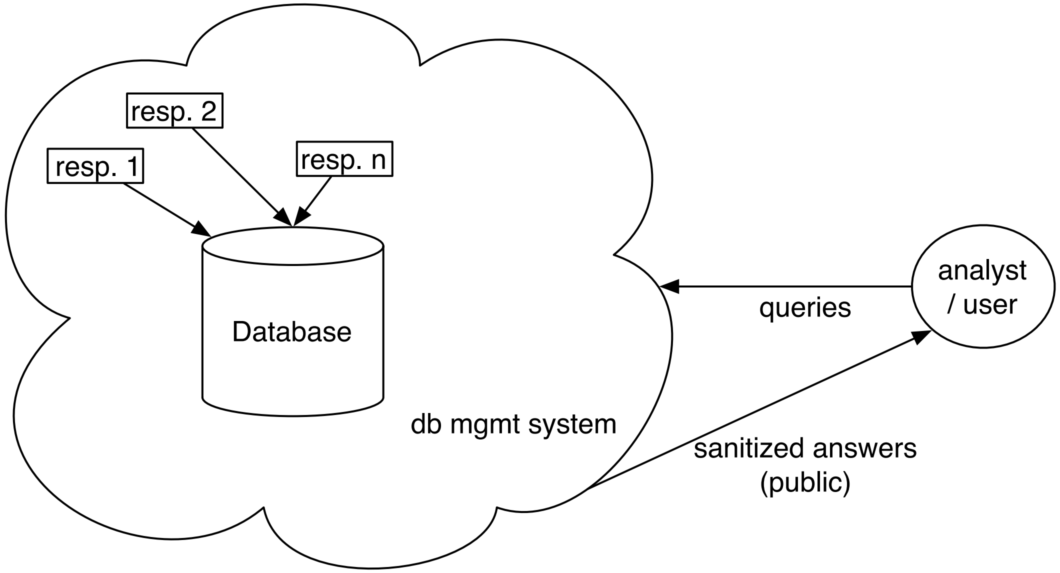

- Key Idea:
  - A differentially private mechanism randomly perturbs the answers to a query so that the output distribution over answers does not vary much if any given individual participates or not
  - Hard to infer if the data of any individual was used or not to answer the query

Differential Privacy, informally

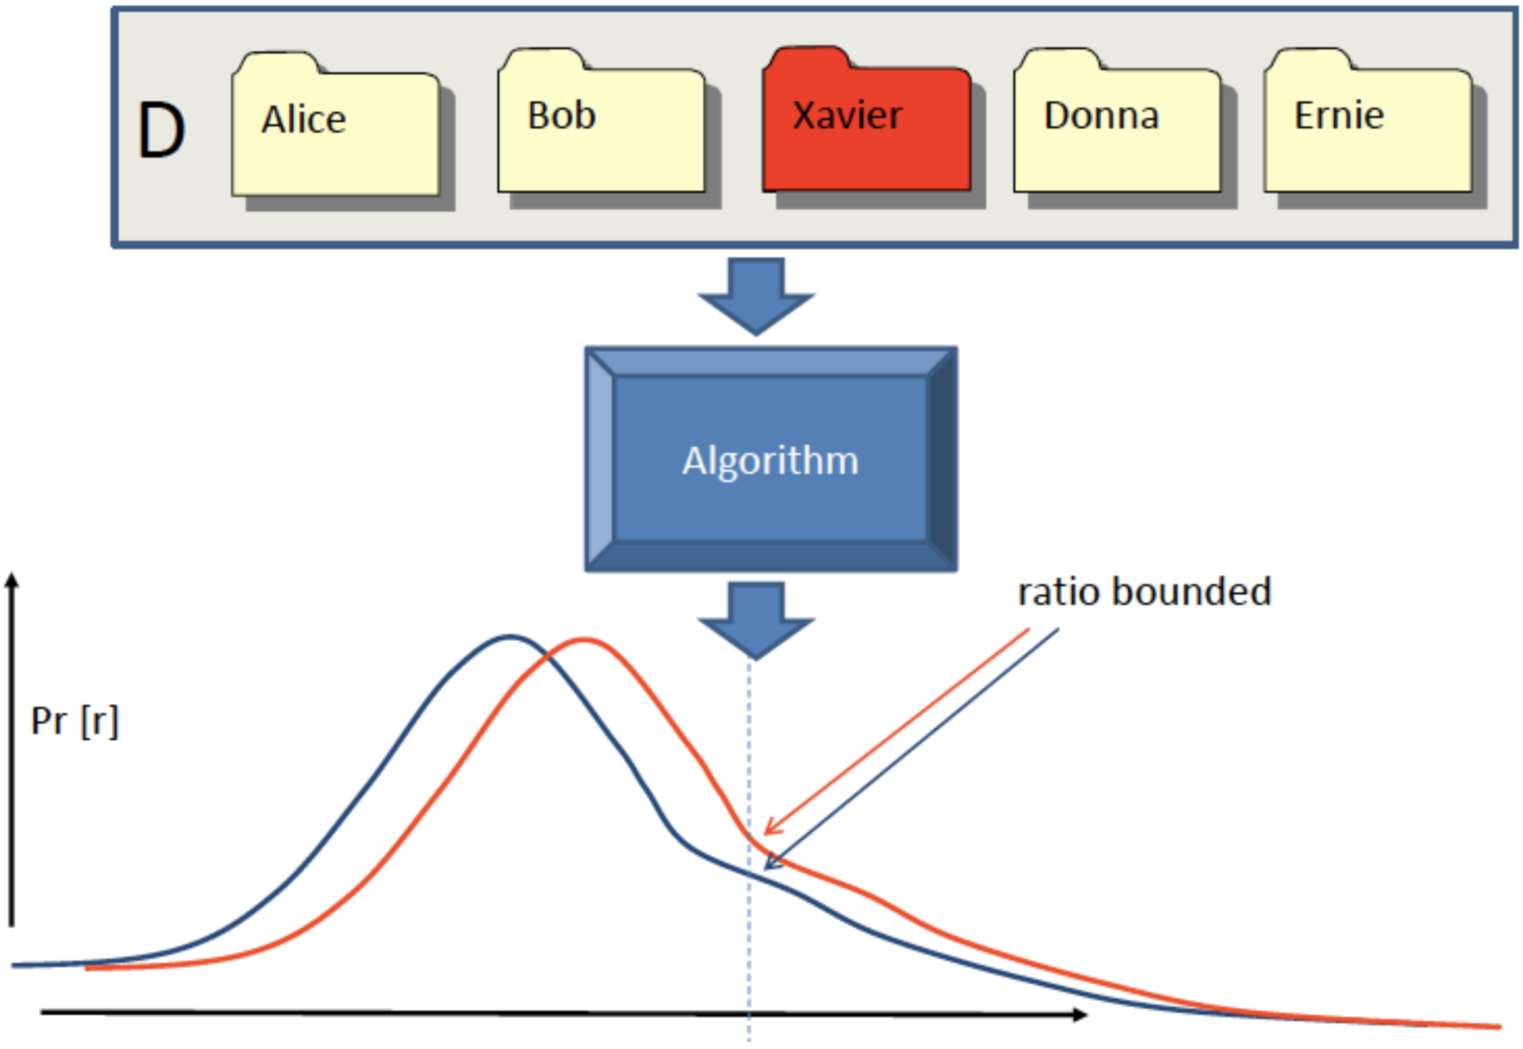

Differential Privacy, formally

- Set-up:

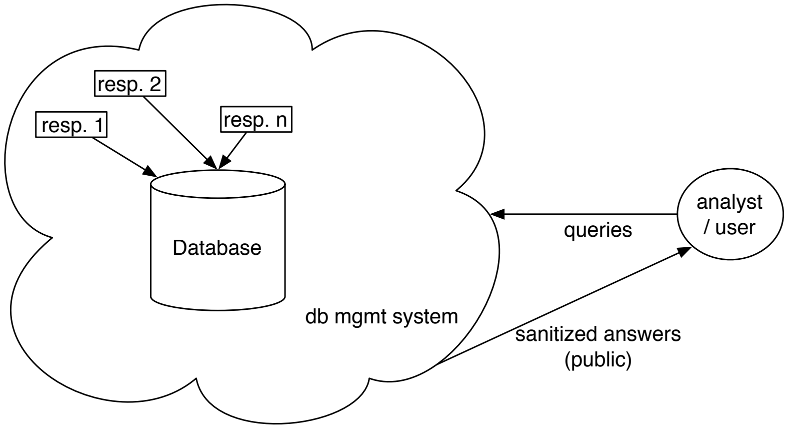

- Formally
  - Adj(d,d') a symmetric binary relation on the set D of databases
  - Adjacent databases differ by the data of a single individual
  - A mechanism $M : D \times \Omega \to (R, \mathcal{M})$ is (ε,δ)-DP if

  for all sets $\forall S \in \mathcal{M}$ and all databases d, d' s.t. Adj(d,d'), we have

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

Differential Privacy, formally

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

- Constant ε is typically small (i.e. ~ 0.1) - multiplicative error
- Constant δ is very small (i.e. ~0.01) - additive error
- If δ=0 then we have (ε,0)-DP or simply ε-DP
- Privacy definition depends on adjacency relation

Two Basic Differentially Private Mechanisms

- Database of CPS PI salaries d = [d_1,…d_n]
- Adjacency: $\mathrm{Adj}(d, d')$ iff for some i, $|d_i - d'_i| \leq \rho_i$

  $$d_j = d'_j, j \neq i$$

  - R = max {ρ_i}
- Analyst Query:

  $$q(d) = \frac{1}{n} \sum_{i=1}^{n} d_i$$

- The mechanism M(d) = q(d) + Lap(R/ε) is ε-DP
- The mechanism M(d) = q(d) + N(0,(κ(δ,ε) R)^2) is (ε,δ)-DP

$$\kappa(\delta, \epsilon) \leq 2\sqrt{2 \ln(2/\delta)}/\epsilon \qquad \left(\mathrm{Lap(b)\ pdf:}\ \frac{1}{2b} e^{-|x|/b}\right)$$

Resilience to post-processing

If mechanism M(d) is (ε,δ)-differentially private and f is an arbitrary function, then f (M(d)) is also (ε,δ)-differentially private

- f as the adversaries : Models arbitrary auxiliary or side information information the adversary may have. Privacy guarantee holds no matter what adversary does.
- f as our algorithm: If we access the database in a differentially private way, we don't have to worry about how our algorithm post-processes the result.

Differential privacy in computer science

• Data mining with differential privacy	A. Friedman 2010
• Combinatorial optimization with privacy	A. Roth et al 2010
• Mechanism design via differential privacy	M. Kearns et al 2012
• Privacy aware learning	J. Duchi 2013
• Streaming continuous data	C. Dwork et al 2010

Differential Private Filtering

Jerome Le Ny and George J. Pappas
Differentially private filtering
IEEE Transactions on Automatic Control
February 2014.

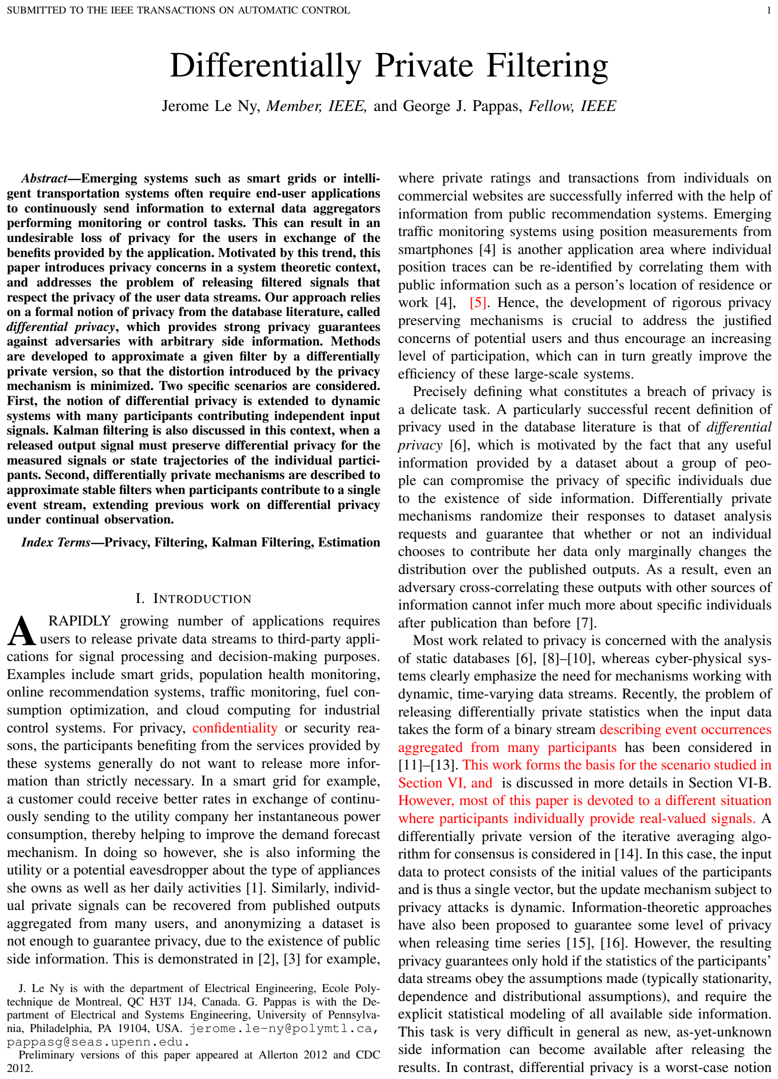

Differentially Private Filtering

Jerome Le Ny, *Member, IEEE,* and George J. Pappas, *Fellow, IEEE*

***Abstract*—Emerging systems such as smart grids or intelligent transportation systems often require end-user applications to continuously send information to external data aggregators performing monitoring or control tasks. This can result in an undesirable loss of privacy for the users in exchange of the benefits provided by the application. Motivated by this trend, this paper introduces privacy concerns in a system theoretic context, and addresses the problem of releasing filtered signals that respect the privacy of the user data streams. Our approach relies on a formal notion of privacy from the database literature, called *differential privacy*, which provides strong privacy guarantees against adversaries with arbitrary side information. Methods are developed to approximate a given filter by a differentially private version, so that the distortion introduced by the privacy mechanism is minimized. Two specific scenarios are considered. First, the notion of differential privacy is extended to dynamic systems with many participants contributing independent input signals. Kalman filtering is also discussed in this context, when a released output signal must preserve differential privacy for the measured signals or state trajectories of the individual participants. Second, differentially private mechanisms are described to approximate stable filters when participants contribute to a single event stream, extending previous work on differential privacy under continual observation.**

***Index Terms*—Privacy, Filtering, Kalman Filtering, Estimation**

I. Introduction

A RAPIDLY growing number of applications requires users to release private data streams to third-party applications for signal processing and decision-making purposes. Examples include smart grids, population health monitoring, online recommendation systems, traffic monitoring, fuel consumption optimization, and cloud computing for industrial control systems. For privacy, confidentiality or security reasons, the participants benefiting from the services provided by these systems generally do not want to release more information than strictly necessary. In a smart grid for example, a customer could receive better rates in exchange of continuously sending to the utility company her instantaneous power consumption, thereby helping to improve the demand forecast mechanism. In doing so however, she is also informing the utility or a potential eavesdropper about the type of appliances she owns as well as her daily activities [1]. Similarly, individual private signals can be recovered from published outputs aggregated from many users, and anonymizing a dataset is not enough to guarantee privacy, due to the existence of public side information. This is demonstrated in [2], [3] for example, where private ratings and transactions from individuals on commercial websites are successfully inferred with the help of information from public recommendation systems. Emerging traffic monitoring systems using position measurements from smartphones [4] is another application area where individual position traces can be re-identified by correlating them with public information such as a person's location of residence or work [4], [5]. Hence, the development of rigorous privacy preserving mechanisms is crucial to address the justified concerns of potential users and thus encourage an increasing level of participation, which can in turn greatly improve the efficiency of these large-scale systems.

Precisely defining what constitutes a breach of privacy is a delicate task. A particularly successful recent definition of privacy used in the database literature is that of *differential privacy* [6], which is motivated by the fact that any useful information provided by a dataset about a group of people can compromise the privacy of specific individuals due to the existence of side information. Differentially private mechanisms randomize their responses to dataset analysis requests and guarantee that whether or not an individual chooses to contribute her data only marginally changes the distribution over the published outputs. As a result, even an adversary cross-correlating these outputs with other sources of information cannot infer much more about specific individuals after publication than before [7].

Most work related to privacy is concerned with the analysis of static databases [6], [8]–[10], whereas cyber-physical systems clearly emphasize the need for mechanisms working with dynamic, time-varying data streams. Recently, the problem of releasing differentially private statistics when the input data takes the form of a binary stream describing event occurrences aggregated from many participants has been considered in [11]–[13]. This work forms the basis for the scenario studied in Section VI, and is discussed in more details in Section VI-B. However, most of this paper is devoted to a different situation where participants individually provide real-valued signals. A differentially private version of the iterative averaging algorithm for consensus is considered in [14]. In this case, the input data to protect consists of the initial values of the participants and is thus a single vector, but the update mechanism subject to privacy attacks is dynamic. Information-theoretic approaches have also been proposed to guarantee some level of privacy when releasing time series [15], [16]. However, the resulting privacy guarantees only hold if the statistics of the participants' data streams obey the assumptions made (typically stationarity, dependence and distributional assumptions), and require the explicit statistical modeling of all available side information. This task is very difficult in general as new, as-yet-unknown side information can become available after releasing the results. In contrast, differential privacy is a worst-case notion

J. Le Ny is with the department of Electrical Engineering, Ecole Polytechnique de Montreal, QC H3T 1J4, Canada. G. Pappas is with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA. jerome.le-ny@polymtl.ca, pappasg@seas.upenn.edu.

Preliminary versions of this paper appeared at Allerton 2012 and CDC 2012.

Private Filtering : Event Streams Counters

- Processing binary input signal (0-1 events) – DP linear filter approximation
  - Event-level privacy: each user contributes a unique event: two inputs differing by a single event must be hard to distinguish
  - Ex: Counter [Dwork et al., 2010], [Chan et al., 2011] (unstable stable filter)
  - Ex: Certain stable filters with slowly decreasing impulse response [Bolot et al., 2011]
  - Complicated algorithms (non-recursive, not finite memory), hard to generalize, poor performance of the approach for the approximation of stable filters

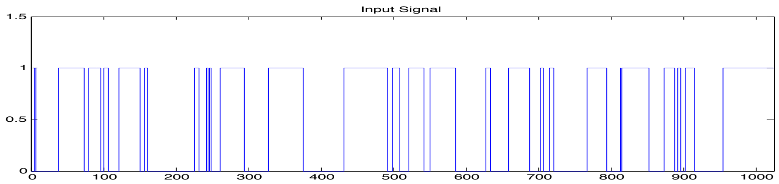

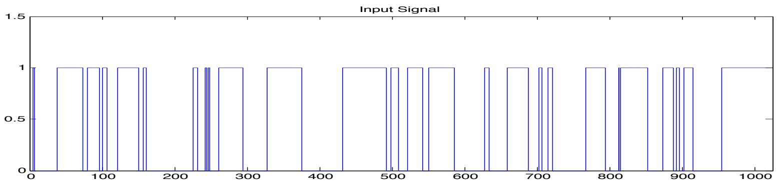

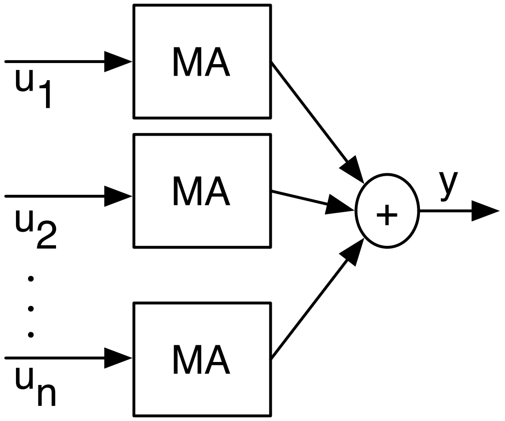

$$y_t = \sum_{i=1}^{n} \frac{1}{l} \sum_{k=t-l+1}^{t} u_{i,k}$$

Differentially Private Filtering

- Approximate filter $y = \sum_{i=1}^{n} G_i u_i$ by a differentially private version

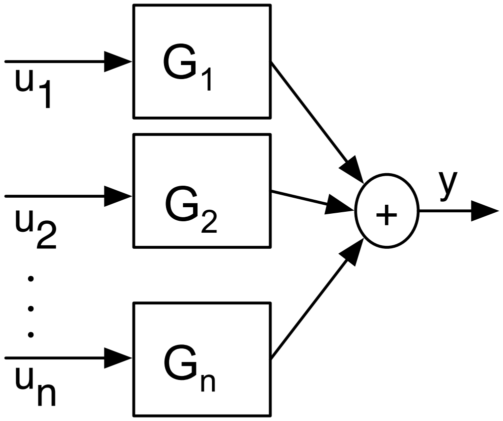

- Adjacency relation

$$\mathrm{Adj}^b(u,u') \text{ iff for some } i, \|u_i - u_i'\|_2 \leq B, \text{and } u_j = u_j' \text{ for all } j \neq i.$$

Differentially Private Filtering

- Approximate filter $y = \sum_{i=1}^{n} G_i u_i$ by a differentially private version
  - Adjacency relation

$$\mathrm{Adj}^b(u, u') \text{ iff for some } i, \|u_i - u'_i\|_2 \leq B, \text{and } u_j = u'_j \text{ for all } j \neq i.$$

Theorem : For (ε,δ)-differential privacy, can add white Gaussian noise proportional to the maximum incremental gain with respect to the input channels

  - Incremental gain: $\|(Gu)_{0:T} - (Gu')_{0:T}\|_2 \leq \gamma \|u_{0:T} - u'_{0:T}\|_2, \; \forall u, u', \forall T$
  - Generalizes mechanism of [Dwork et al., 2006] to dynamic setting with continuous streams of real-time data
  - System and control theoretic tools can be used to design differentially private mechanisms for continuous data streams

Differentially Private Filtering

Approximate filter $y = \sum_{i=1}^{n} G_i u_i$ by a differentially private version

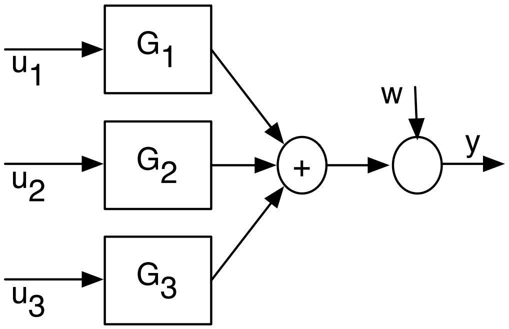

$$\begin{aligned} \mathrm{Adj}^b(u, u') \text{ iff for some } i, \|u_i - u_i'\|_{r_i} \leq b_i, \\ \text{and } u_j = u_j' \text{ for all } j \neq i. \end{aligned}$$

$$B \geq \max_{1 \leq i \leq n} \{\gamma_{r_i,1}(\mathcal{G}_i)\, b_i\}$$

Theorem: The mechanism M(u)=G(u)+w where w is white noise with

$$w_t \sim (\mathrm{Lap}(B/\varepsilon))^m$$

is ε-differentially private.

Differentially Private Filtering

Approximate filter $y = \sum_{i=1}^{n} G_i u_i$ by a differentially private version

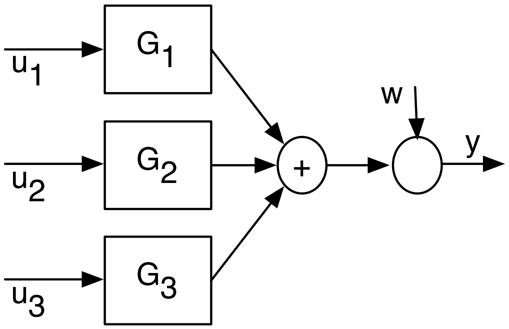

$$\begin{aligned} \mathrm{Adj}^b(u,u') \text{ iff for some } i, \|u_i - u_i'\|_{r_i} \le b_i, \\ \text{and } u_j = u_j' \text{ for all } j \neq i. \end{aligned}$$

$$\sigma \geq \kappa(\delta,\epsilon)\max_{1\leq i\leq n}\{\gamma_{r_i,2}(\mathcal{G}_i)\, b_i\}.$$

Theorem: The mechanism M(u)=G(u)+w where w is white noise with

$$w_t \sim N(0, \sigma^2 I_m)$$

is (ε,δ)-differentially private.

Differential Privacy Architectures

- Two basic architectures for (ε,δ)-differential privacy

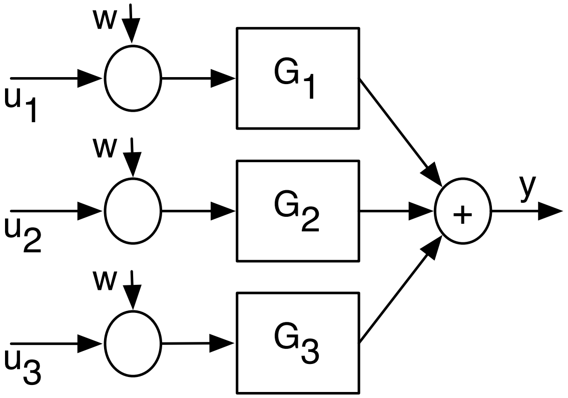

$$w \sim \mathcal{N}(0, \sigma^2)$$

$$\sigma = \kappa(\epsilon, \delta) B$$

$$MSE = \sigma^2 \sum_{i=1}^{n} \|G_i\|_2^2$$

Input perturbation mechanism

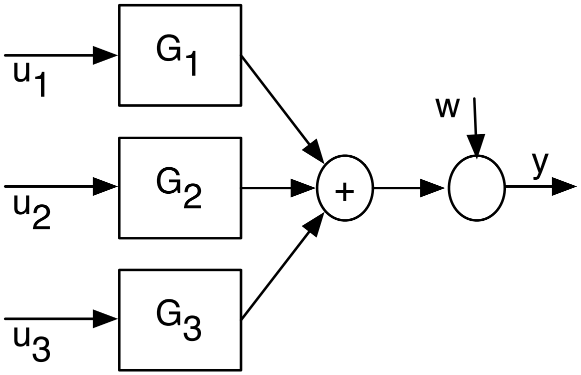

$$w \sim \mathcal{N}(0, \sigma^2)$$

$$\sigma = \kappa(\epsilon, \delta) B \max_{1 \le i \le n} \{\|G_i\|_\infty\}$$

$$MSE = \sigma^2$$

Output perturbation mechanism

Event Streams Counters

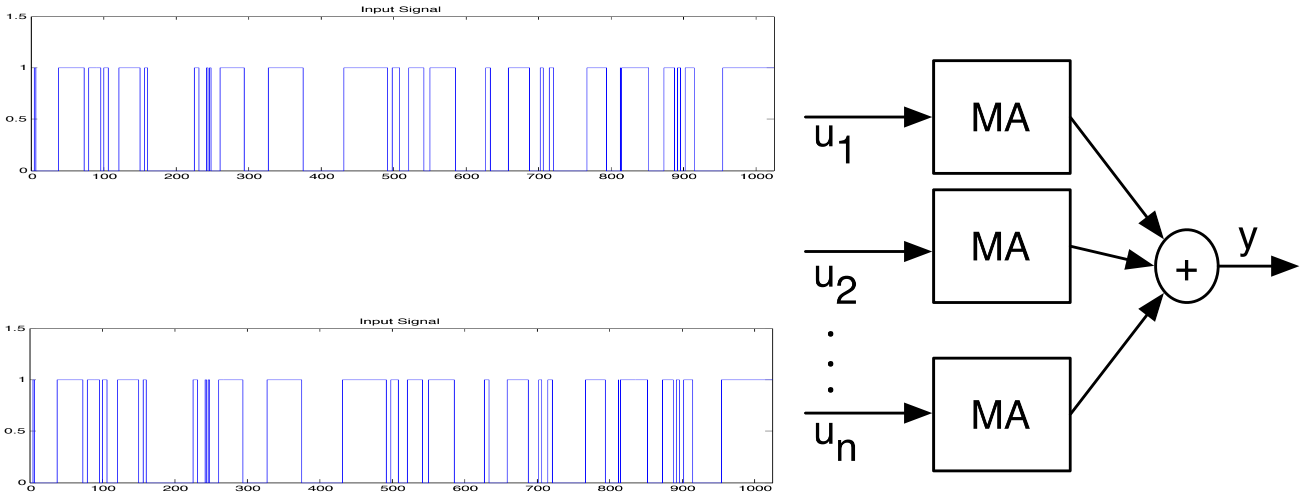

$$y_t = \sum_{i=1}^{n} \frac{1}{l} \sum_{k=t-l+1}^{t} u_{i,k}$$

$$\|G_i\|_2^2 = \frac{1}{l}, \|G_i\|_\infty = 1$$

→ output perturbation better than input perturbation iff $n > l$

Event Streams Counters

$$y_t = \sum_{i=1}^{n} \frac{1}{l} \sum_{k=t-l+1}^{t} u_{i,k}$$

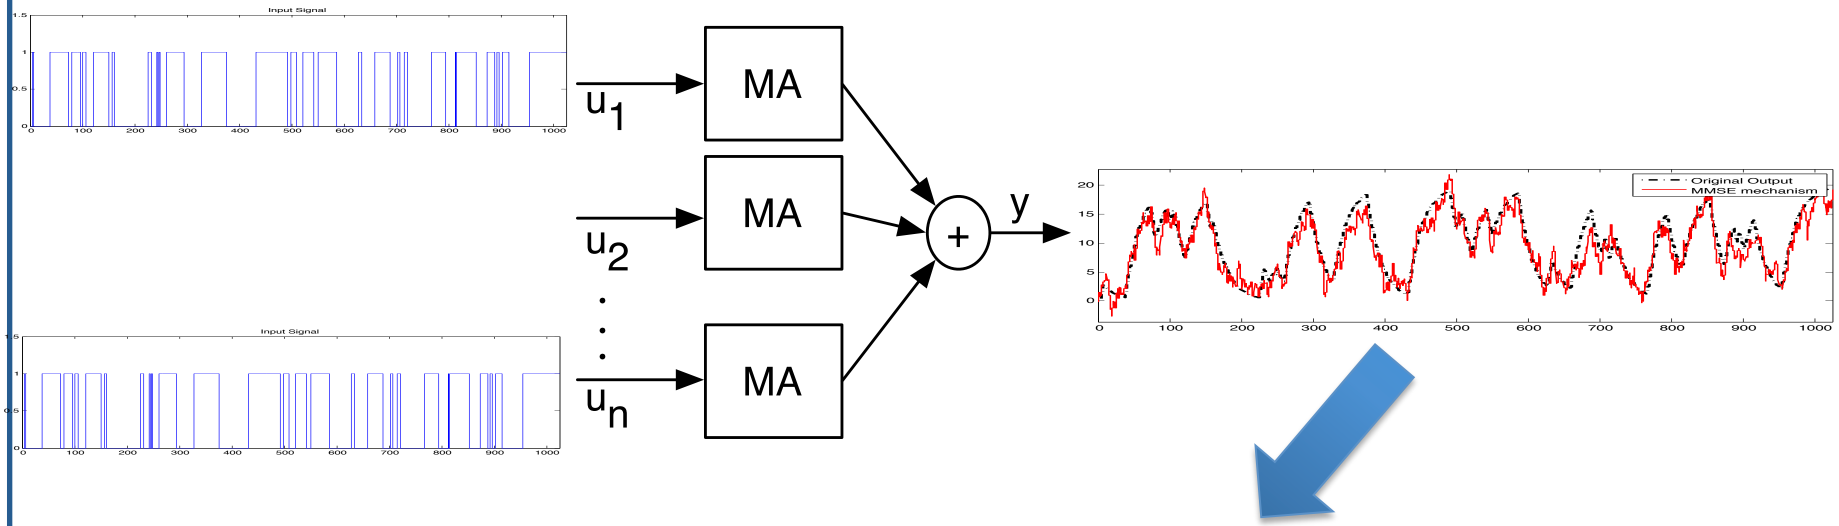

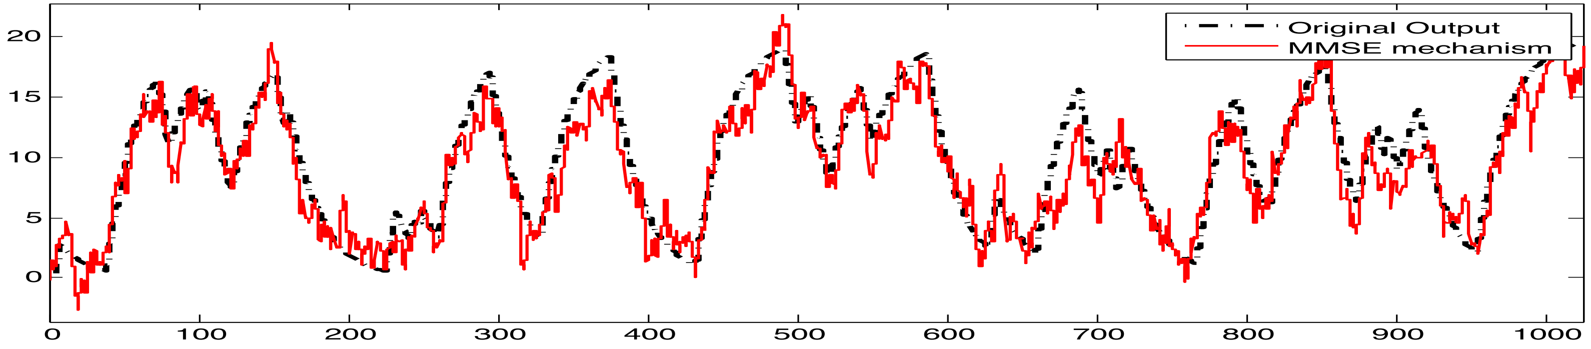

A Traffic Monitoring Example

- Traffic velocity estimation using individual location traces from smartphones (ex: Mobile Millenium, Bayen et al.)

$$x_{i,t+1} = \begin{bmatrix} 1 & T_s \\ 0 & 1 \end{bmatrix} x_{i,t} + \sigma_{i1} \begin{bmatrix} T_s^2/2 & 0 \\ T_s & 0 \end{bmatrix} w_{i,t},$$

$$y_{i,t} = \begin{bmatrix} 1 & 0 \end{bmatrix} x_{i,t} + \sigma_{i2} \begin{bmatrix} 0 & 1 \end{bmatrix} w_{i,t}$$

$$\text{Estimate } \frac{1}{n} \sum_{i=1}^{n} x_{i,2,t}$$

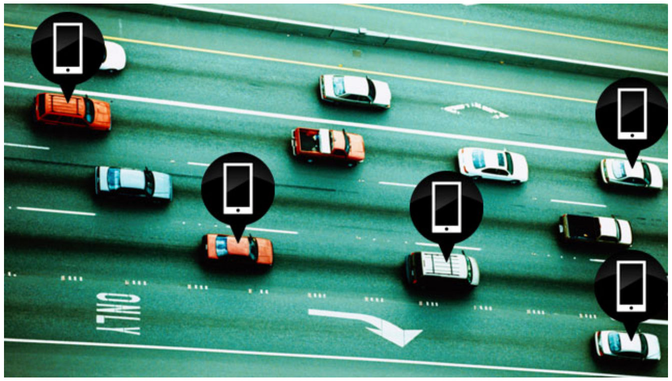

Differentially Private Kalman Filtering

- For Kalman Filtering, we have additional public information about the dynamics generating the user signals

$$\begin{aligned} x_{i,t+1} &= A_i x_{i,t} + B_i w_{i,t} \\ y_{i,t} &= C_i x_{i,t} + D_i w_{i,t} \end{aligned}$$

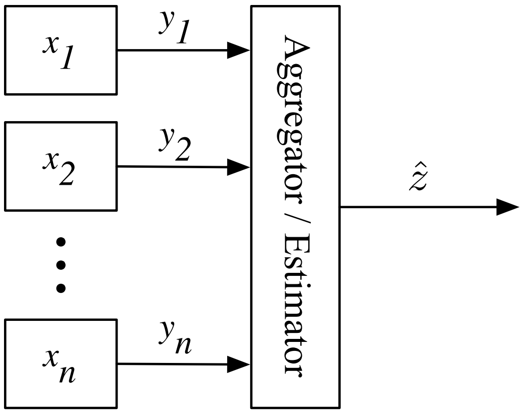

$$z_t = \sum_{i=1}^{n} L_i x_{i,t}$$

- Estimation objective: $z_t = \sum_{i=1}^{n} L_i x_{i,t} \qquad \min_{\hat{z}} \lim_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}\left[\|z_t - \hat{z}_t\|_2^2\right]$

- Adjacency relation:

$$\mathrm{Adj}^{\rho}(x, x') \text{ iff for some } i, \|x_i - x'_i\|_2 \leq \rho_i, \text{and } x_j = x'_j \text{ for all } j \neq i.$$

  - Cannot distinguish between two sufficiently close state trajectories of a user

Differentially Private Kalman Filtering

$$x_{i,t+1} = A_i x_{i,t} + B_i w_{i,t}$$
$$y_{i,t} = C_i x_{i,t} + D_i w_{i,t}$$

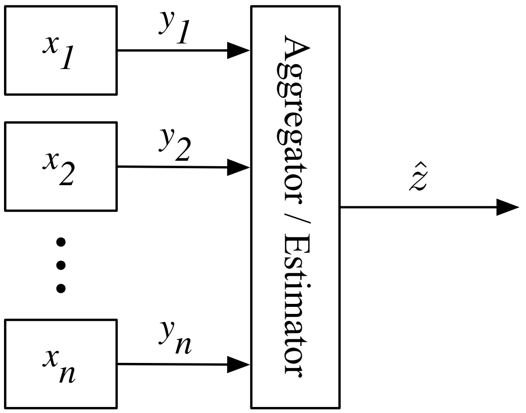

$$z_t = \sum_{i=1}^{n} L_i x_{i,t}$$

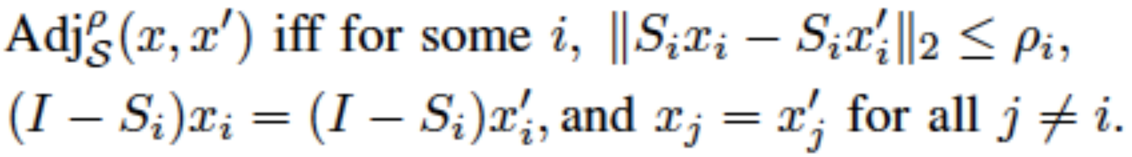

$\mathrm{Adj}_S^{\rho}(x, x')$ iff for some i, $\|S_i x_i - S_i x_i'\|_2 \le \rho_i$, $(I - S_i)x_i = (I - S_i)x_i'$, and $x_j = x_j'$ for all $j \neq i$.

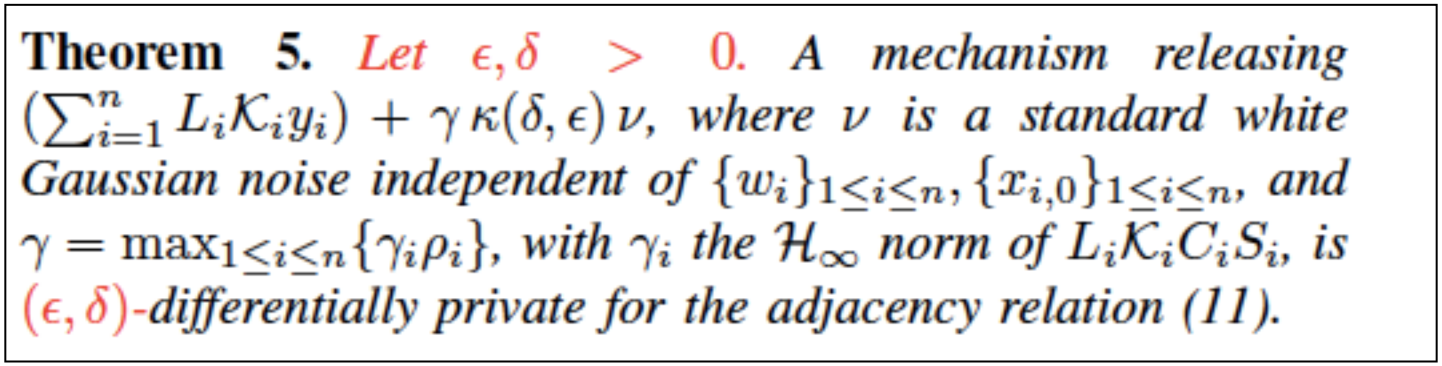

Theorem 5. *Let $\epsilon, \delta > 0$. A mechanism releasing $(\sum_{i=1}^{n} L_i \mathcal{K}_i y_i) + \gamma\, \kappa(\delta, \epsilon)\, \nu$, where ν is a standard white Gaussian noise independent of $\{w_i\}_{1 \le i \le n}, \{x_{i,0}\}_{1 \le i \le n}$, and $\gamma = \max_{1 \le i \le n}\{\gamma_i \rho_i\}$, with γ_i the $\mathcal{H}_\infty$ norm of $L_i \mathcal{K}_i C_i S_i$, is (ϵ, δ)-differentially private for the adjacency relation (11).*

Kalman Filtering - Differentially Private Mechanisms

- Can add the previous input and output perturbation schemes to the standard Kalman filter

$$x_{i,t+1} = A_i x_{i,t} + B_i w_{i,t}$$
$$y_{i,t} = C_i x_{i,t} + D_i w_{i,t}$$

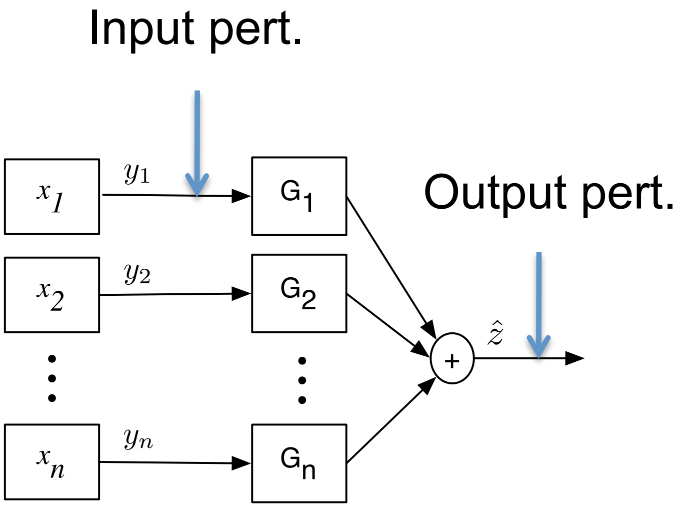

- For the input perturbation scheme, can take into account the additional privacy-preserving noise in the redesign of the KF

Filter Redesign for Output Perturbation Scheme

$$x_{i,t+1} = A_i x_{i,t} + B_i w_{i,t}$$
$$y_{i,t} = C_i x_{i,t} + D_i w_{i,t}$$

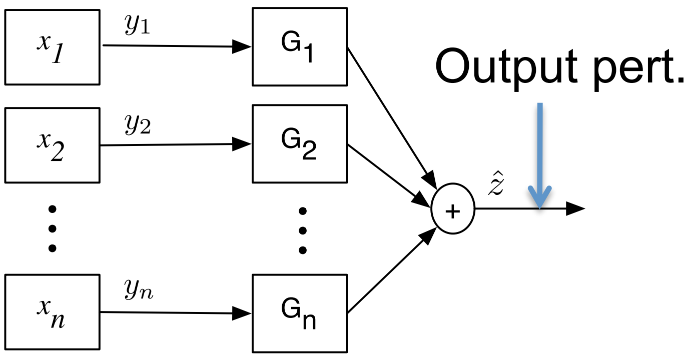

- For the output perturbation scheme, can redesign the filter to trade-off the estimation error and the H_∞ norm of the filter
  - Overall MSE is

$$\left(\sum_{i=1}^{n} \|TF(w_i \to e_i)\|_2^2\right) + \kappa(\delta,\epsilon)^2 \max_{1\leq i\leq n} \{\rho_i^2 \|TF(x_i \to \hat{z}_i)\|_\infty^2\}$$

- Multi-objective H_2/H_∞ optimization problem
  - Lyapunov shaping using Linear Matrix Inequalities
  - Distinguish between stable and unstable dynamics

A Traffic Monitoring Example

- Traffic velocity estimation using individual location traces from smartphones (ex: Mobile Millenium, Bayen et al.)

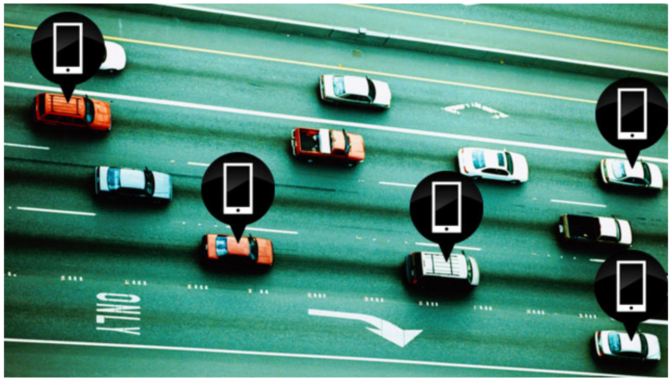

$$x_{i,t+1} = \begin{bmatrix} 1 & T_s \\ 0 & 1 \end{bmatrix} x_{i,t} + \sigma_{i1} \begin{bmatrix} T_s^2/2 & 0 \\ T_s & 0 \end{bmatrix} w_{i,t},$$

$$y_{i,t} = \begin{bmatrix} 1 & 0 \end{bmatrix} x_{i,t} + \sigma_{i2} \begin{bmatrix} 0 & 1 \end{bmatrix} w_{i,t}$$

$$\text{Estimate } \frac{1}{n} \sum_{i=1}^{n} x_{i,2,t}$$

$$\boxed{\begin{aligned} \rho &= 100\text{m} \\ \epsilon &= \ln 3 \\ \delta &= 0.05, \\ T_s &= 1\text{s} \\ \sigma_{i1} &= \sigma_{i2} = 1 \\ n &= 200 \end{aligned}}$$

Traffic Monitoring – input versus output architectures

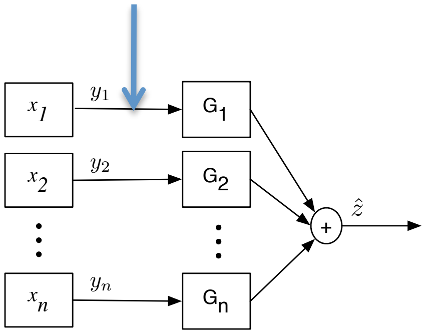

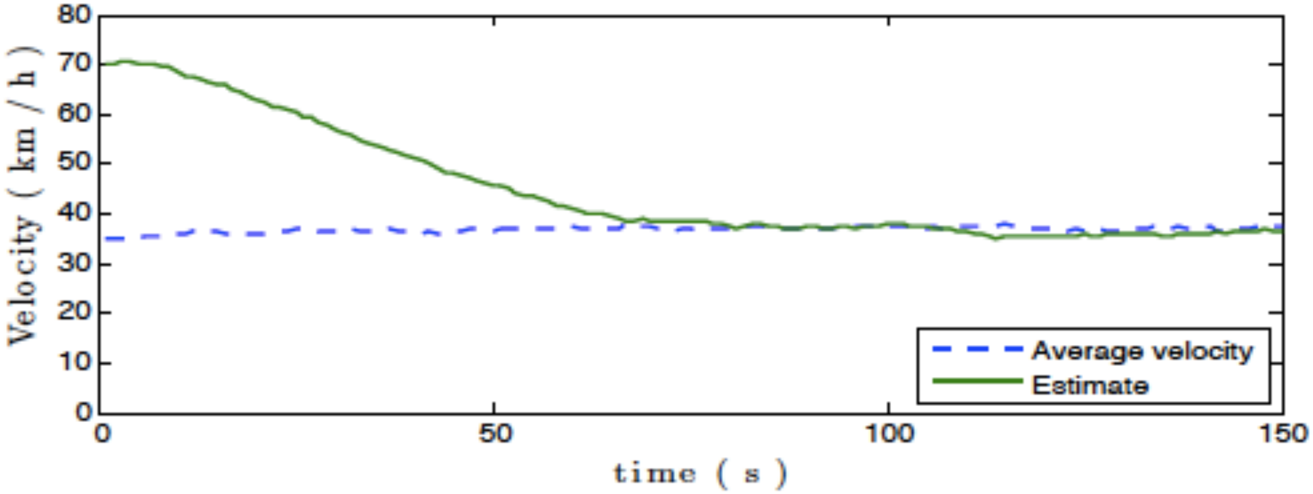

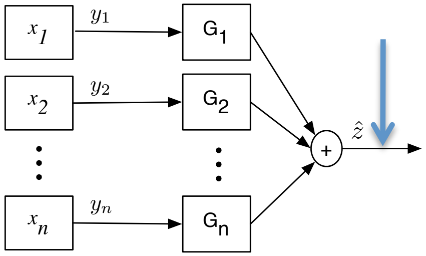

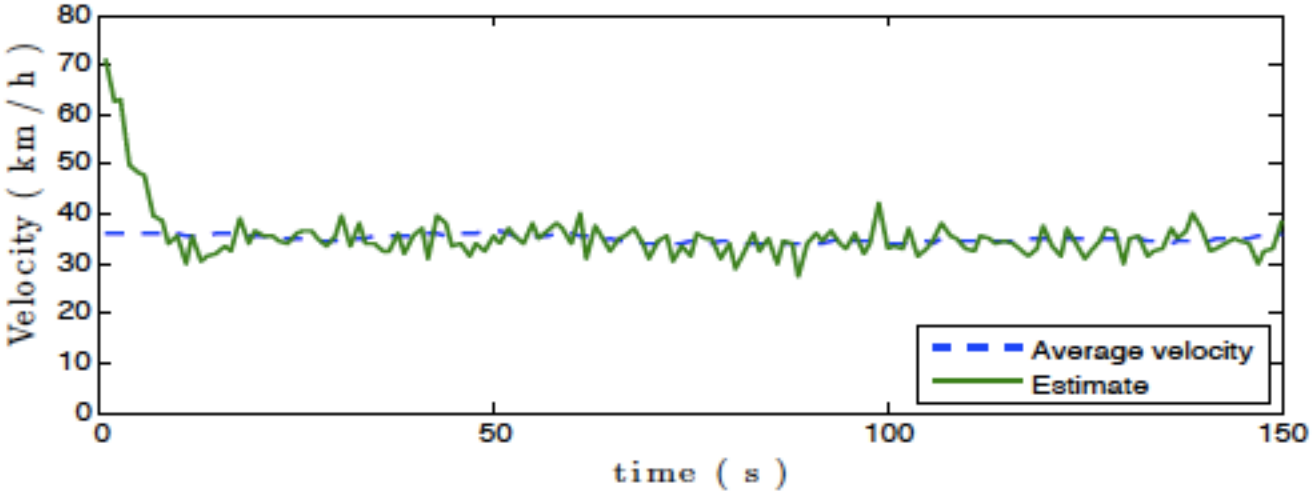

Traffic Monitoring – utility/privacy tradeoff

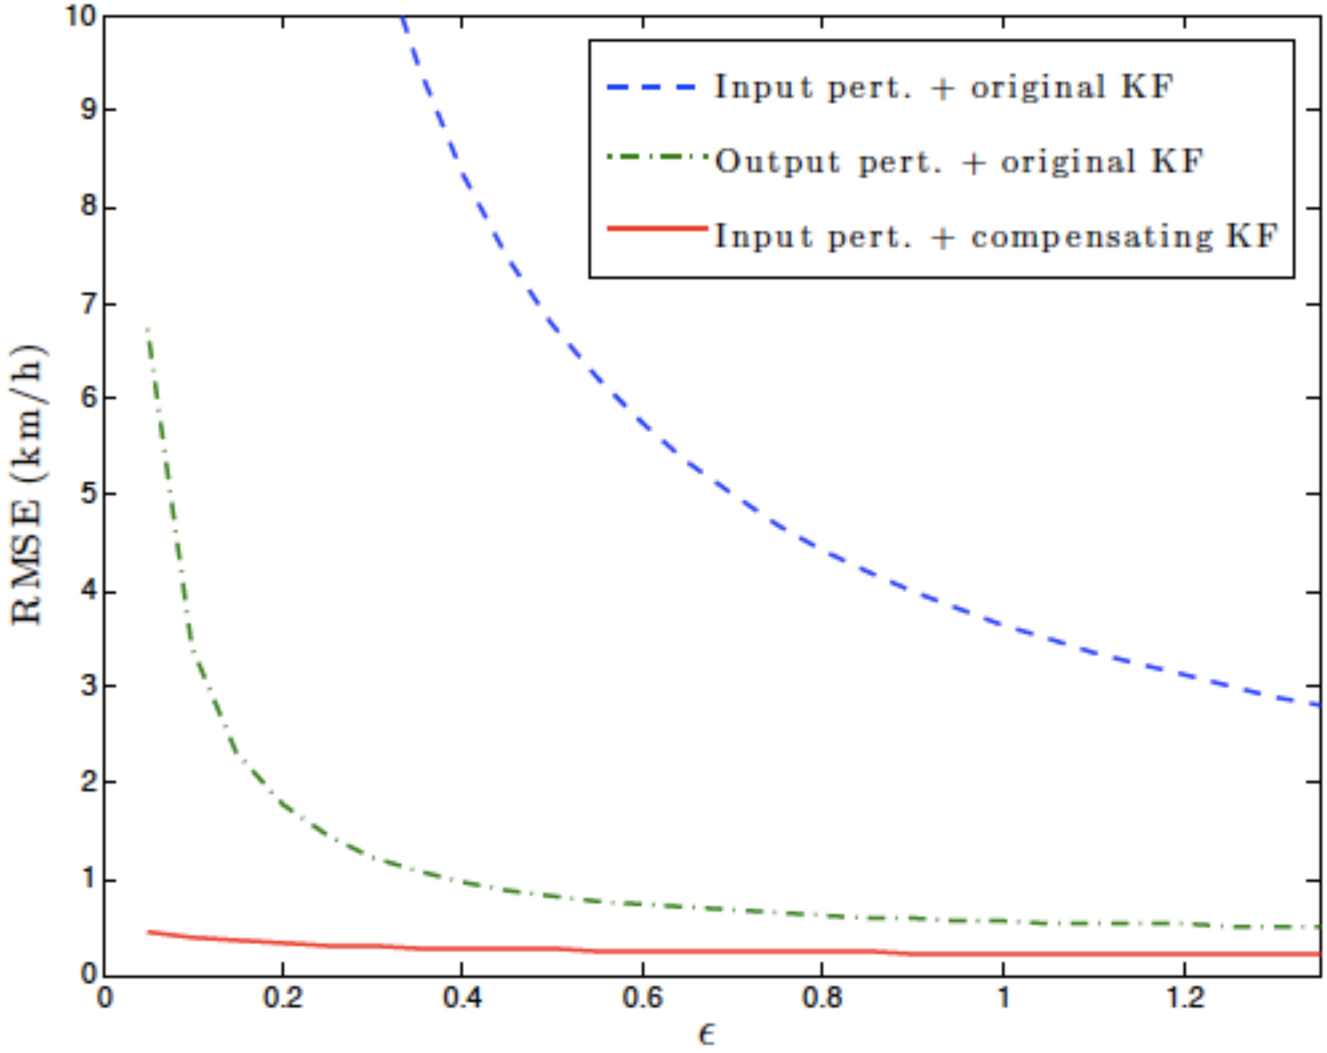

Traffic Monitoring – convergence/privacy tradeoff

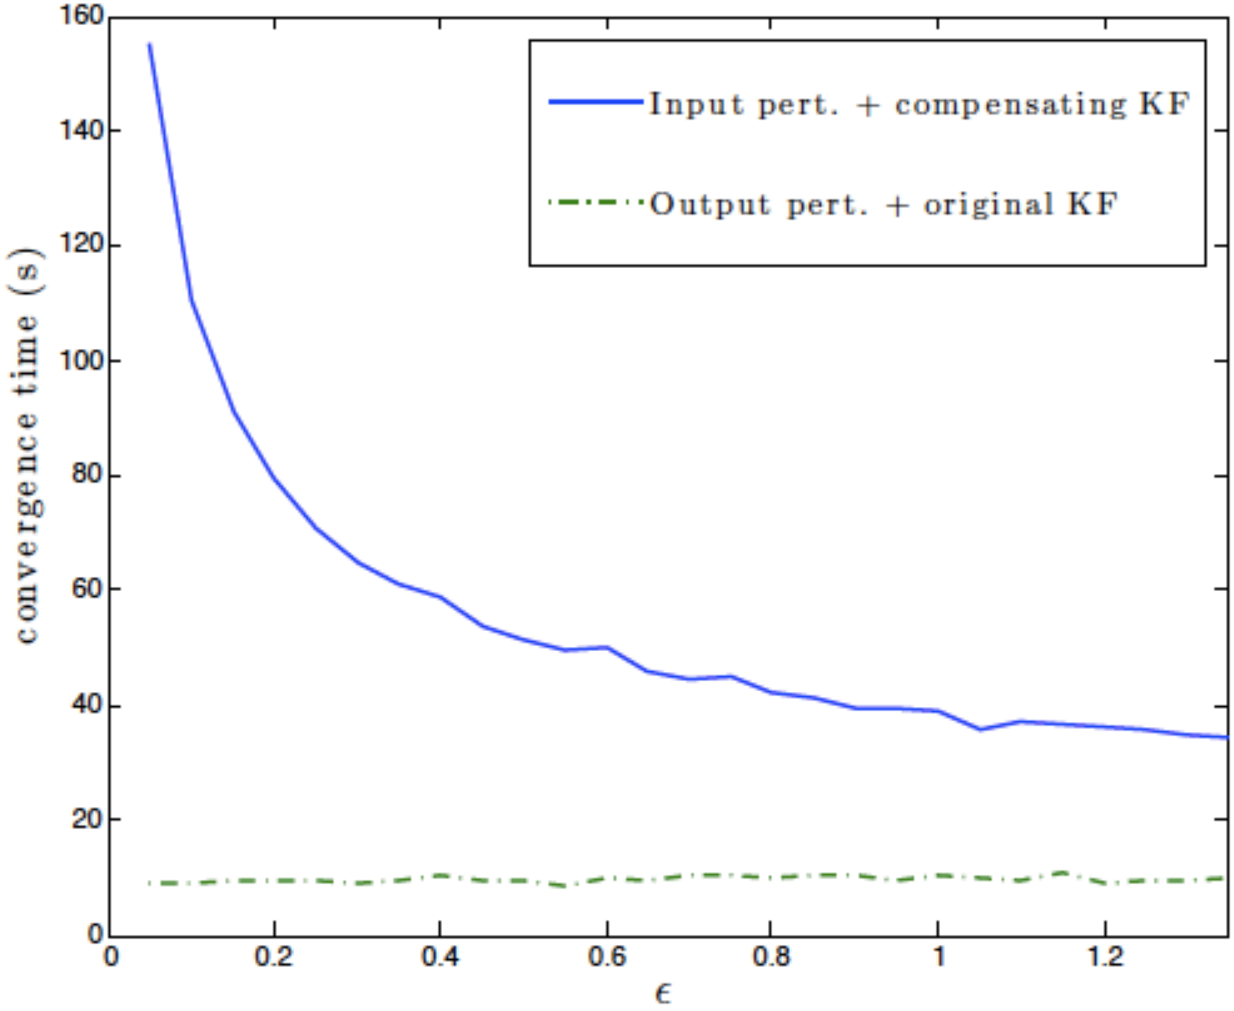

Summary

- Many cyber-physical applications raise privacy concerns that need to be addressed to encourage user participation

- Need privacy-preserving mechanisms for various types of dynamic systems and data

- Characterizing privacy-utility tradeoff requires a quantitative definition of privacy

- System and control theoretic tools (optimal estimators, system gains) can be used to design differentially private mechanisms

Next steps

- Fundamental limits between privacy and performance
- Scaling laws
- Apply to Smart Meter models and data
- Protect controllers and strategy not data
- Privacy pricing

A science of CPS privacy

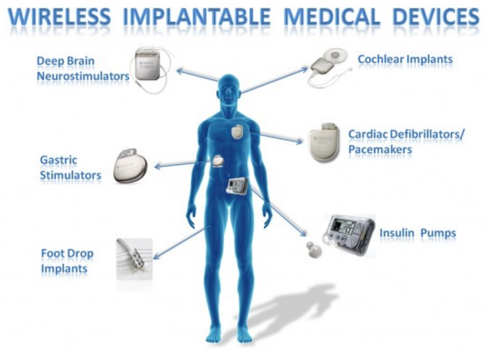
WIRELESS IMPLANTABLE MEDICAL DEVICES
Deep Brain Neurostimulators
Cochlear Implants
Cardiac Defibrillators/ Pacemakers
Gastric Stimulators
Insulin Pumps
Foot Drop Implants

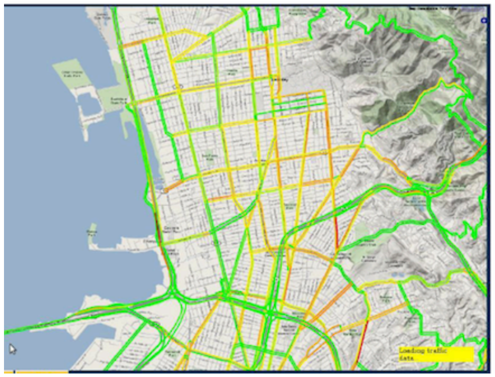

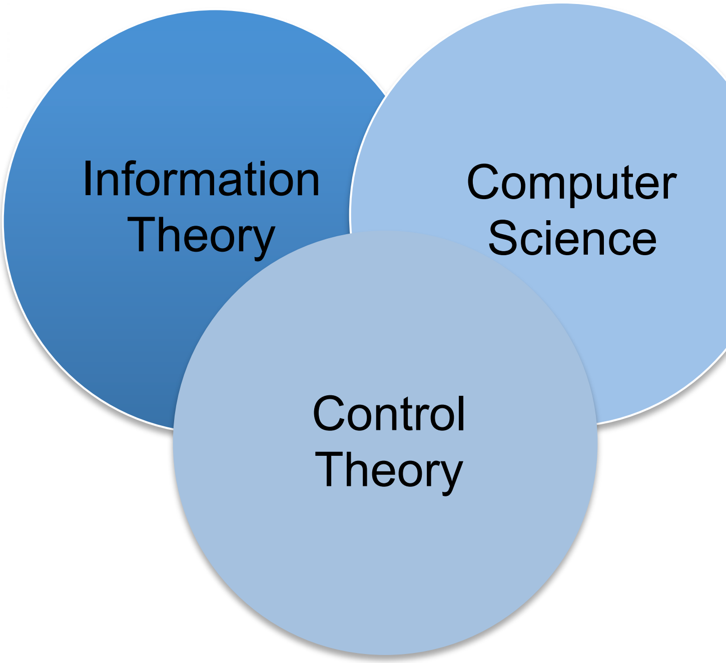

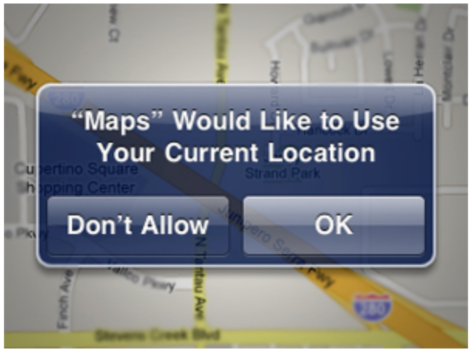

"Maps" Would Like to Use Your Current Location