Differentially Private Filtering
Application to Traffic Estimation

Jerome Le Ny
Polytechnique Montreal and GERAD

IPAM Workshop on Traffic Estimation
UCLA, October 13 2015

Based partly on joint work with Hubert André, Meisam Mohammady (Poly), George Pappas (UPenn), Ahmed Touati

GERAD
GROUP FOR RESEARCH IN DECISION ANALYSIS

POLYTECHNIQUE MONTRÉAL

PRIVACY CONCERNS IN CYBER-PHYSICAL APPLICATIONS

POLYTECHNIQUE MONTRÉAL

[image: Inst. Mech. Eng.]

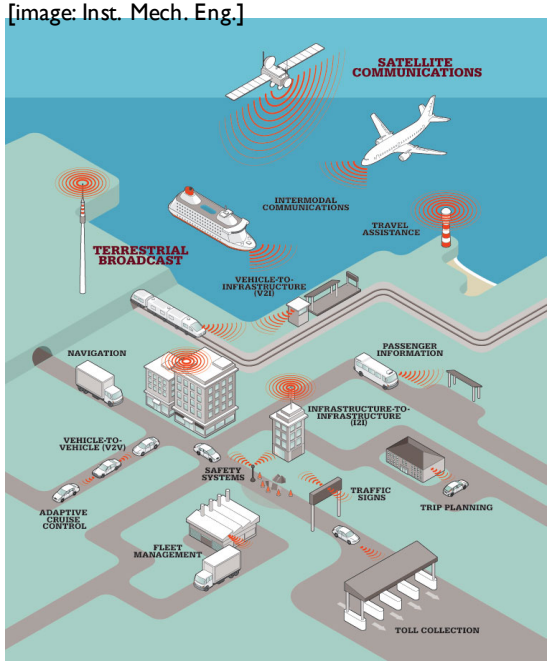

Intelligent Transportation Systems

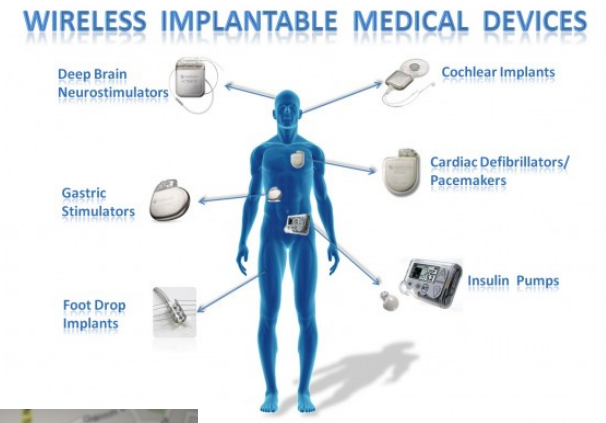

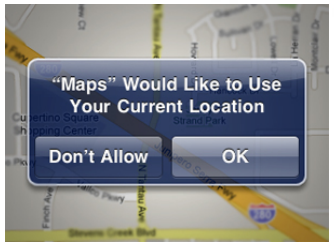

Location based services

[UCR]

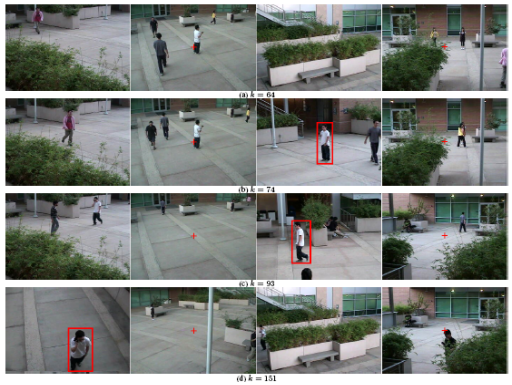

Camera Networks

smart meters

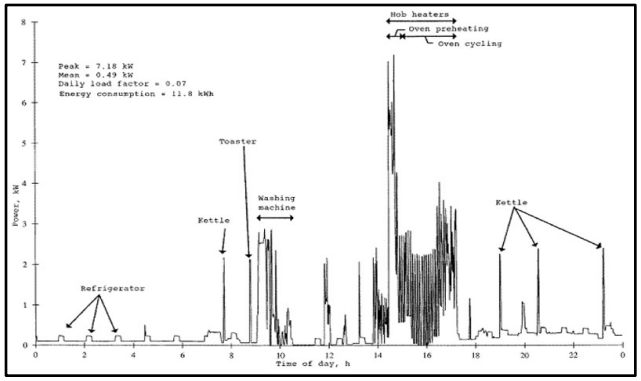

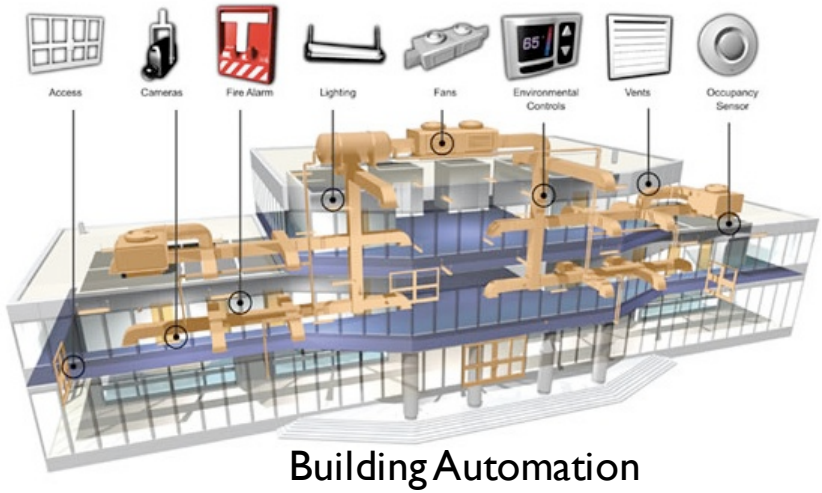

Building Automation

TRACKING IN SMART BUILDINGS

POLYTECHNIQUE MONTRÉAL

Some benefits:

- HVAC optimization
- Emergency evacuation
- …

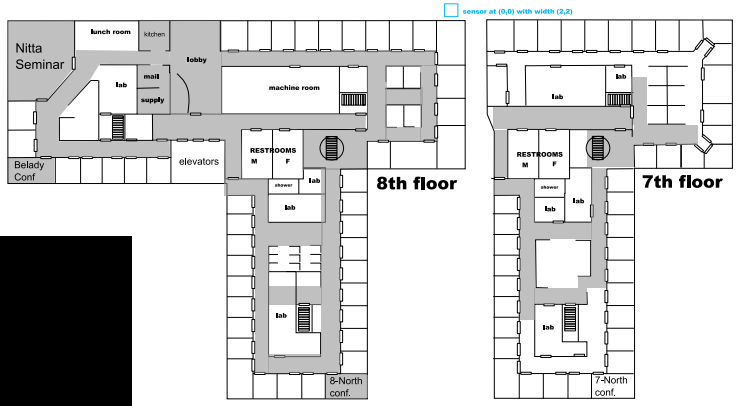

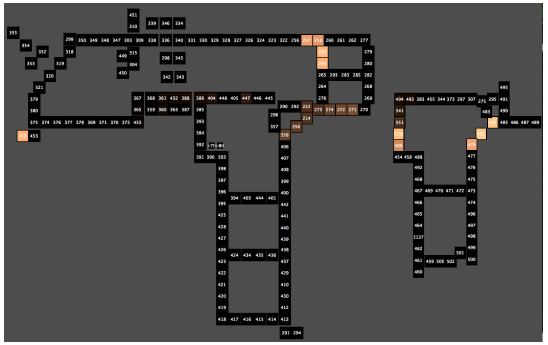

PRIVACY CONCERNS IN MODERN TRAFFIC ESTIMATION SYSTEMS

POLYTECHNIQUE MONTRÉAL

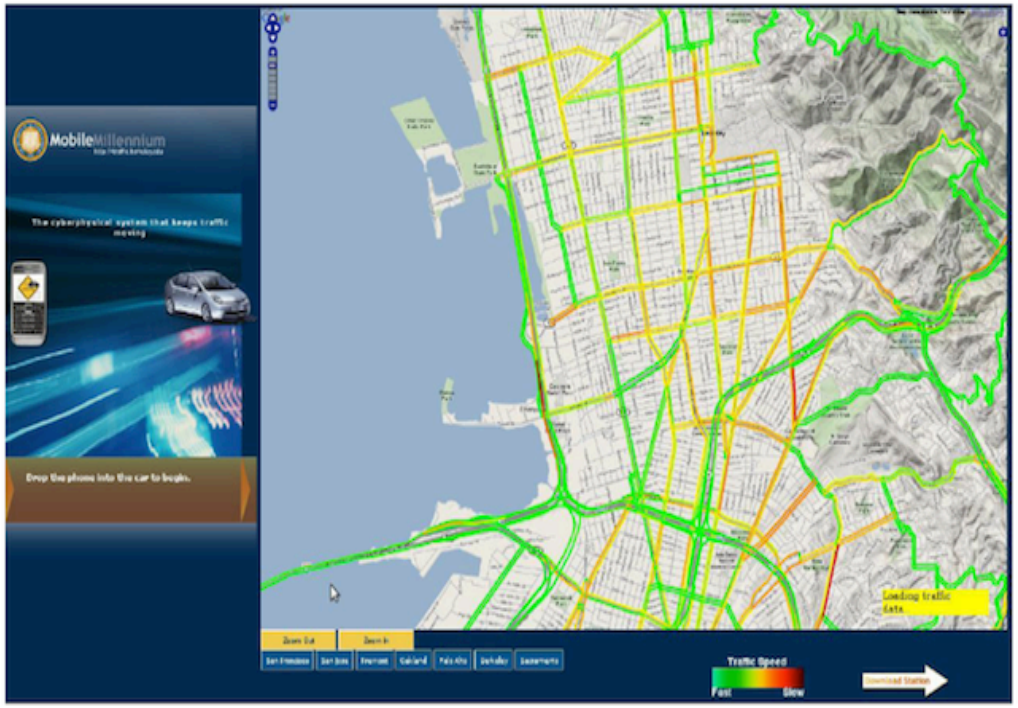

Virtual Trip Lines for Distributed Privacy-Preserving Traffic Monitoring

Baik Hoh, Marco Gruteser
WINLAB / ECE Dept., Rutgers Univ.
Piscataway, NJ USA
baikhoh,gruteser@winlab.rutgers.edu

Ryan Herring, Jeff Ban[a], Dan Work, Juan-Carlos Herrera, Alexandre Bayen
Civil Engineering Dept., UC Berkeley
Berkeley, CA USA
ryanherring,dbwork,jcherrera,bayen@berkeley.edu,[a]xban@calccit.org

Murali Annavaram, Quinn Jacobson
Nokia Research Center
Palo Alto, CA USA
annavara@usc.edu,quinn.jacobson@nokia.com

ABSTRACT

Automotive traffic monitoring using probe vehicles with Global Positioning System receivers promises significant improvements in cost, coverage, and accuracy. Current approaches, however, raise privacy concerns because they require participants to reveal their positions to an external traffic monitoring server. To address this challenge, we propose a system based on virtual trip lines and an associated cloaking technique. Virtual trip lines are geographic markers that indicate where vehicles should provide location updates. These markers can be placed to avoid particularly privacy sensitive locations. They also allow aggregating and cloaking sev-

General Terms

Algorithms, Design, Experimentation, Security

Keywords

Privacy, GPS, Traffic

1. INTRODUCTION

Automotive navigation systems enable the effective delivery and presentation of fine-grained traffic information to drivers and have

[Hoh et al., MobiSys'08]

PRIVACY ATTACK ON THE MOBILE CENTURY DATASET

POLYTECHNIQUE MONTRÉAL

A Model-based Framework for Privacy and Security Analysis of Traffic Monitoring Systems

Edward S. Canepa, *Member, IEEE,* and Christian G. Claudel, *Member, IEEE*

Abstract

Most large scale traffic information systems rely on fixed sensors (e.g. loop detectors, cameras) and user generated data, the latter in the form of GPS traces sent by smartphones or GPS devices onboard vehicles. While this type of data is relatively inexpensive to gather, it can pose multiple security and privacy risks, even if the location tracks are anonymous. In particular, creating bogus location tracks and sending them to the system is relatively easy. This bogus data could perturb traffic flow estimates, and disrupt the transportation system whenever these estimates are used for actuation. Another issue could be the possibility for an attacker to infer user location tracks from anonymous location data, which affects users privacy. In this article, we propose a new framework for solving a variety of privacy and cybersecurity problems arising in transportation systems. The state of traffic is modeled by the Lighthill-Whitham-Richards traffic flow model, which is a first order scalar conservation law with a concave flux function. Given a set of traffic flow data, we show that the constraints resulting from this partial differential equation are mixed integer linear inequalities for some decision variable. The resulting framework is very flexible, and can in particular be used to detect spoofing attacks in real time, or to carry out attacks on location tracks. Numerical implementations are performed on experimental data from the *Mobile Century* experiment.

I. Introduction

The convergence of mobile sensing, communication and computing has led to the rise of a

[Canepa and Claudel, HiCoNS'13]

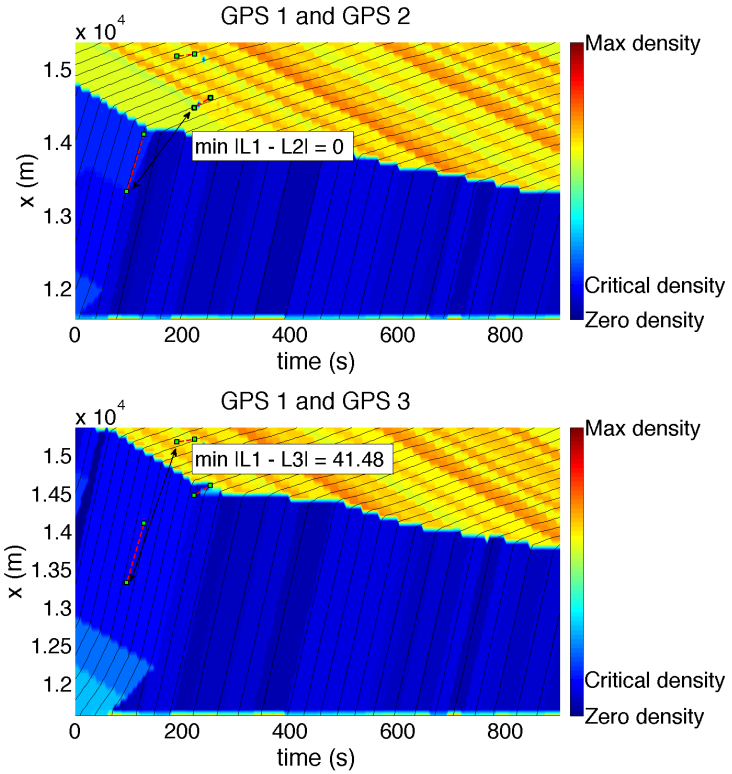

Figure 3: Example of reidentification. The corresponding scenario is decribed in Figure 2. Top: Solution to the reidentification problem (15), with an objective $|L_1 - L_2|$. Bottom: Solution to the same problem with an objective $|L_1 - L_3|$. A nonzero optimum means that both tracks cannot be generated by the same vehicle, according to both the model and the available data.

PRIVACY IS NOT ANONYMITY

POLYTECHNIQUE MONTRÉAL

- Privacy breaches generally due to existence of side information
  - Mass. GIC medical DB linked with voter registration DB [Sweeney, 1997]
  - Netflix prize with IMDB [Narayanan & Shmatikov, 2008]
  - Individual online transactions with changes in public recommendation systems [Calandrino et al., 2011]
  - "Anonymity" in location based services is very hard to provide!
- Very hard to know what the adversary knows, or might know in the future

NYC TAXI DATASET

POLYTECHNIQUE MONTRÉAL

Riding with the Stars: Passenger Privacy in the NYC Taxicab Dataset

[Anthony Tockar]

SEPTEMBER 15, 2014 BY ATOCKAR 56 COMMENTS

In my previous post, Differential Privacy: The Basics, I provided an introduction to differential privacy by exploring its definition and discussing its relevance in the broader context of public data release. In this post, I shall demonstrate how easily privacy can be breached and then counter this by showing how differential privacy can protect against this attack. I will also present a few other examples of differentially private queries.

The Data

There has been a lot of online comment recently about a dataset released by the New York City Taxi and Limousine Commission. It contains details about every taxi ride (yellow cabs) in New York in 2013, including the pickup and drop off times, locations, fare and tip amounts, as well

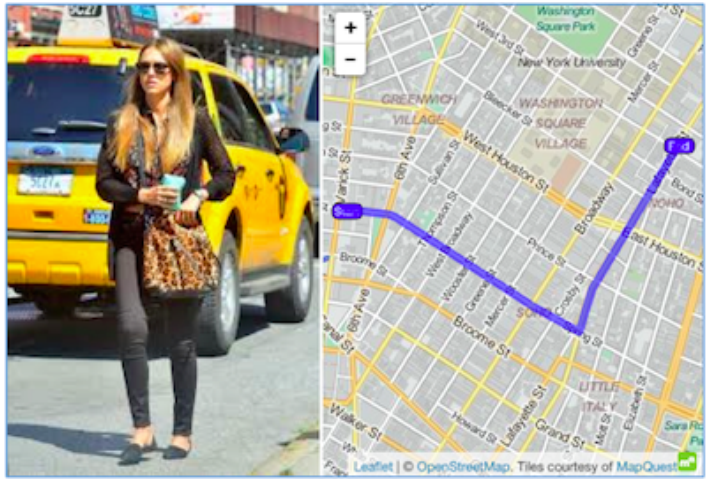

Jessica Alba (Click to Explore)

Anonymization of Location Data Does Not Work: A Large-Scale Measurement Study

Hui Zang
Sprint
1 Adrian Ct
Burlingame, CA 94010, USA
hui.zang@sprint.com

Jean Bolot*
Technicolor
735 Emerson St
Palo Alto, CA 94301, USA
jean.bolot@technicolor.com

ABSTRACT

We examine a very large-scale data set of more than 30 billion call records made by 25 million cell phone users across all 50 states of the US and attempt to determine to what extent anonymized location data can reveal private user in

Categories and Subject Descriptors

C.2.m [**Computer-Communication Networks**]: [Miscellaneous]; H.4 [**Information Systems Applications**]: Miscellaneous

PRIVACY THREATS AND OBJECTIVES

POLYTECHNIQUE MONTRÉAL

- The data collection process (communications, third party aggregator) is not the only risk
- Risks associated to data collection can typically be addressed by anonymization, cryptographic means and other system design tricks (encrypted communications, proxy servers, homomorphic encryption if the aggregator is not trusted, etc.)
- Still, a useful system cannot avoid a fundamental disclosure of information (ex: real-time traffic density map, travel time estimates, etc.). Want to make sure this does not leak too much information about individuals
- Positive feedback loop: offering privacy guarantees encourages user adoption which helps with performance and privacy

PRIVACY CHALLENGES

POLYTECHNIQUE MONTRÉAL

- What is privacy, formally?
- Developing practical privacy-preserving mechanisms
- What are the tradeoffs between privacy and utility for a given system?
- Here in particular, privacy-preserving real-time information processing
  - Estimation / monitoring
  - Control / decision making

TRAFFIC INFORMATION WITH FORMAL PRIVACY GUARANTEES ?

POLYTECHNIQUE MONTRÉAL

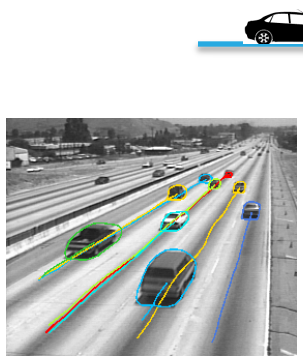

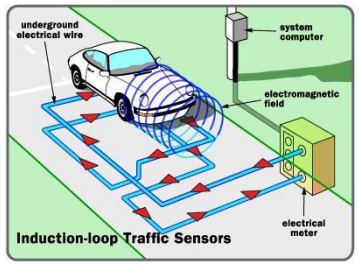

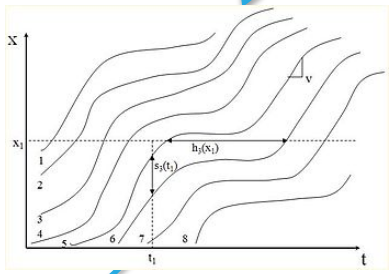

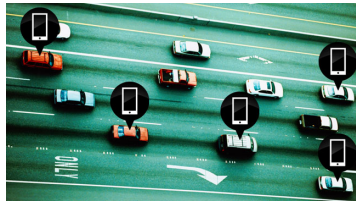

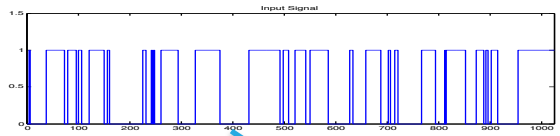

$$x_{i,t+1} = \begin{bmatrix} 1 & T_s \\ 0 & 1 \end{bmatrix} x_{i,t} + \sigma_{i1} \begin{bmatrix} T_s^2/2 & 0 \\ T_s & 0 \end{bmatrix} w_{i,t},$$

$$\frac{\partial \rho}{\partial t} + \frac{\partial q}{\partial x} = 0$$

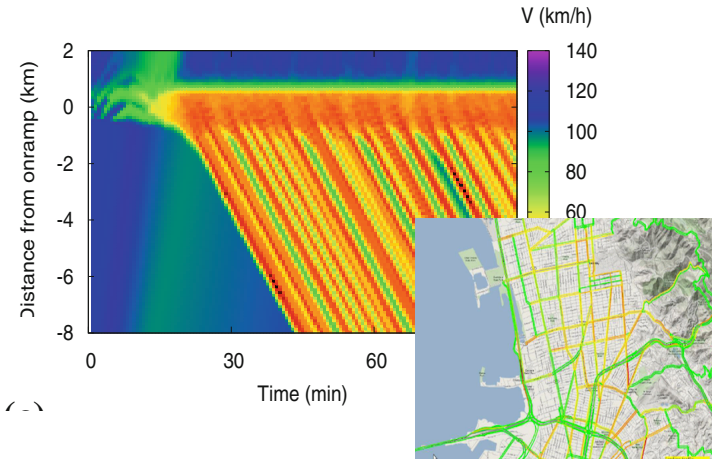

Need privacy-preserving data-assimilation procedures for heterogeneous data sources and dynamic models

OUTLINE

POLYTECHNIQUE MONTRÉAL

1. Differential Privacy

2. Example: Differentially Private Kalman Filtering

3. Differentially Private Traffic Estimation with Hydrodynamic Models

OUTLINE

POLYTECHNIQUE MONTRÉAL

1. Differential Privacy

2. Example: Differentially Private Kalman Filtering

3. Differentially Private Traffic Estimation with Hydrodynamic Models

DEFINITIONS OF PRIVACY

POLYTECHNIQUE MONTRÉAL

- In statistics
  - Risk associated to small entries in table counts
  - Techniques: cell suppression, clustering, perturbations…
  - [Duncan & Lambert, 1986], [Reiter, 2005]
- In information theory
  - Lower bound conditional entropy H(private info|public info) while still publishing useful information
  - [Sankar et al., 2010], [Venkitasubramaniam, 2013]
- In computer science
  - K-anonymity [Sweeney, 1998] (used by [Hoh et al.], but per speed sample)
  - Differential privacy [Dwork et al., 2006]

Preserving Privacy requires some form of obfuscation of the data published

DIFFERENTIAL PRIVACY, INFORMALLY

POLYTECHNIQUE MONTRÉAL

- Set-up:

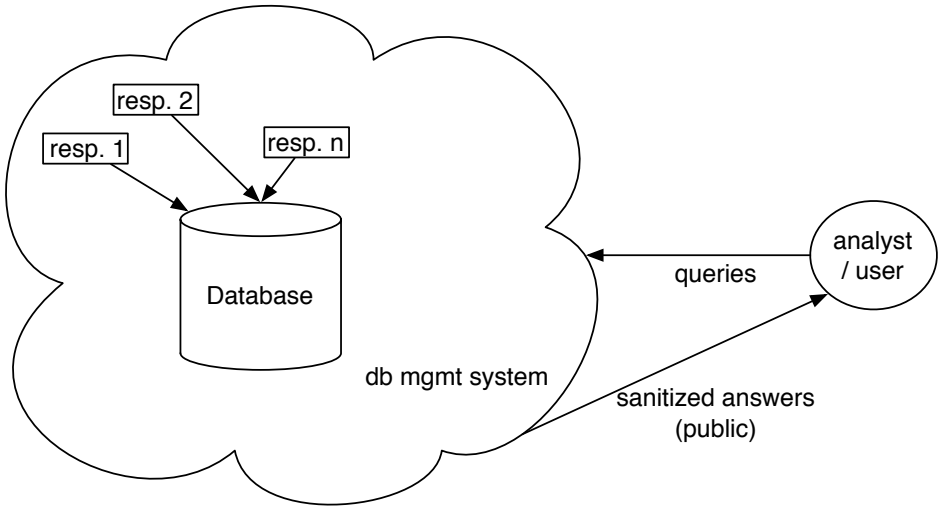

- Key Idea:
  - A differentially private mechanism randomly perturbs its answer to a query so that the output distribution over answers does not vary much if any given individual changes its data (or even participates or not)
  - Hard to infer if the specific data of any individual was used or not to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

DIFFERENTIAL PRIVACY, INFORMALLY

POLYTECHNIQUE MONTRÉAL

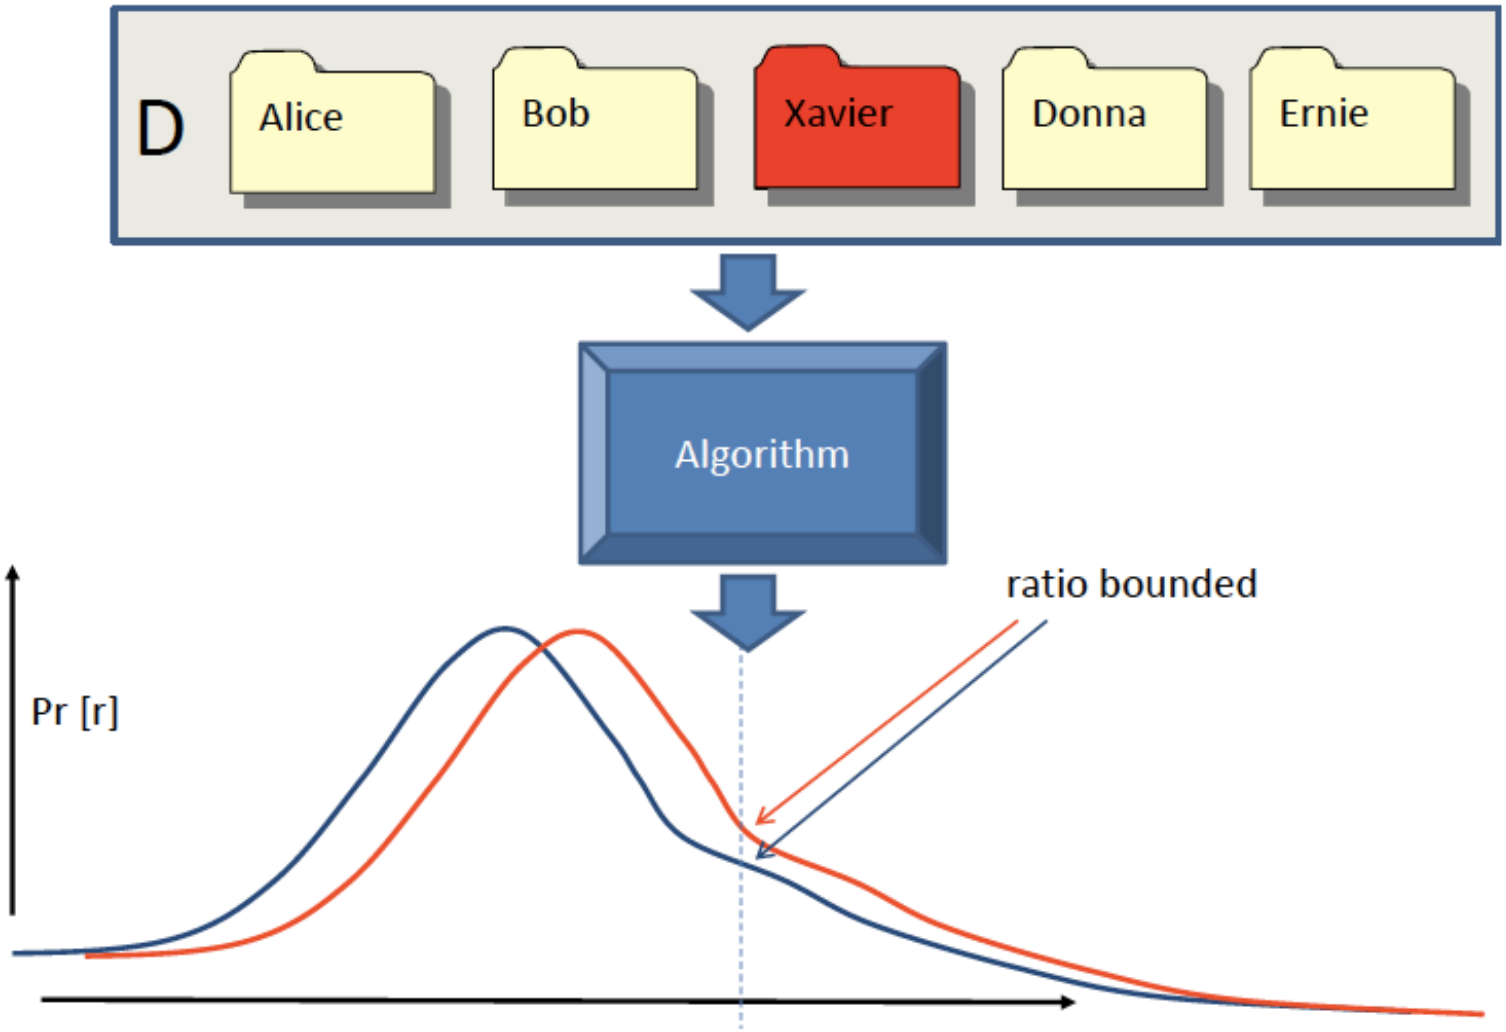

DIFFERENTIAL PRIVACY, FORMALLY

POLYTECHNIQUE MONTRÉAL

- Set-up:

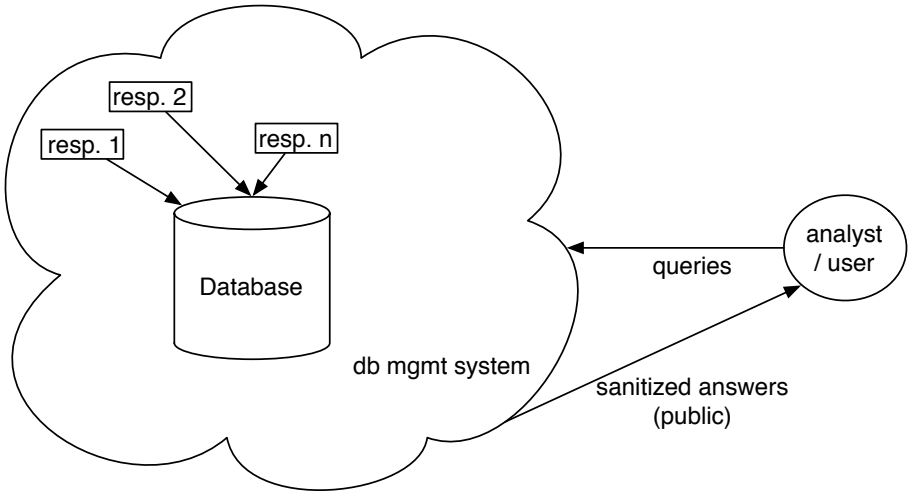

- Adj(d,d') a **symmetric** relation on the set D of datasets
- Adjacent datasets differ by the data of a single individual
- A mechanism $M : D \times \Omega \to (R, \mathcal{M})$ is (ε,δ)-DP if for all sets $S \in \mathcal{M}$ and all databases d, d' s.t. Adj(d,d'), we have

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

(N.B.: we can interchange d and d' in (*) by symmetry of Adj)

DIFFERENTIAL PRIVACY, FORMALLY

POLYTECHNIQUE MONTRÉAL

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

- Constant ε is typically small (i.e. ~ 0.1, ln(2), …)
  - multiplicative error
- Constant δ is very small (i.e. ~0.01, 0.05)
  - additive error
- If δ=0 then we have (ε,0)-DP or simply ε-DP
- Privacy definition depends on adjacency relation: characterizes the pairs of datasets that we want to make hard to distinguish

AN OPERATIONAL INTERPRETATION OF DIFFERENTIAL PRIVACY

POLYTECHNIQUE MONTRÉAL

- Equivalent definition in terms of hypothesis testing

  [Wasserman and Zhu, 2010], [Oh and Viswanath, 2013]

  - H0: the dataset is d
  - H1: the dataset is d', adjacent to d
  - S a rejection region to design
  - Pba of false alarm (H0 true but rejected): $P_{fa}^{S,d,d'} = \mathbb{P}(M(d) \in S)$
  - Pba of missed detection (H0 false but retained): $P_{md}^{S,d,d'} = \mathbb{P}(M(d') \in \bar{S})$

- Theorem: M is differentially private iff for all Adj(d,d') and all S:

$$P_{fa}^{S,d,d'} + e^{\epsilon} P_{md}^{S,d,d'} \geq 1 - \delta$$

$$\text{and } e^{\epsilon} P_{fa}^{S,d,d'} + P_{md}^{S,d,d'} \geq 1 - \delta$$

It is impossible to get both small P_{fa} et P_{md} from data obtained via a DP mechanism

RESILIENCE TO POST-PROCESSING

POLYTECHNIQUE MONTRÉAL

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

- f as the adversaries : models arbitrary auxiliary or side 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, **as long as it does not re-access the database**

OUTLINE

POLYTECHNIQUE MONTRÉAL

1. Differential Privacy

2. Example: Differentially Private Kalman Filtering

3. Differentially Private Traffic Estimation with Hydrodynamic Models

DIFFERENTIALLY PRIVATE FILTERING

POLYTECHNIQUE MONTRÉAL

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*—Estimation, filtering, Kalman filtering, privacy.**

I. Introduction

A rapidly growing number of applications require 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.

Manuscript received September 04, 2012; revised April 16, 2013; accepted August 22, 2013. Date of publication September 23, 2013; date of current version January 21, 2014. This work was supported in part by the Natural Sciences and Engineering Research Council under Grant RGPIN-435905-13 and in part by TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA. Preliminary versions of this paper appeared in [1] and [2]. Recommended by Associate Editor L. Schenato.

J. Le Ny is with the Department of Electrical Engineering, Ecole Polytechnique de Montreal, Montreal, QC H3T 1J4, Canada (e-mail: jerome.le-ny@polymtl.ca).

G. J. Pappas is with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA (e-mail: pappasg@seas.upenn.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TAC.2013.2283096

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 [3]. 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 [4], [5] 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 [6] 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 [6], [7]. 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* [8], 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 [9].

Most work related to privacy is concerned with the analysis of static databases [8], [10]–[12], 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 [13]–[15]. This work forms the basis for the scenario studied in Section VI, and is discussed in more details in Section VI-C. 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 [16]. In this case, the input data to protect consists of the initial values of the participants and is

TWO BASIC DIFFERENTIALLY PRIVATE MECHANISMS

POLYTECHNIQUE MONTRÉAL

- Dataset d contains n numbers (ex: salaries): $d_1, \ldots, d_n$
- Query q, ex: $q(d) = \frac{1}{n} \sum_{i=1}^{n} d_i$

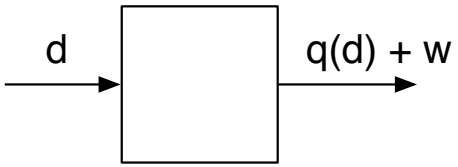

- L_p-sensitivity of a query q: $R \to R^m$

$$\Delta_p q := \max_{d,d':\mathrm{Adj}(d,d')} \|q(d) - q(d')\|_p$$

- Laplace mechanism $q(d) + \mathrm{Lap}(\Delta_1 q/\epsilon)$ is ε-DP

$\mathrm{Lap(b)}$ pdf: $\frac{1}{2b} e^{-|x|/b}$, std. dev. $\sqrt{2}b$, iid on each component

- Gaussian mechanism $q(d) + \kappa_{\delta,\epsilon}\ \Delta_2 q\ \mathcal{N}(0, I_m)$ is (ε,δ)-DP

$$\kappa_{\delta,\epsilon} \in O(\sqrt{\ln(1/\delta)}/\epsilon)$$

GENERALIZATION TO DISCRETE-TIME DYNAMIC SYSTEMS

POLYTECHNIQUE MONTRÉAL

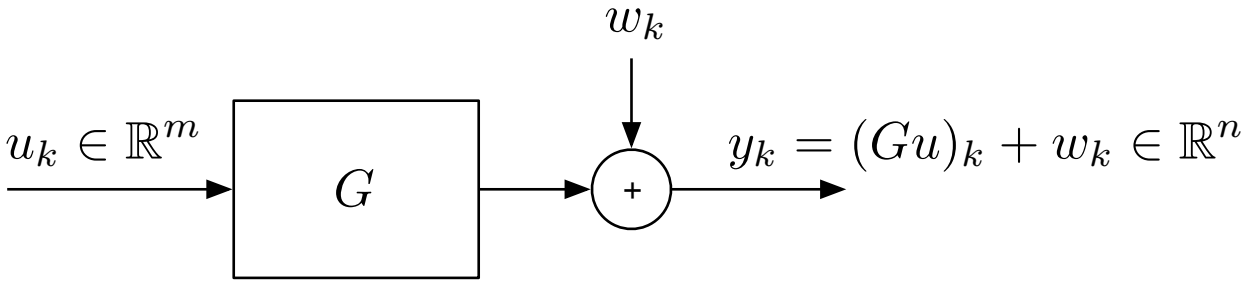

- Adjacency Adj(u,u') between input signals
- Query ↔ system/filter G
- L_p-sensitivity of G: $\Delta_p G = \max_{\text{Adj(u,u')}} \|Gu - Gu'\|_p$

$$\text{where } \|x\|_p = \left(\sum_{k=0}^{\infty} |x_k|_p^p \right)^{1/p}, \; |v|_p = \left(\sum_{i=1}^{n} |v_i|^p \right)^{1/p}$$

- Then the signal $y = Gu + w$ is
  - ε-DP if w is Laplace white noise vector with $w_{k,i} \sim \text{Lap}(\Delta_1 G/\varepsilon)$
  - (ε,δ)-DP if w is Gaussian white noise with $\Sigma = (\kappa_{\delta,\epsilon} \Delta_2 G)^2 I_n$

EXAMPLE

POLYTECHNIQUE MONTRÉAL

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

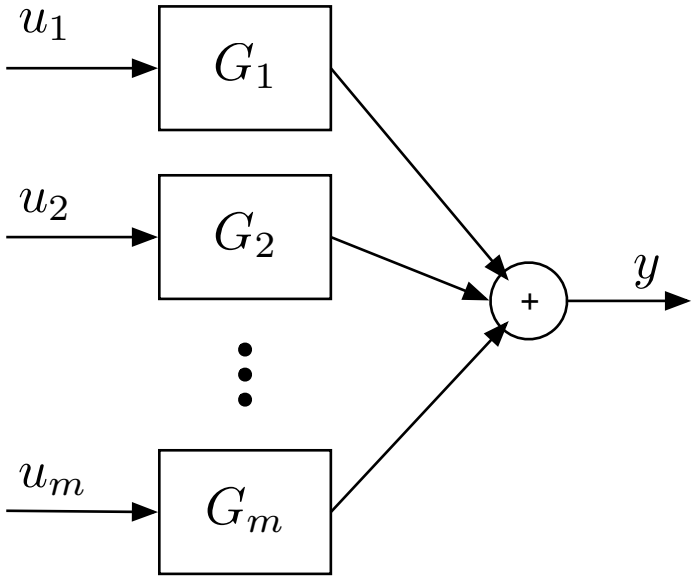

- Adjacency relation
  - n input signals, l per user

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

EXAMPLE (CONT.)

POLYTECHNIQUE MONTRÉAL

$$y = \sum_{i=1}^{m} G_i u_i$$

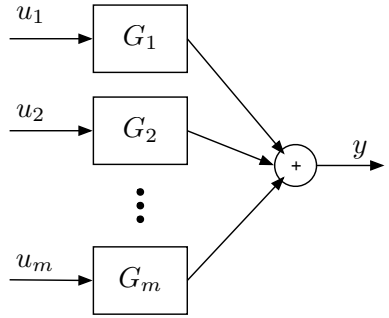

$\mathrm{Adj}^b(u, u')$ iff for some i, $\|u_i - u'_i\|_2 \leq B$, and $u_j = u'_j$ for all $j \neq i$.

- For this adjacency relation:

$$\Delta_2 G = \max_i \left\{ \sup_{\|u_i - u'_i\|_2 \leq B} \|G_i u_i - G_i u'_i\|_2 \right\}$$

$$\Delta_2 G \leq \max_i \{\gamma_2(G_i)\}\ B$$

- $\gamma_2(G_i)$ is the $l_2 \longrightarrow l_2$ incremental gain of G_i
  - $\|G_i\|_\infty$ if G_i is linear

- $y = Gu + w$ with w a WGN s.t. $\Sigma = \sigma^2 I_n,\ \sigma = \kappa_{\delta,\epsilon}\, B\, \max_i \{\gamma_2(G_i)\}$ is (ε,δ)-DP
  - Output perturbation mechanism

INPUT AND OUTPUT PERTURBATION

POLYTECHNIQUE MONTRÉAL

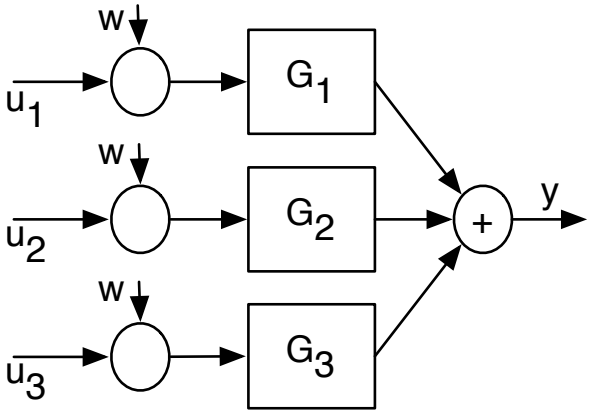

$$w \sim \mathcal{N}(0, \sigma^2)$$

$$\sigma = \kappa_{\delta,\epsilon} B$$

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

Input perturbation mechanism

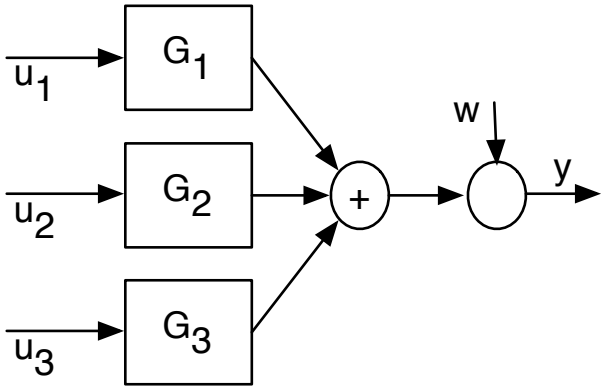

$$w \sim \mathcal{N}(0, \sigma^2)$$

$$\sigma = \kappa_{\delta,\epsilon} B \max_{1 \le i \le m} \{\|G_i\|_\infty\}$$

$$MSE = \sigma^2$$

Output perturbation mechanism

INPUT AND OUTPUT PERTURBATION: EXAMPLE

POLYTECHNIQUE MONTRÉAL

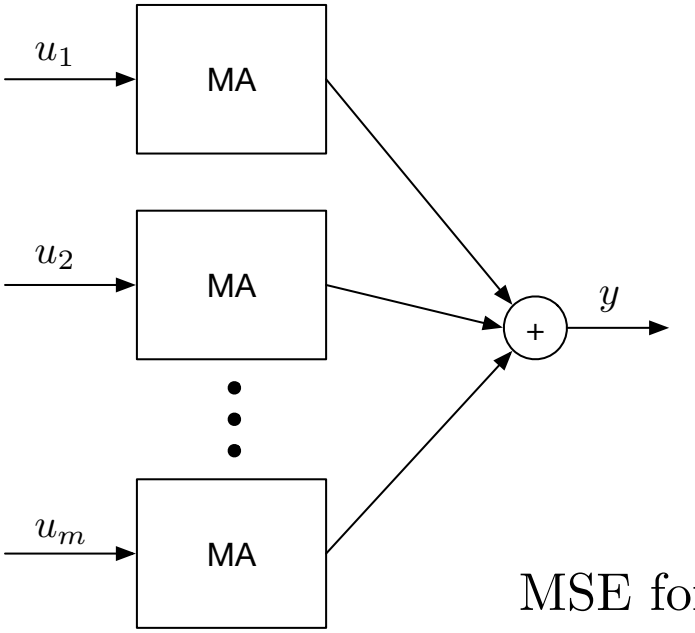

m users

Average over l steps

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

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

MSE for input perturbation: $\kappa_{\delta,\epsilon}^2 B^2 \frac{m}{l}$

MSE for output perturbation: $\kappa_{\delta,\epsilon}^2 B^2$

Output perturbation better than input perturbation when $m > l$

A MORE GENERAL APPROXIMATION ARCHITECTURE

POLYTECHNIQUE MONTRÉAL

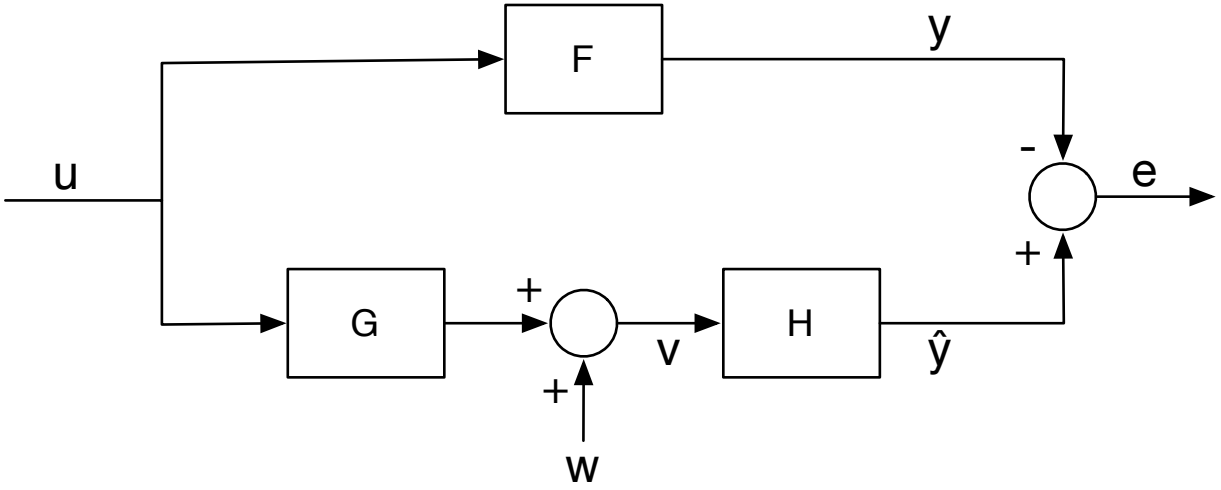

- Approximate F by a differentially private system
- Add noise w proportional to sensitivity of pre-filter G for privacy (sensor/channel design)
- Post-processing with H: estimation (equalization) problem
- Design Variables: G, H (w fixed by ΔG)
- Related work:
  - [Li and Miklau 2010], [Tanaka, Kim, Parrilo and Mitter 2015]

MODEL BASED ESTIMATION: AVERAGE SPEED MONITORING EXAMPLE

POLYTECHNIQUE MONTRÉAL

- Often we have a (public) model of the input signals
- Ex: Average velocity estimation using individual location traces

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

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

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

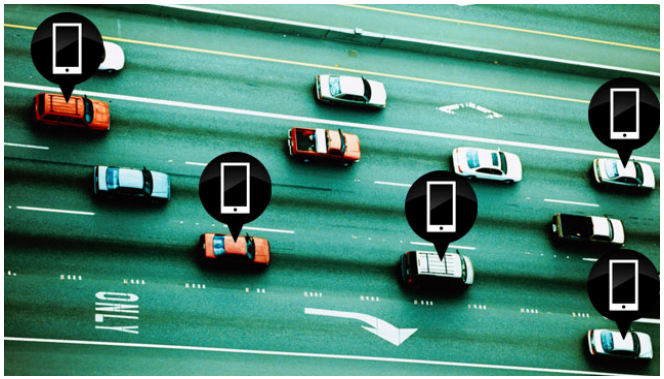

- Preserve privacy of measured location signals $y_{i,t}$ or the actual 2D trajectories (position & speed) $x_{i,t}$

DIFFERENTIALLY PRIVATE KALMAN FILTERING

POLYTECHNIQUE MONTRÉAL

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

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

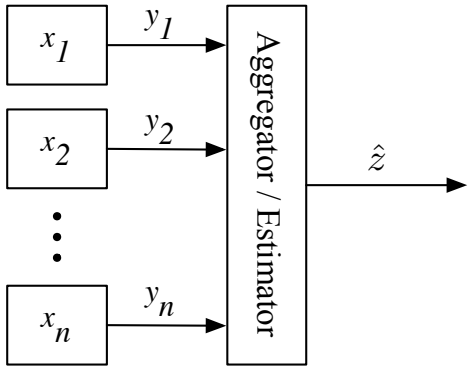

- Estimation objective: $z_t = \sum_{i=1}^{n} L_i x_{i,t}$ $\qquad \min_{\hat{z}} \lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}\left[\|z_t - \hat{z}_t\|_2^2\right]$
- Adjacency relation:

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

  - Hard to distinguish between two sufficiently close state trajectories of a user

KALMAN FILTERING
DIFFERENTIALLY PRIVATE MECHANISMS

POLYTECHNIQUE MONTRÉAL

- Input or output perturbation to the standard Kalman filter

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

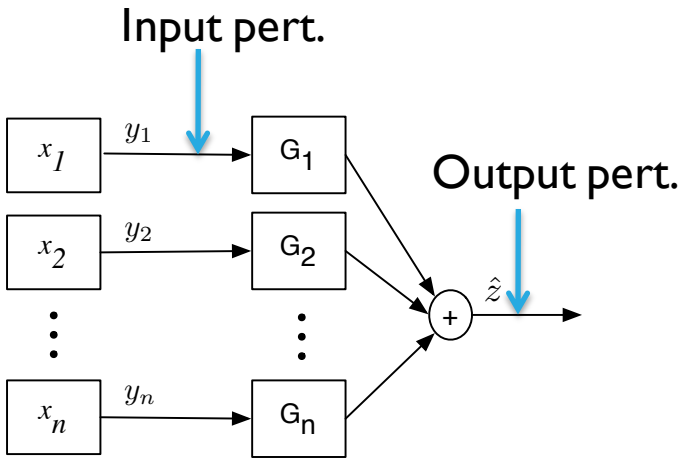

- Adjacency relation:

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

- For the input perturbation scheme, can take into account the additional privacy-preserving noise in the redesign of the KF
  - But this slows down convergence

FILTER REDESIGN FOR OUTPUT PERTURBATION SCHEME

POLYTECHNIQUE MONTRÉAL

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

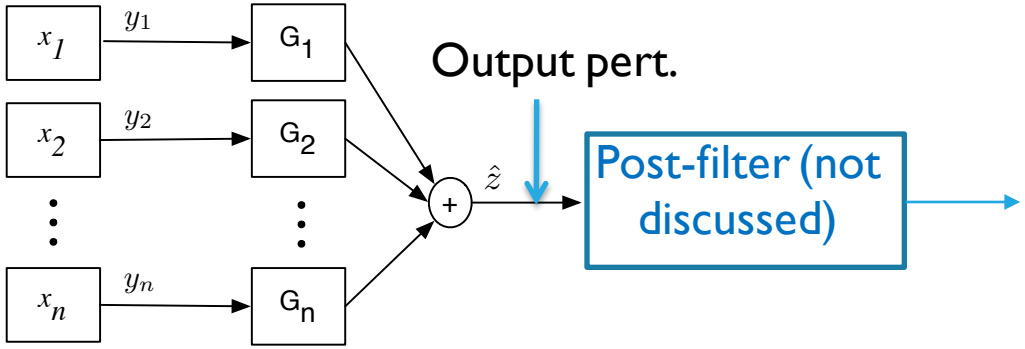

- For the output perturbation scheme, can redesign the filter to trade-off the estimation error and the H_∞ norm of the filter
  - Overall MSE is

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

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

AVERAGE SPEED MONITORING EXAMPLE

POLYTECHNIQUE MONTRÉAL

- Average velocity estimation using individual location traces

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

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

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

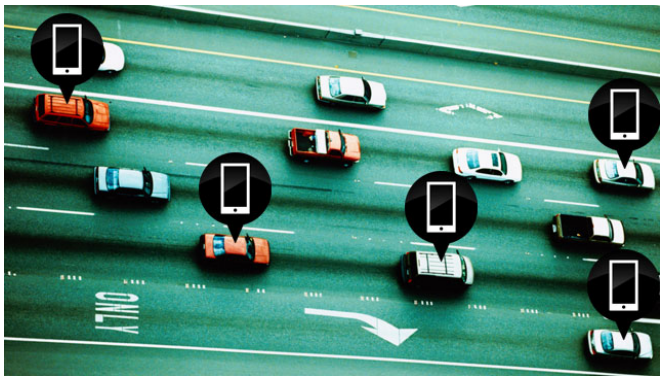

$$\begin{aligned} \rho &= 100\text{m} \\ \epsilon &= \ln 3 \\ \delta &= 0.05, \\ T_s &= 1\text{s} \\ \sigma_{i1} &= \sigma_{i2} = 1 \\ n &= 200 \end{aligned}$$

SPEED MONITORING
INPUT VERSUS OUTPUT ARCHITECTURE

POLYTECHNIQUE MONTRÉAL

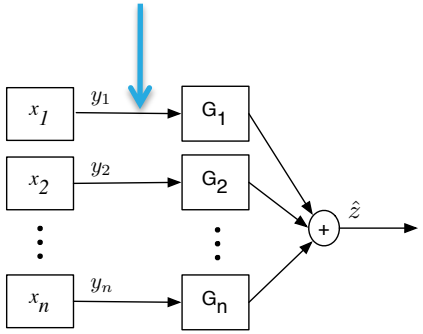

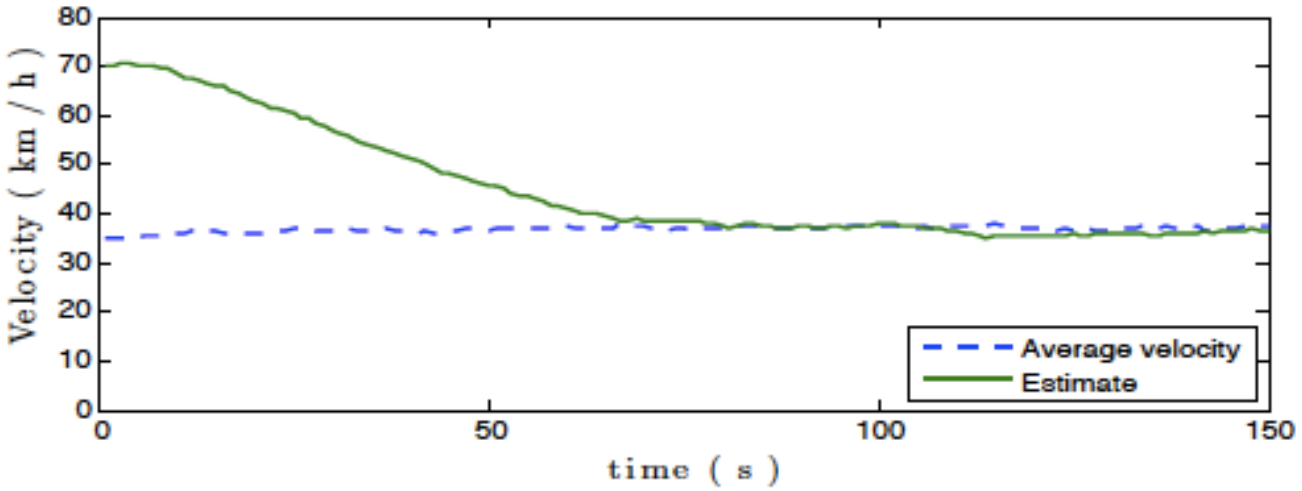

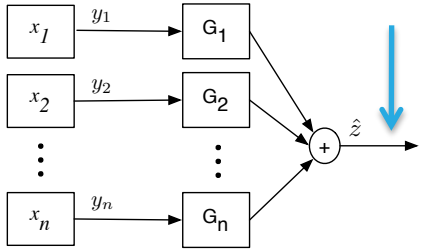

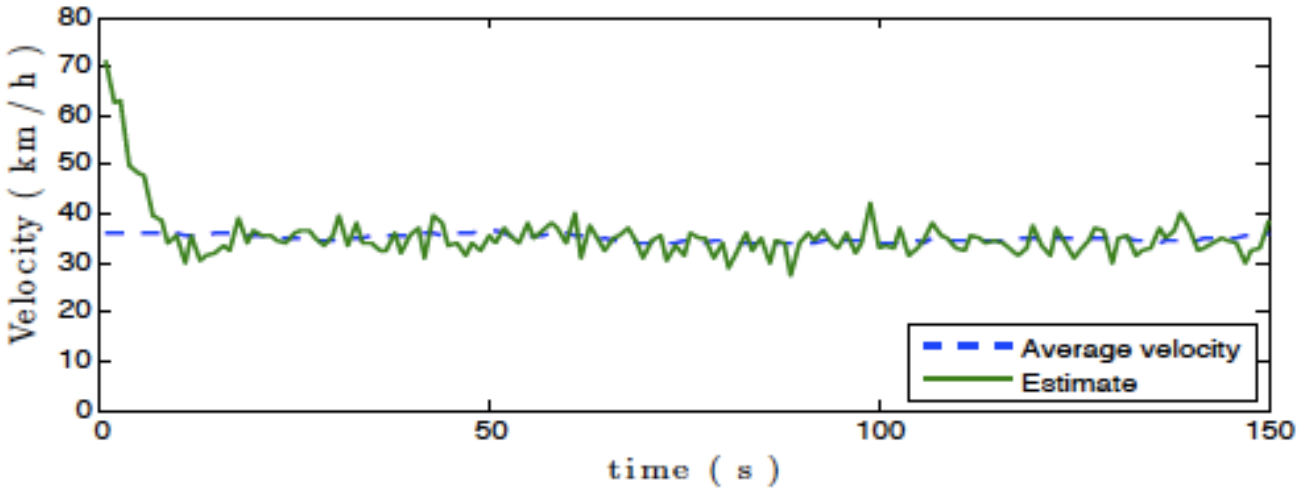

SPEED MONITORING UTILITY/PRIVACY TRADEOFF

POLYTECHNIQUE MONTRÉAL

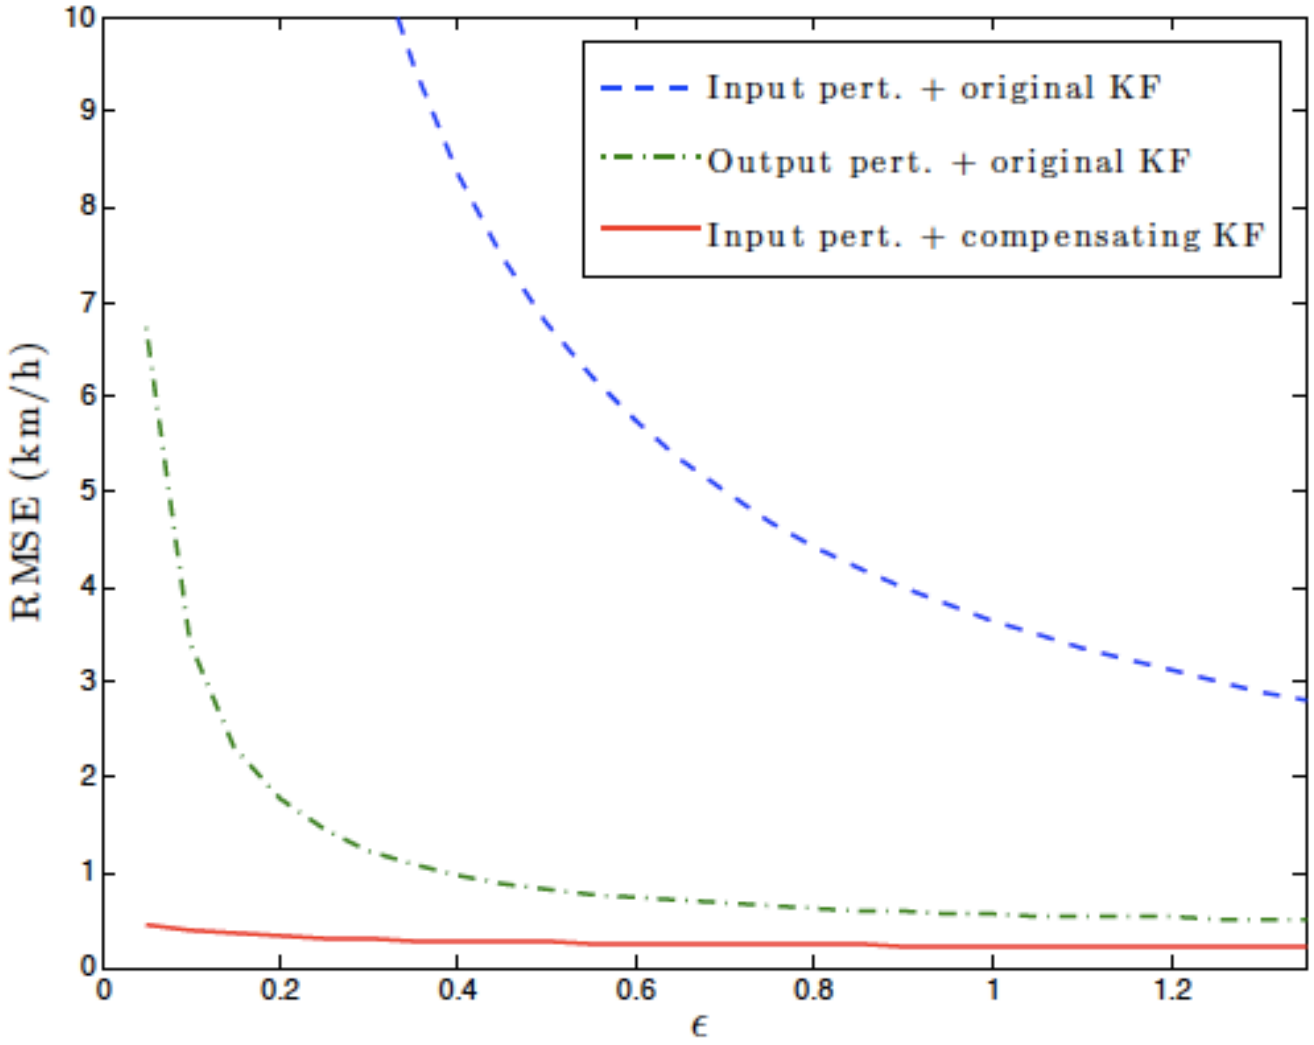

SPEED MONITORING
CONVERGENCE/PRIVACY TRADEOFF

POLYTECHNIQUE MONTRÉAL

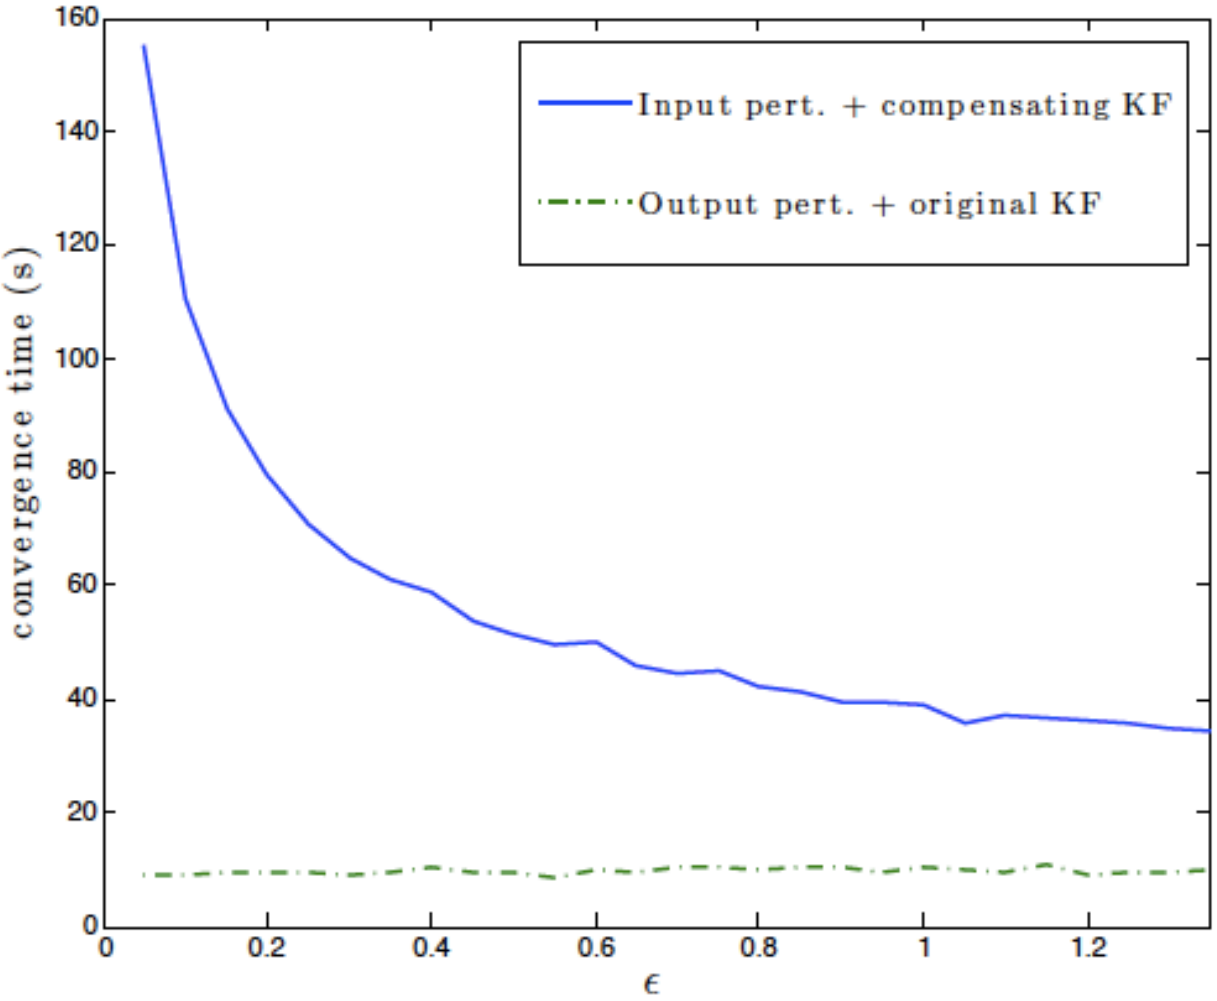

OUTLINE

POLYTECHNIQUE MONTRÉAL

1. Differential Privacy
2. Example: Differentially Private Kalman Filtering
3. Differentially Private Traffic Estimation with Hydrodynamic Models

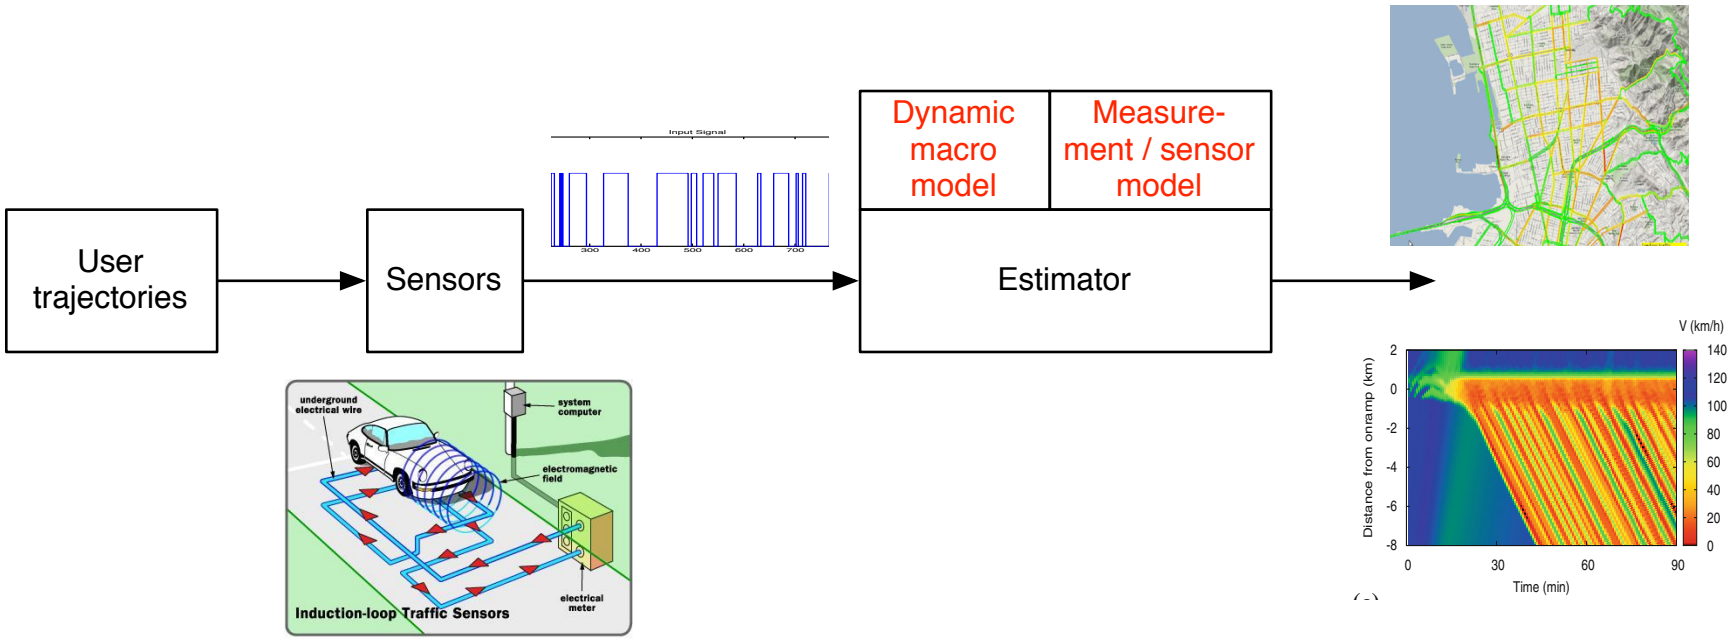

OUTLINE

POLYTECHNIQUE MONTRÉAL

Differentially Private Traffic Estimation with Hydrodynamic Models

A. Dynamic Model for Density Estimation

B. Measurement models and sanitization of observations (input perturbation mechanisms)

C. EKF-based designs

MACROSCOPIC TRAFFIC FLOW MODEL

POLYTECHNIQUE MONTRÉAL

- Based on a conservation law (for vehicles)

$$\frac{d}{dt}\int_{x_1}^{x_2} \rho_{\text{tot}}(x,t)dx = q_{\text{tot}}(x_1,t) - q_{\text{tot}}(x_2,t), \quad \forall x_1, x_2, t,$$

  - Flow = traffic velocity x density $q = v\rho$

- Numerical discretization gives finite dimensional dynamic model for ρ

$$\rho_{k+1}^i = \rho_k^i + \frac{\tau}{L_i}\left(\frac{\lambda^{i-1}}{\lambda^i} f_k^{i-1} - f_k^i\right), \quad \text{for } i = 1, \ldots I.$$

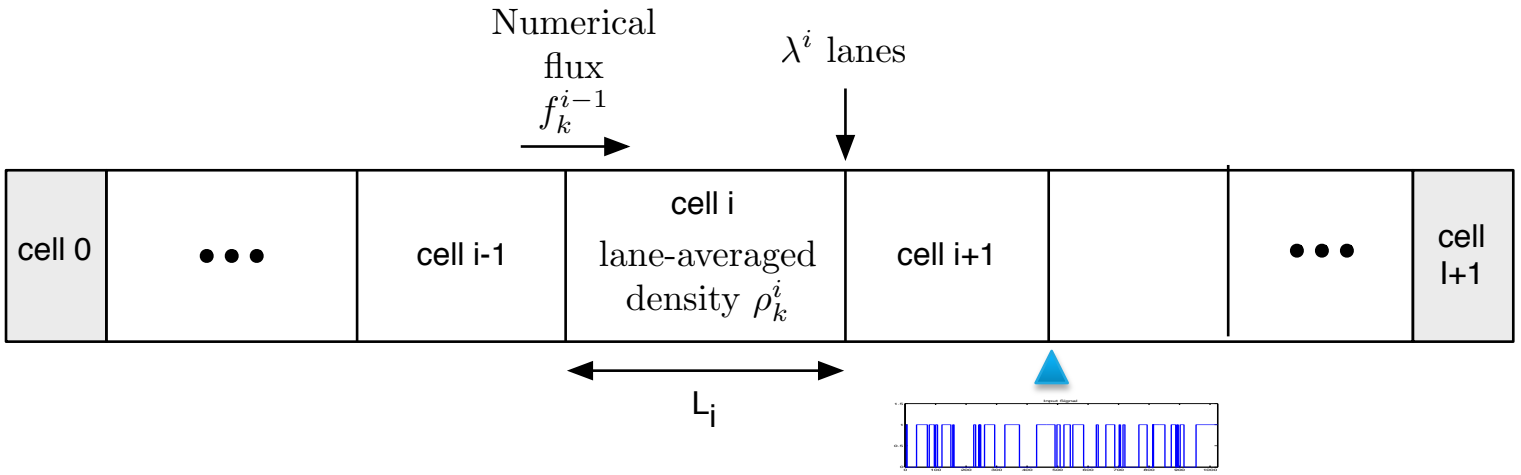

MACROSCOPIC FLOW MODEL: NUMERICAL FLUX

POLYTECHNIQUE MONTRÉAL

- First-order models (LWR = Lighthill, Whitham, Richards):
  - Assume a static relation between traffic flow or velocity and density = "fundamental diagram" $q(\rho)$
  - Ex: Cell-transmission model (CTM), triangular f. d.

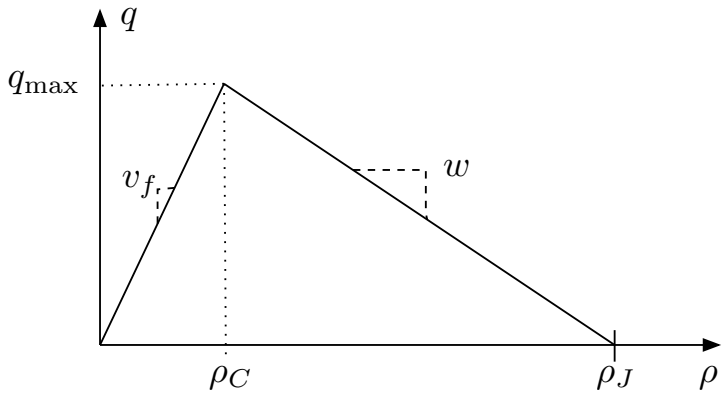

- Godunov numerical method with triangular f.d.

$$f_k^i = F(\rho_k^i, \rho_k^{i+1}) := \min\{S(\rho_k^i), R(\rho_k^{i+1})\}$$
$$= \min\left\{\rho_k^i v_f, q_{\max}, w(\rho_J - \rho_k^{i+1})\right\}.$$

FUNDAMENTAL DIAGRAM CALIBRATION

POLYTECHNIQUE MONTRÉAL

[Mobile Millenium dataset]

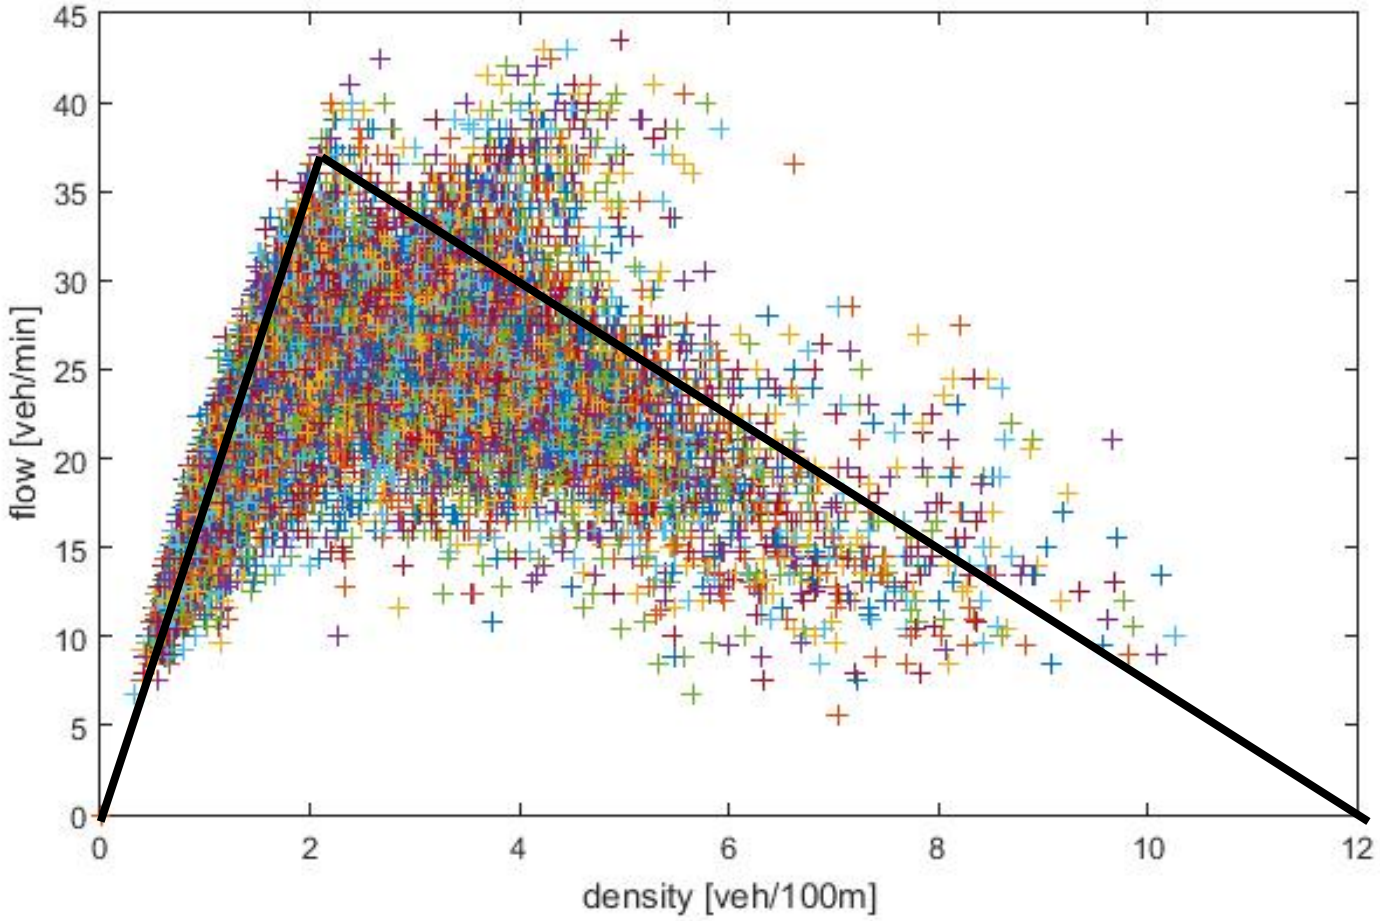

STOCHASTIC DYNAMIC MACROSCOPIC MODEL

POLYTECHNIQUE MONTRÉAL

- Set-up the dynamic model for estimation (here with an Extended Kalman Filter)

- Introduce random noise to account for imperfections of the deterministic model

$$\rho_{k+1}^i = \rho_k^i + \frac{\tau}{L_i}\left(\frac{\lambda^{i-1}}{\lambda^i}F(\rho_k^{i-1},\rho_k^i) - F(\rho_k^i,\rho_k^{i+1})\right) + \xi_k^i$$

- White noise driven boundary conditions
  - State of the ghost cells to estimate as well in general

$$\rho_{k+1}^0 = \rho_k^0 + \xi_k^0, \quad \rho_{k+1}^{I+1} = \rho_k^{I+1} + \xi_k^{I+1},$$

OUTLINE

POLYTECHNIQUE MONTRÉAL

Differentially Private Traffic Estimation with Hydrodynamic Models

A. Dynamic Model for Density Estimation

B. Measurement models and sanitization of observations (input perturbation mechanisms)

C. EKF-based designs

MEASUREMENT MODELS

POLYTECHNIQUE MONTRÉAL

- Single-loop detectors report every T seconds for each lane:
  - Vehicle count c(t) for the last period
  - Percentage occupancy o(t) for the last period
- → Pseudo-measurements (j = lane #):

$$q_j(t) \approx \frac{c_j(t)}{T}, \quad \rho_j(t) \approx \frac{o_j(t)}{g}, \quad v_j(t) \approx g\frac{c_j(t)}{o_j(t)T}$$

- g = g-factor (average vehicle length), requires calibration (in fact g(x,t))
- Various measurement models possible (i = sensor location #)

$$\phi_k^i := \frac{1}{T\lambda^i}\sum_{j=1}^{\lambda^i} c_{j,k}^i = F(\rho_k^i, \rho_k^{i+1}) + \nu_k^i$$ (flow measurement model)

$$y_k^i := \frac{1}{g\lambda^i}\sum_{j=1}^{\lambda^i} o_{j,k}^i = \rho_k^i + \nu_k^i$$ (density measurement model #1, problematic for privacy)

MEASUREMENT MODELS: GPS DATA

POLYTECHNIQUE MONTRÉAL

- GPS data can be used to simulate static sensors, similarly to the virtual trip lanes of Hoh. et al.
- Allows the same algorithms to be used, although now we can have speed measurements from GPS
- Sensibility analysis is similar to that of o (or o/c)
- Other schemes might be possible but require a careful analysis of how to spend the privacy budget

SANITIZATION OF FLOW MEASUREMENTS

POLYTECHNIQUE MONTRÉAL

- Sensitivity computation for flow measurements
- 2 sets of trajectories differing by the trajectory of a single user

$$\|\boldsymbol{\phi} - \tilde{\boldsymbol{\phi}}\|_2^2 = \sum_{k=0}^{\infty}\sum_{i=1}^{M} |\phi_k^i - \tilde{\phi}_k^i|^2 = \sum_{i=1}^{M}\sum_{k=0}^{\infty} |\phi_k^i - \tilde{\phi}_k^i|^2$$

$$\sum_{k=0}^{\infty} |\phi_k^i - \tilde{\phi}_k^i|^2 = \frac{1}{T^2(\lambda^i)^2} \sum_{k=0}^{\infty} \left| \sum_{j=1}^{\lambda^i} (c_{j,k}^i - \tilde{c}_{j,k}^i) \right|^2 \leq \frac{2}{T^2(\lambda^i)^2}$$

$$\Rightarrow \|\boldsymbol{\phi} - \tilde{\boldsymbol{\phi}}\|_2 \leq \frac{\sqrt{2}}{T} \sqrt{\sum_{i=1}^{M} \frac{1}{(\lambda^i)^2}} =: \Delta_f$$

- → add Gaussian white noise with variance $(\kappa(\delta,\epsilon)\ \Delta_f)^2$ to each single-loop detector count signal to get a set of (ε,δ)-differentially private signals
- Noise variance grows linearly with number of sensor locations (trip lines). But valid for arbitrary variation in a user's trajectory (strong)

SANITIZATION OF OCCUPANCY MEASUREMENTS

POLYTECHNIQUE MONTRÉAL

$$y_k^i := \frac{1}{g\lambda^i} \sum_{j=1}^{\lambda^i} o_{j,k}^i = \rho_k^i + \nu_k^i$$

- To use directly the occupancy measurements, let us weaken the guarantee by strengthening the conditions for adjacency. Impose a limit on the impact that one trajectory can have on the occupancy measurements
- Adj(d,d') if for each sensor location i, there is at most two pairs (k, j) or (time, lane) such that

$$\frac{\left|o_{j,k}^i(d) - o_{j,k}^i(d')\right|}{\min\left(o_{j,k}^i(d),\ o_{j,k}^i\ (d')\right)} \leq \gamma$$

- Bound relative variations of occupancy by parameter ɣ
- Then $\tilde{o} = \lambda o$ with λ log-N(0,$\kappa_{\delta,\epsilon}\sqrt{2M}$) for M sensors is (ϵ, δ)-DP (note: sign-preserving mechanism)
- How to set ɣ ?

MODELING THE EFFECT OF ONE CAR VARIATION ON OCCUPANCY MEASUREMENTS

POLYTECHNIQUE MONTRÉAL

$G(o) = \max\{ |o-o'| / \min(o,o') | \}$ for measured occupancy o

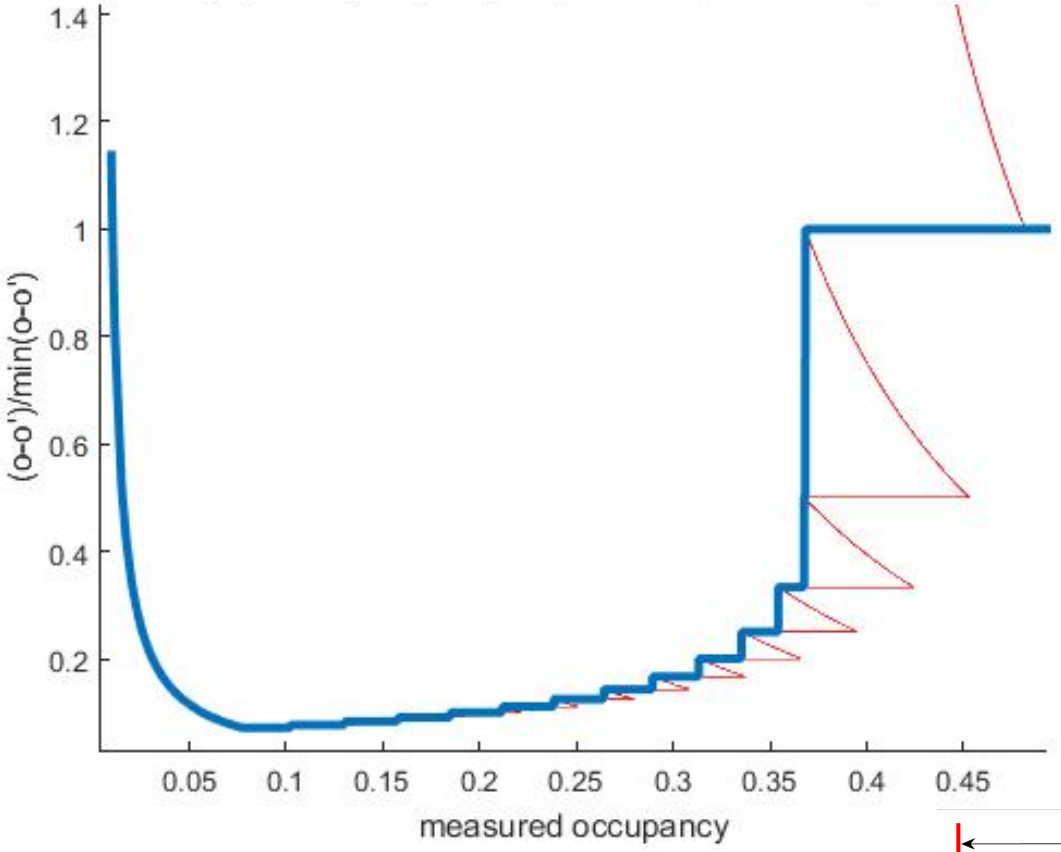

- How to set ɣ ?

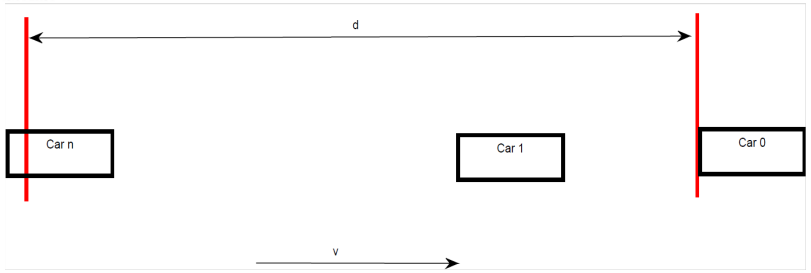

RELATIVE VARIATION OF OCCUPANCY MOBILE CENTURY DATASET

POLYTECHNIQUE MONTRÉAL

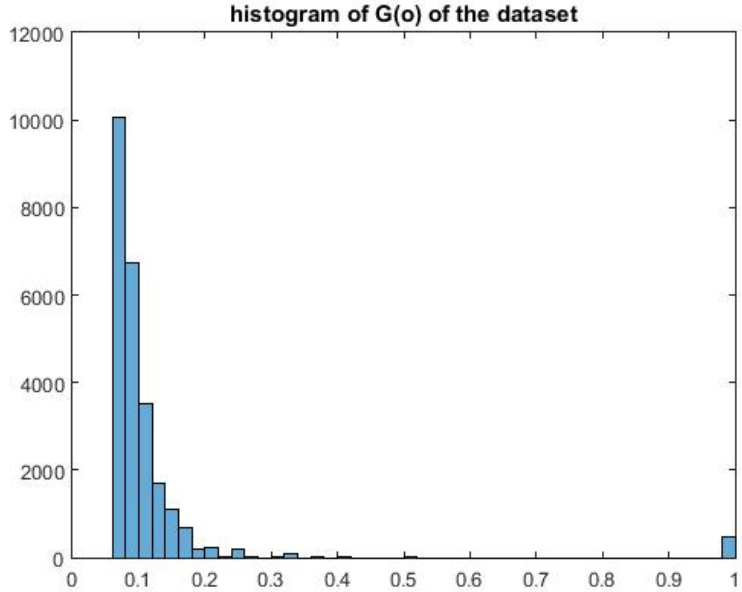

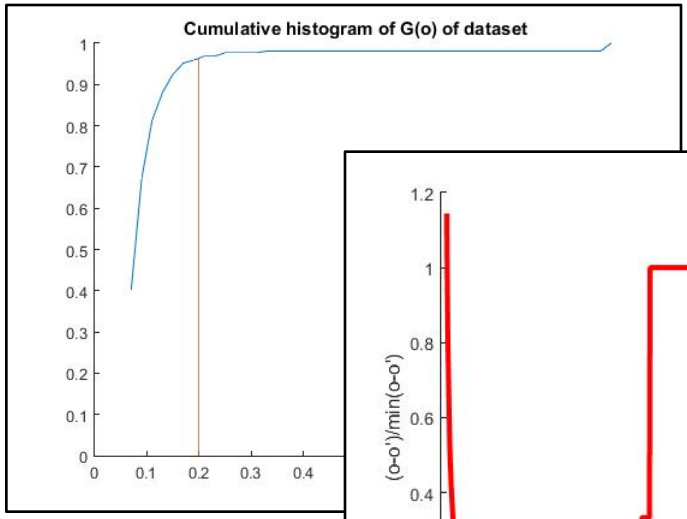

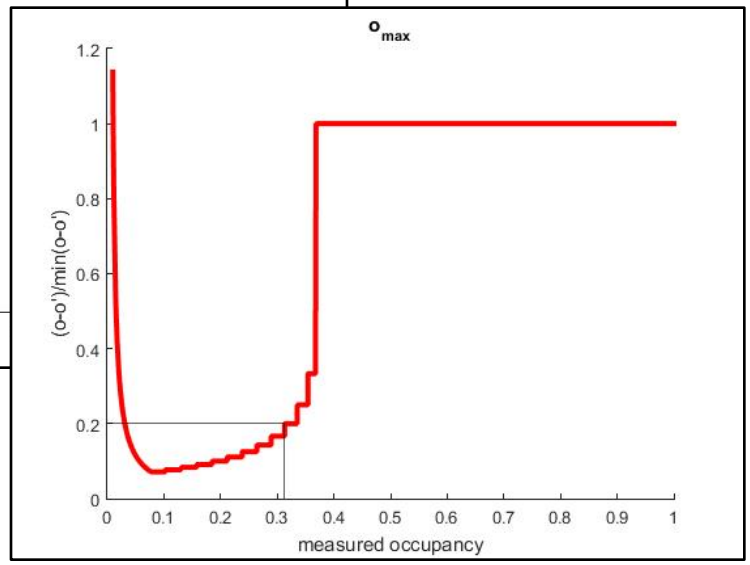

- $\frac{|o_{k,t}(d)-o_{k,t}(d')|}{\min\left(o_{k,t}(d),o_{k,t}(d')\right)} = 20\%$ satisfying 96% of data under this model

  $\Rightarrow o_{max}$=0.3132
- Limitation of occupancies for 100% of data protection under this model: $o \leftarrow \min(o, o_{max})$

AN ALTERNATIVE DENSITY MEASUREMENT MODEL

POLYTECHNIQUE MONTRÉAL

- An alternative density measurement model for the initial adjacency relation (change one trajectory arbitrarily)
  - Use of y_k^i directly is problematic for differential privacy

$$z_k^i = \rho_k^i + \eta_k^i \qquad \text{(density pseudo-measurement model)}$$

$$\text{avec} \quad z_k^i = \begin{cases} \frac{\phi_k^i}{v_f}, & \text{if } m_k^i = \text{Free} \\ \rho_J - \frac{\phi_k^i}{w}, & \text{if } m_k^i = \text{Congested} \end{cases}$$

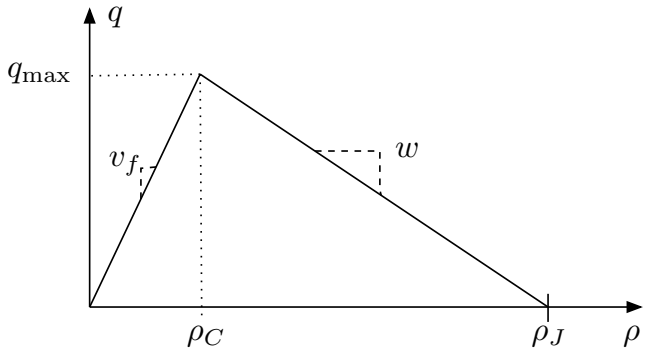

- Requires a traffic mode estimate m_k^i

$$m_k^i = \text{Free if and only if } y_k^i < \rho_c$$

  - Tolerates less precise measurements y_k^i
  - Mode estimate can be improved with HMM filter

A DIFFERENTIALLY PRIVATE MECHANISM FOR GENERAL DATA TYPES

POLYTECHNIQUE MONTRÉAL

- Gaussian mechanism not adapted to non numerical outputs such as mode estimate m_k^i (Congested or Free)

- Exponential mechanism:
  - Generic output space R for the possible answers to a query
  - Define a score function s(d,r) for the response r when the dataset is d
  - Pick the answer r with probability distribution

$$\frac{\exp(\epsilon s(d,r))dr}{\int_{\mathsf{R}} \exp(\epsilon s(d,r))dr}.$$

  - This mechanism is $2\epsilon\Delta s$ - differentially private, with

$$\Delta s = \sup_{r \in \mathsf{R}} \sup_{\mathrm{Adj}(d,d')} |s(d,r) - s(d',r)|$$

APPLICATION TO TRAFFIC MODE ESTIMATION

POLYTECHNIQUE MONTRÉAL

- We define the score function (high for high density)

$$s_k^i(\mathsf{C}) = y_k^i/\rho_c \ \ \text{(high for high density)}$$

$$s_k^i(\mathsf{F}) = 2 - y_k^i/\rho_c \ \ \text{(high for low density)}$$

$$\left(y_k^i = \frac{1}{g\lambda^i} \sum_{j=1}^{\lambda^i} o_{j,k}^i \right)$$

- Theorem: A mechanism publishing the traffic mode m_k^i at each period, by generating these outputs randomly and independently with distribution

$$\mathbb{P}(m_k^i = \mathsf{C}) = \frac{\exp(\epsilon s_k^i(\mathsf{C}))}{\exp(\epsilon s_k^i(\mathsf{C})) + \exp(\epsilon s_k^i(\mathsf{F}))}$$

is $4\epsilon \left(\sum_{i=1}^{M} \frac{1}{\lambda^i} \right)$ - differentially private.

- Pf idea: a single user can change the score function at most once at each sensor location, and by at most $2/\lambda_i$ (this requires bounding the occupation time of any car by $g\rho_c$)

OUTLINE

POLYTECHNIQUE MONTRÉAL

Differentially Private Traffic Estimation with Hydrodynamic Models

A. Dynamic Model for Density Estimation

B. Measurement models and sanitization of observations (input perturbation mechanisms)

C. EKF-based designs

EXAMPLE OF DP OUTPUT
DIRECT OCCUPANCY MEASUREMENTS

POLYTECHNIQUE MONTRÉAL

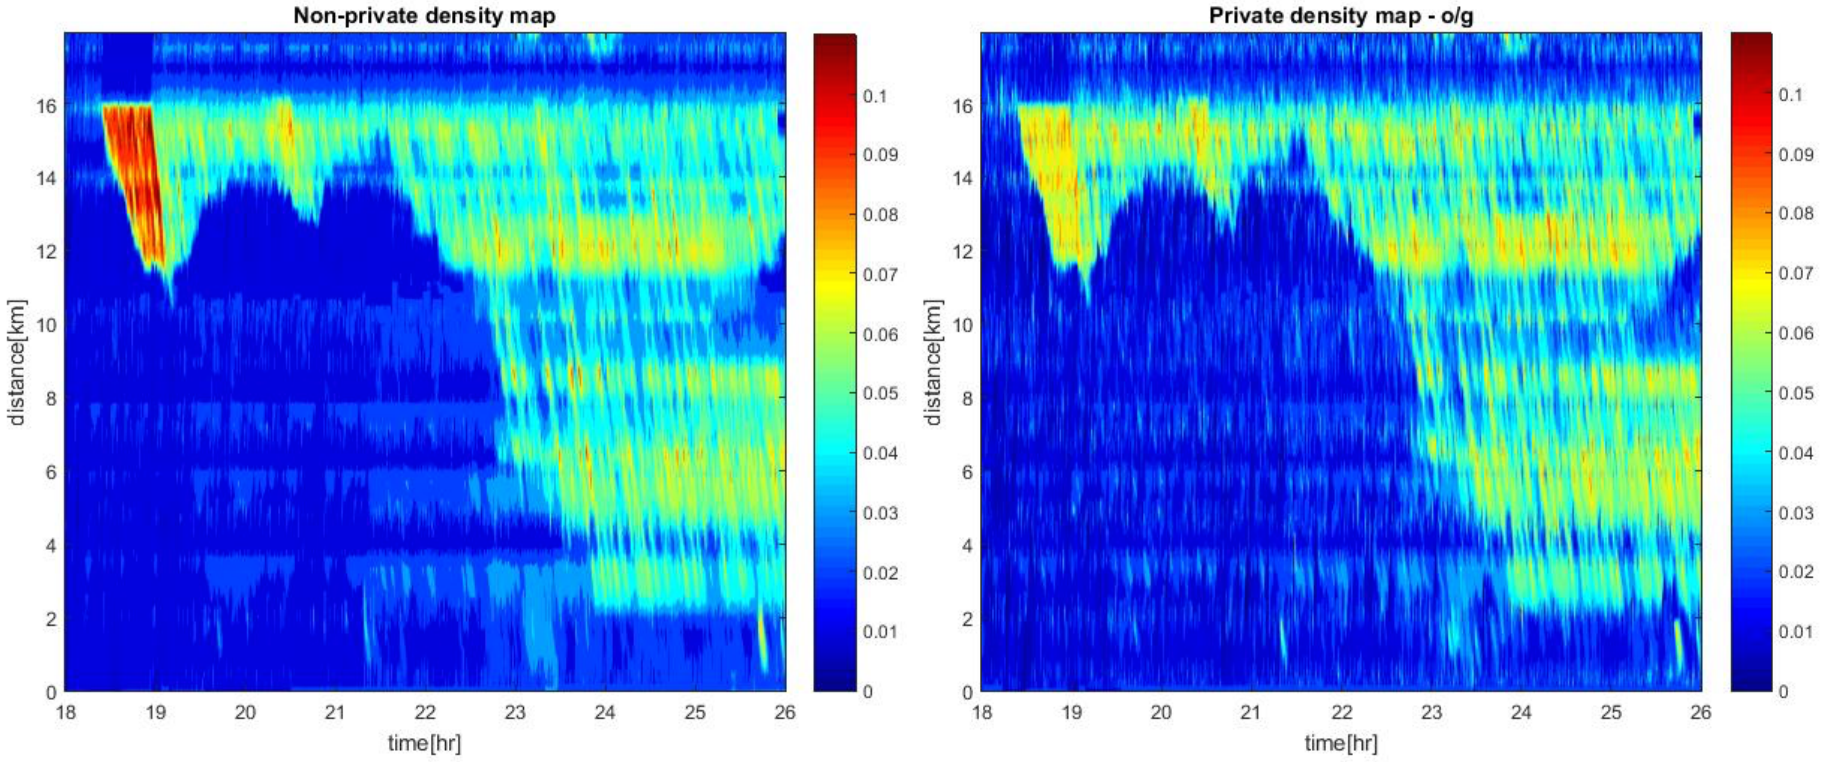

$$\epsilon = \ln 2\,, \delta = 0.05$$

FINAL ARCHITECTURE: FLOW INVERSION MODEL

POLYTECHNIQUE MONTRÉAL

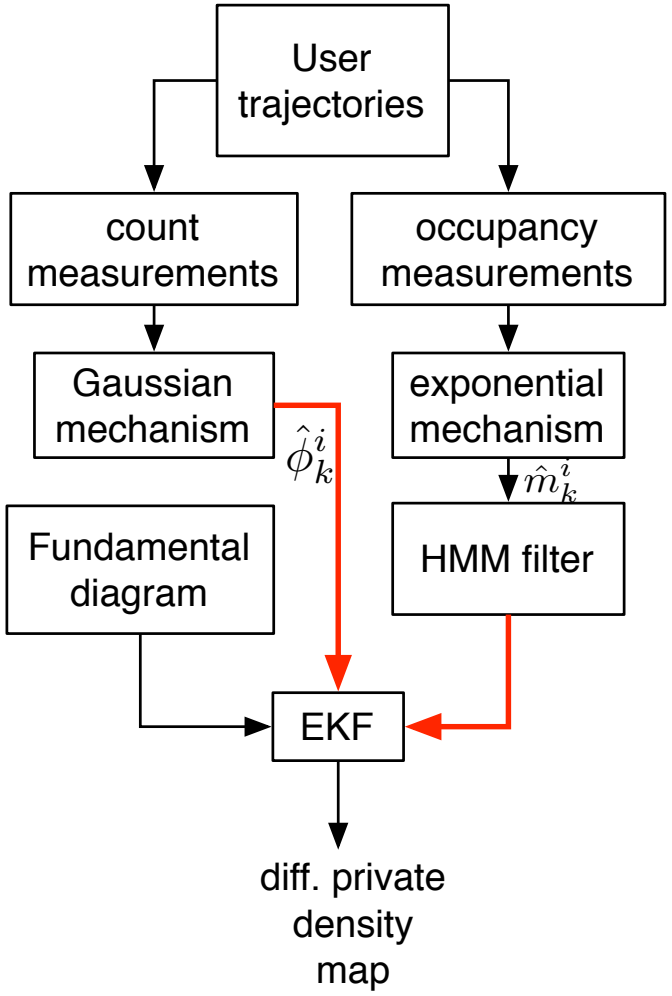

EXAMPLE OF DP OUTPUT FLOW INVERSION MODEL

POLYTECHNIQUE MONTRÉAL

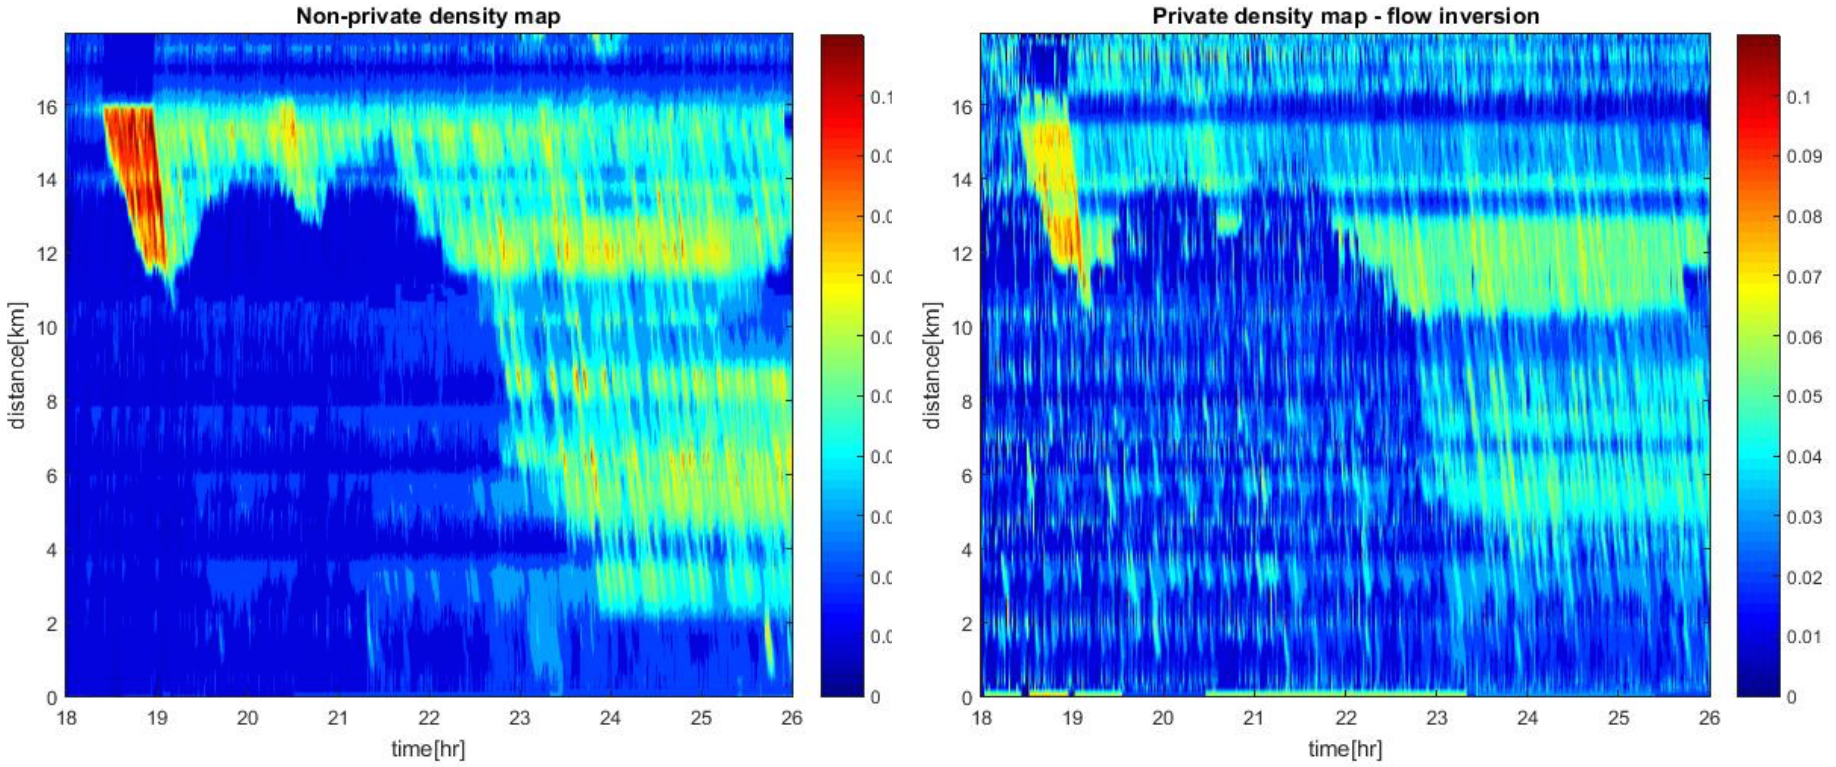

$$\epsilon = \ln 2\,, \delta = 0.05$$

SUMMARY

POLYTECHNIQUE MONTRÉAL

- Many cyber-physical applications raise privacy concerns that need to be addressed to encourage user participation
- Traffic information systems are an example of such sensitive systems
- Characterizing privacy-utility tradeoffs requires a quantitative definition of privacy
- Need privacy-preserving mechanisms for model-based data assimilation with various types of dynamic systems and data
- Further work needed to provide quantitative privacy guarantees for traffic information with reasonable performance impact