Data-Driven Dynamic Robust Resource Allocation for Efficient Transportation

Fei Miao
Postdoc, Electrical and Systems Engineering
University of Pennsylvania

Penn Engineering | GRASP Laboratory
General Robotics, Automation, Sensing & Perception Lab
PRECISE
PENN RESEARCH IN EMBEDDED COMPUTING AND INTEGRATED SYSTEMS ENGINEERING

Cyber-Physical Systems

Tight integration of communication and computation systems with control of physical world

Transportation &
Autonomous vehicles

Science of CPS

Industrial Automation and
Advanced Manufacturing

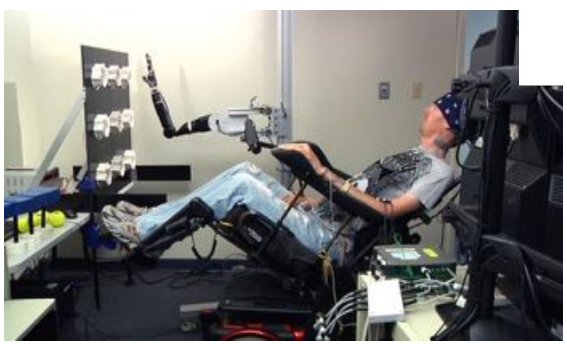

Medical Devices and Syster

Smart Grids

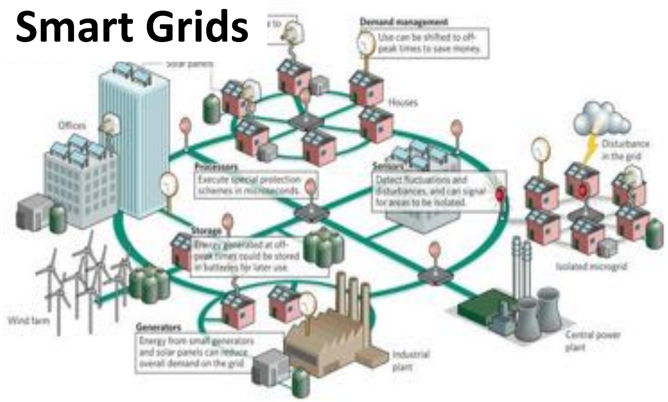

Building Automation

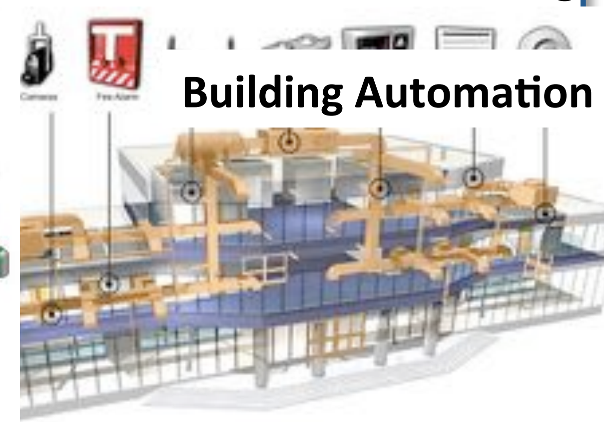

Internet of Things

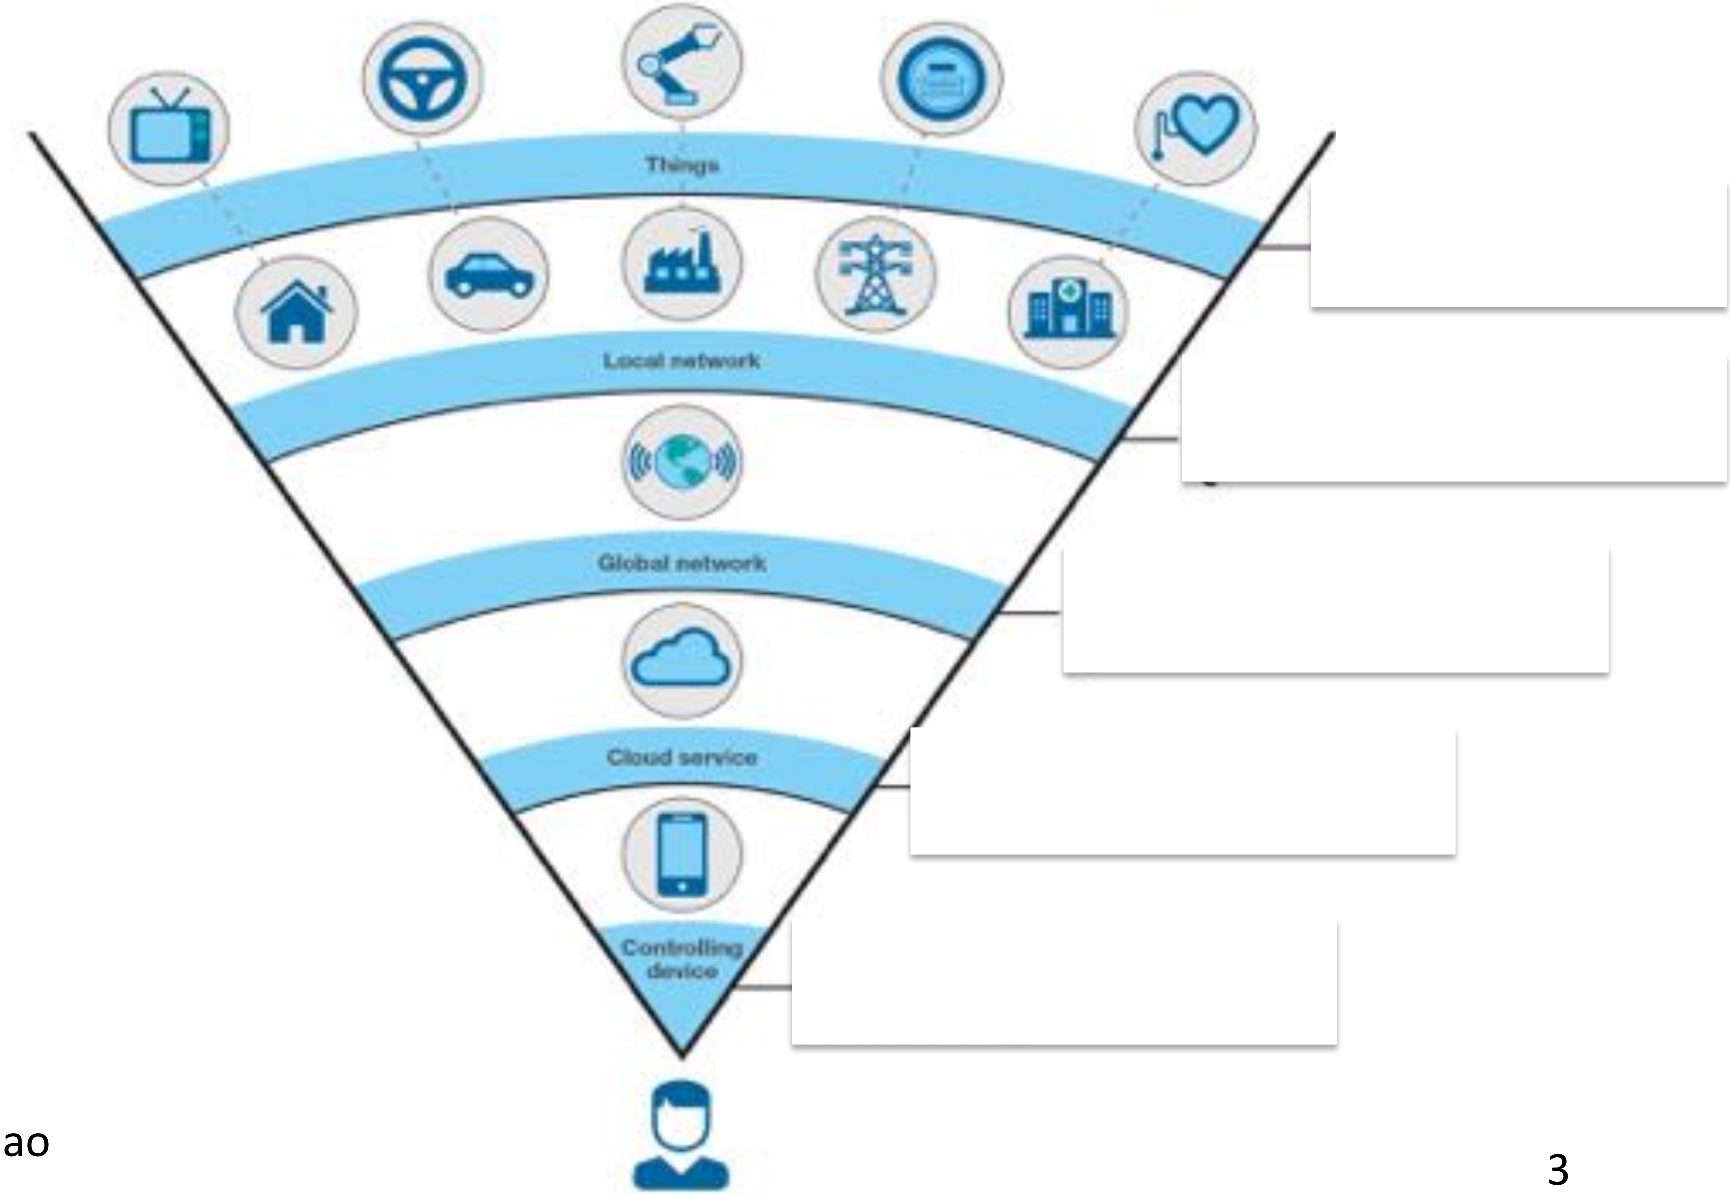

Penn
UNIVERSITY of PENNSYLVANIA

2013

2010

REP

DESIG
FEDER
A
NETWO

Ex
Pre

The White House unveils new US$160 million Smart Cities Initiative

17th September 2015 Tom Teodorczuk

- **$35 million** for **Smart Cities Grants** by **NSF**
  - **$10 million** for **CPS Program** in 2016
- **$70 million** for **Transportation, Energy**, and more, by
  - DoT
  - DoE
  - DHS
  - NIST

Smart Cities

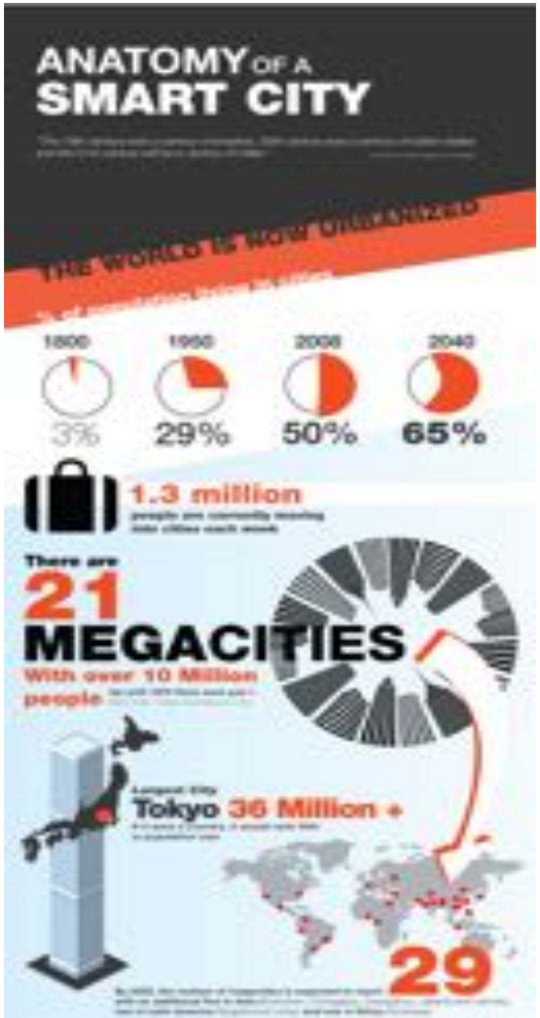

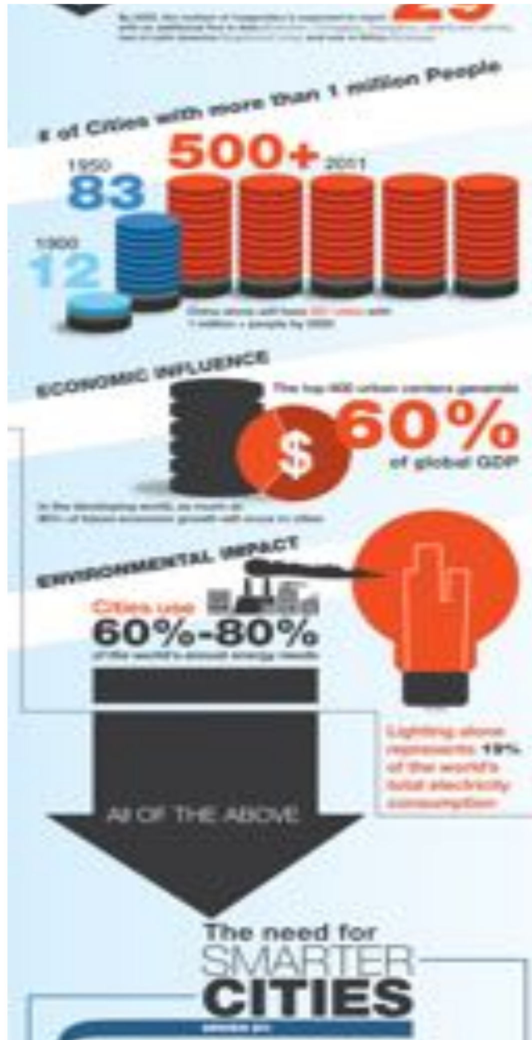

Challenges of CPS

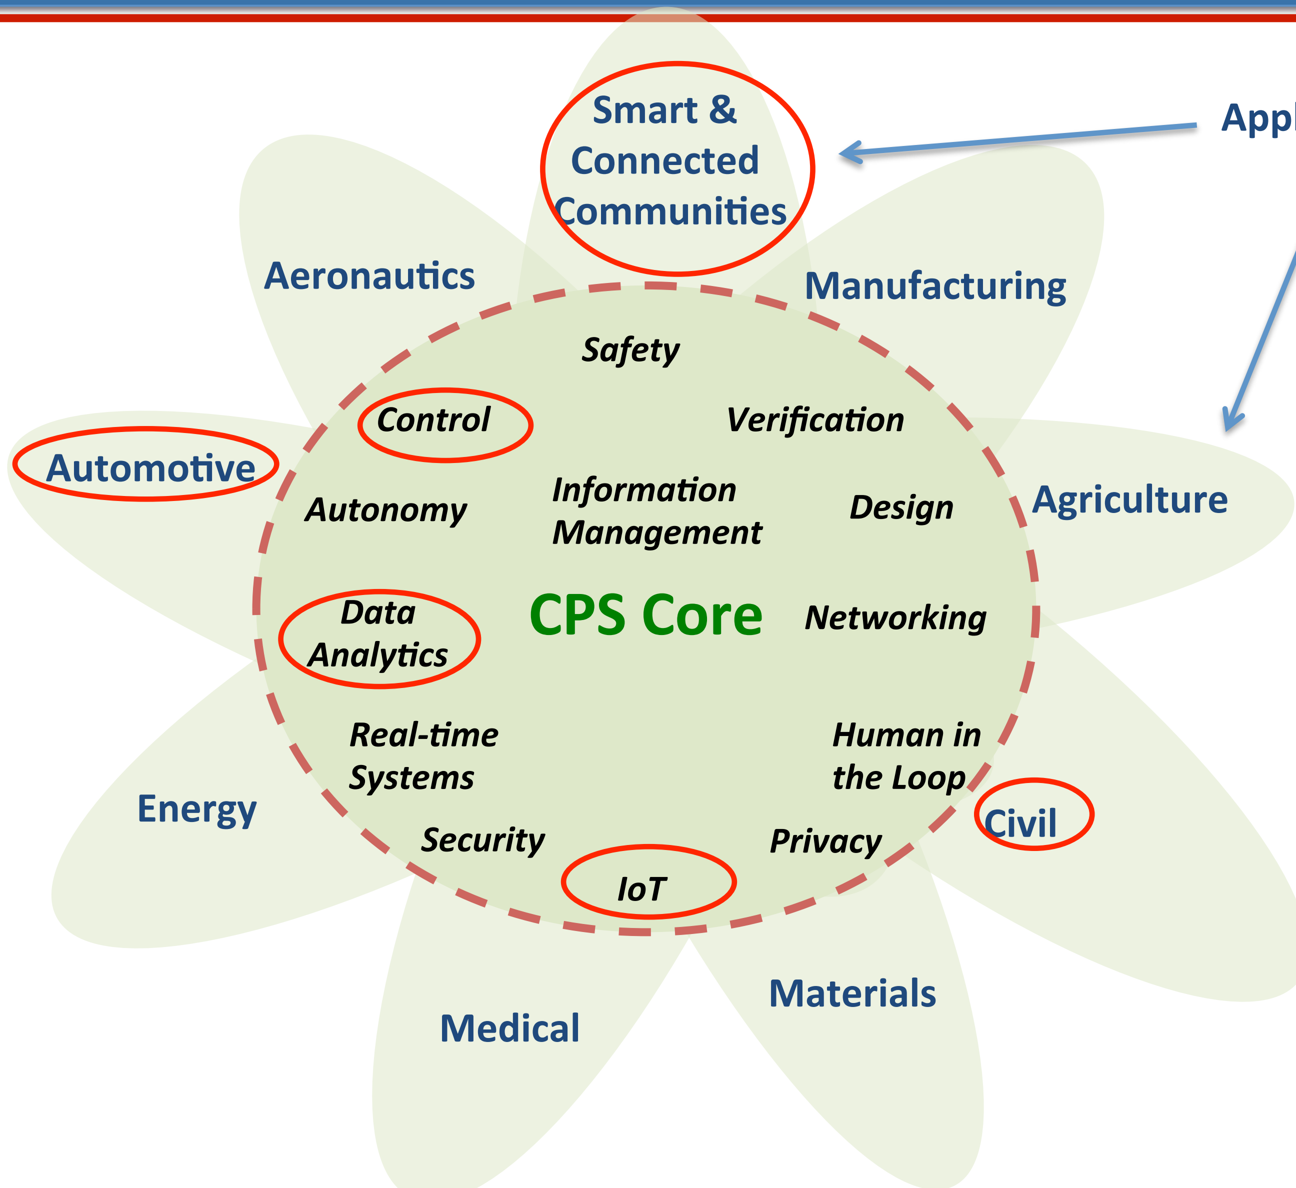

Data-Driven Control, Optimization of CPS

Data to information
Machine learning
Statistical method

Minimize cost or
Maximize profit
Optimization

Information to Action
Control

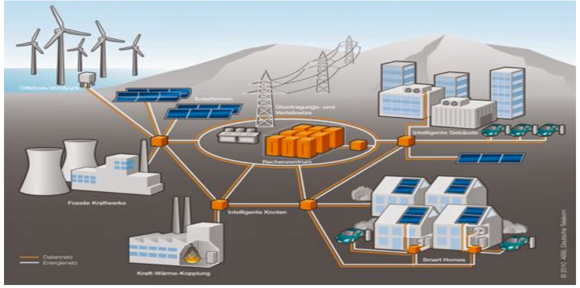

Demand response of smart grid

Intelligent transportation systems: integration research challenges unsolved

Growing **complexity and dynamics**: scale, uncertainty, heterogeneous…

Domain: Transportation Systems

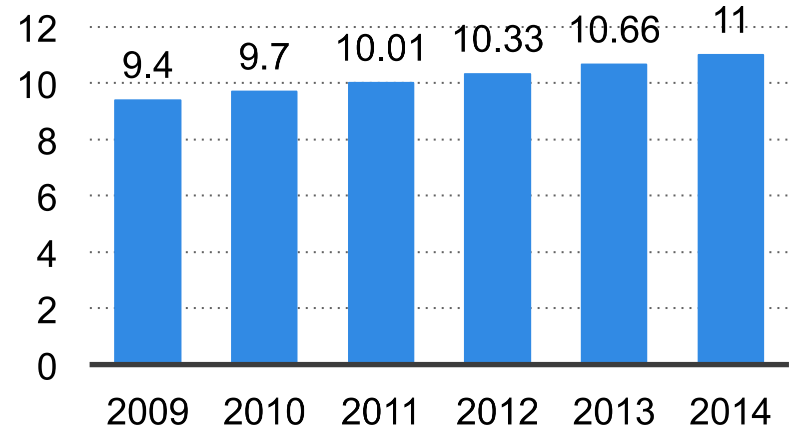

Revenue of US taxi services increases

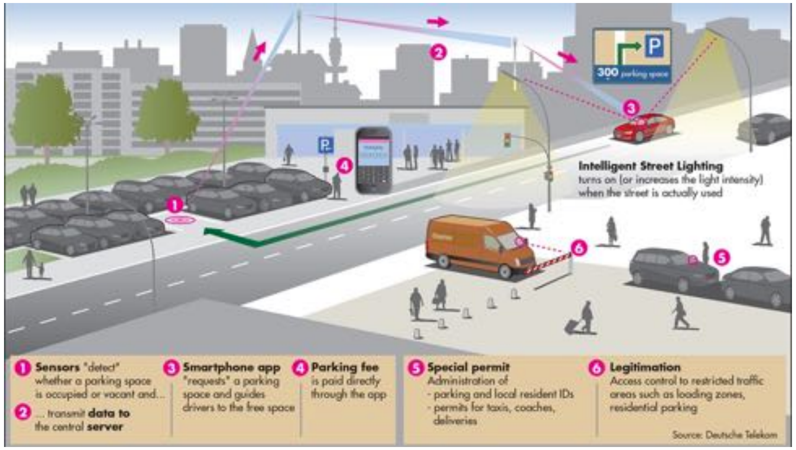

Smart Parking

Self-driving car

State of the Art – *The Dark Side…*

 =20 × 

Cruising Mileage

Idle: 300 million miles/year

Previous Work

- Urban traffic, demand modeling
  - Spatial and temporal patterns in speed prediction: [P. Jaillet et al., 2014]
  - Infer traffic condition based on taxi trip dataset: [Work et al., 2015]

- Coordination and resource allocation in smart transportation
  - Smart parking: [Geng and Cassandras, 2013]
  - Bike redistribution and incentives in sharing:[Morari et.al, 2014]
  - Mobility-on-demand systems: autonomous vehicles/robotics

    Minimize re-balancing number; with future demand : [Pavone et.al, 2014, 2015]

    Dynamic on-demand ride-sharing [Frazzoli, et al., 2016]

    Evaluation metric: number of vehicles needed; simulated waiting time

Novelty of my work

Brings data-driven optimization to system-wide efficient transportation

Motivation: Efficient Transportation

No system-level optimal, proactive resource allocation

Data→Predicted demand →Improved system performance?

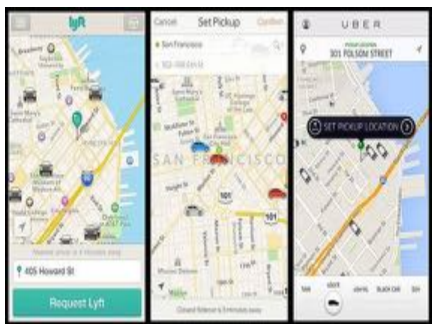

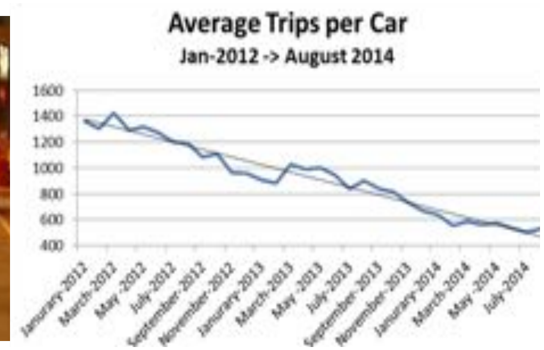

On demand ridesharing service

- Greedy, increase human satisfaction myopically
- Optimal vehicle allocating with accurate, known demand distributions

Problem and Goal

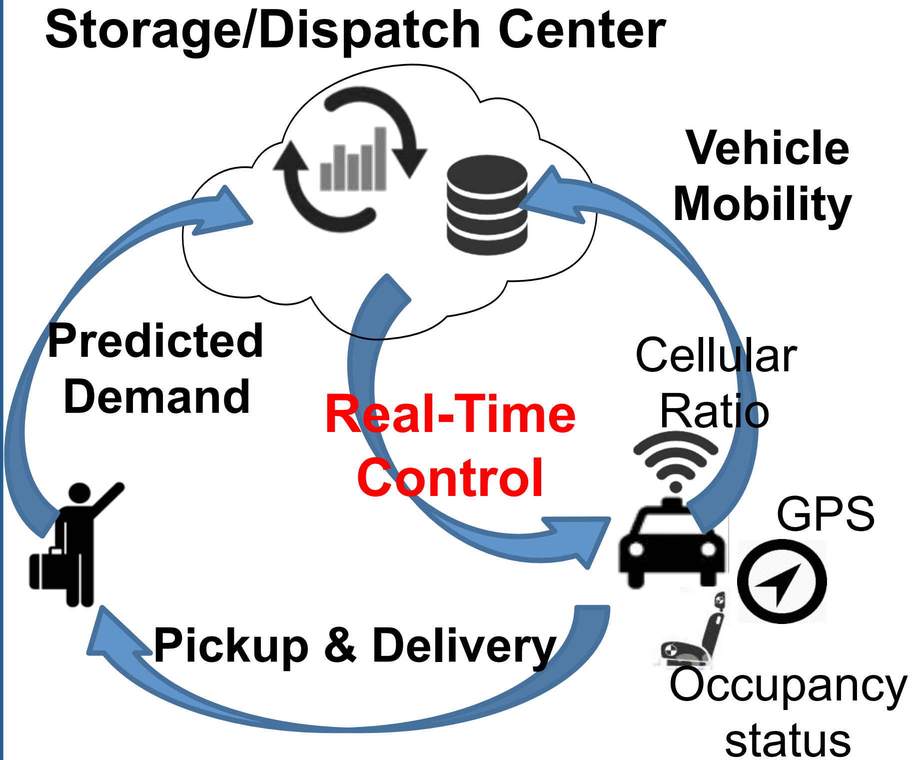

Challenges: scale, uncertainty
Conflict of interests
Limited resource
Customers: reach destination soon
System optimal: minimum idle

System-level balanced supply-demand ratio (fair service) with least total idle distance

Hierarchical

Local controller: heuristic, greedy, matching, etc.

Contributions Outline: Data-Driven Control, Optimization of CPS

Large-scale dynamic decision, receding horizon control

Demand predicted from historical and real-time sensing data

Uncertain demand: computationally tractable, probabilistic cost guarantee

1.Robust optimization for the worst-case resource allocation cost

2.Distributionally robust optimization for the expected cost

Contributions Outline: Data-Driven Control, Optimization

Large-scale dynamic decision, receding horizon control

Demand predicted from historical and real-time sensing data

- Reduce idle mileage of ride-sharing service
- Considering both current and predicted future demand

Receding Horizon Control Hierarchical Vehicle Dispatch

Spatial-temporal data → demand, mobility
e.g. Poisson, vector time series

Dynamically optimize: time t, consider dispatch costs of (t,…, t+τ), execute decision for t

1.Update sensing data, demand prediction

2. System-level: balanced supply towards demand with minimum total idle distance

3.Send decisions to vehicles, local dispatcher

t = t+1

Spatial
Urban Regions

Temporal
Hourly Windows

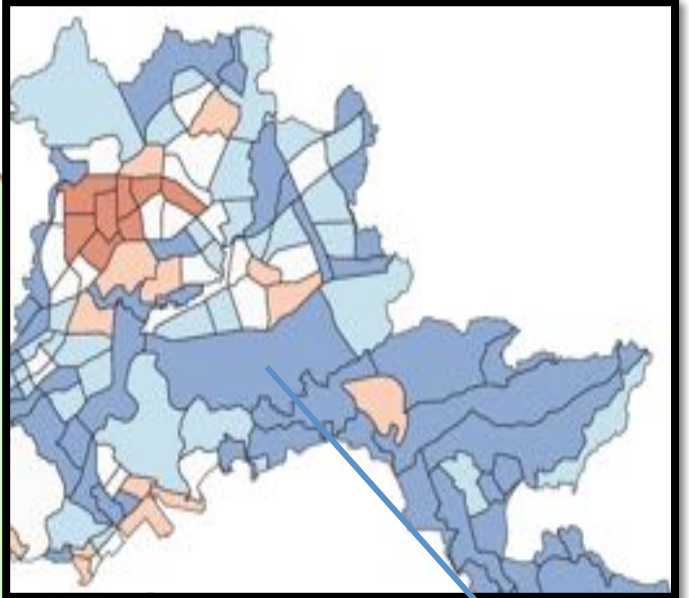

n regions, optimization

One region, local dispatcher

Objective #1: Fair Service → Balanced Supply

- Goal: Fair service, similar average waiting time

Predicted demand — r_i^k

Supply — original number of vacant vehicles before dispatch (GPS): L_i^k

Decision variables $X_{ij}^k \geqslant 0$: number of vacant vehicles from region i to j

- Local supply/demand ratio close to global level: penalty difference

$$J_E = \sum_{k=1}^{\tau} \sum_{i=1}^{n} \left| \underbrace{\frac{\sum_{j=1}^{n} X_{ji}^k - \sum_{j=1}^{n} X_{ij}^k + L_i^k}{r_i^k}}_{\textbf{Local}\text{ s/d ratio}} - \underbrace{\frac{N^k}{\sum_{j=1}^{n} r_j^k}}_{\textbf{Global}\text{ s/d ratio (city)}} \right|$$

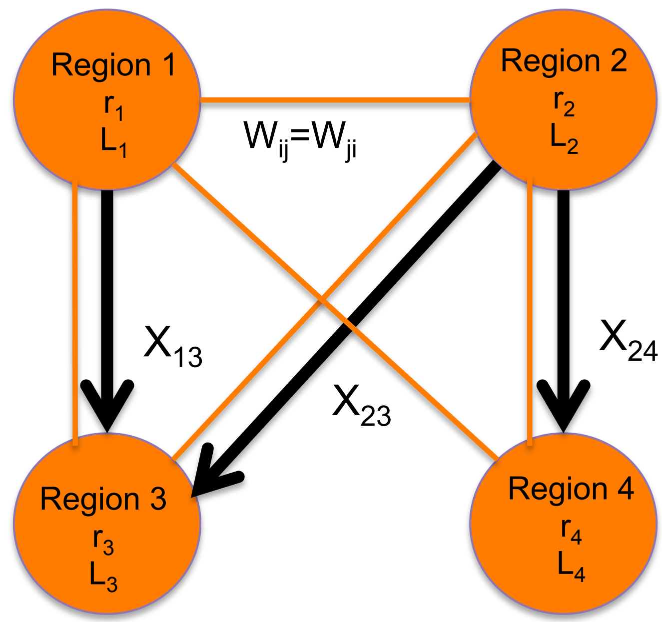

Objective #2: Minimum Cost: Total Idle Distance

- The distance between region i and region j at time k: W_{ij}^k

Decision variables $X_{ij}^k \geqslant 0$: number of vacant vehicles from region i to j

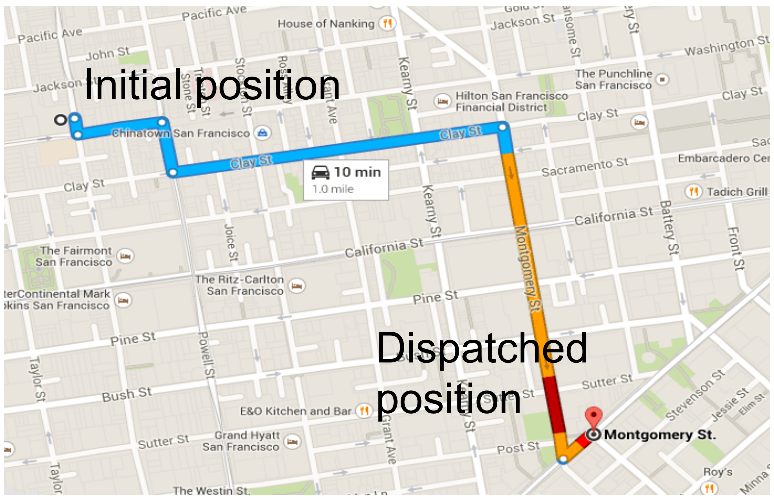

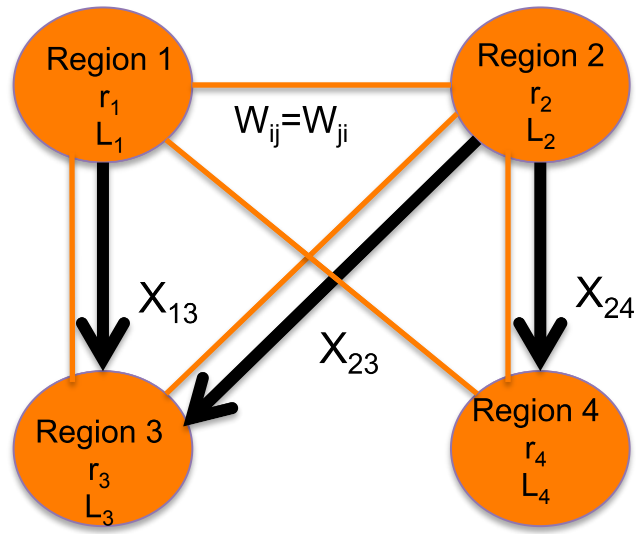

- Total idle distance to meet demand

$$J_D = \sum_{k=1}^{\tau} \sum_{i=1}^{n} \sum_{j=1}^{n} X_{ij}^k W_{ij}^k$$

System-Level Optimal Vehicle Dispatch

Idle distance to meet demand

Balanced supply

$$\min_{X^{1:\tau},L^{2:\tau}} \quad J=\sum_{k=1}^{\tau}\Big(\boxed{J_D(X^k)}+\beta \boxed{J_E(X^k,r^k)}\Big)$$

$$=\sum_{k=1}^{\tau}\sum_{i=1}^{n}\left(\sum_{j=1}^{n}X_{ij}^kW_{ij}+\beta\left|\frac{\sum\limits_{j=1}^{n}X_{ji}^k-\sum\limits_{j=1}^{n}X_{ij}^k+L_i^k}{r_i^k}-\frac{N^k}{\sum\limits_{j=1}^{n}r_j^k}\right|\right)$$

$$\begin{aligned}
\text{s.t.} \quad & (L^{k+1})^T=(\mathbf{1}_n^TX^k-(X^k\mathbf{1}_n)^T+(L^k)^T)P^k, && \text{State dynamics: trip}\\
& \mathbf{1}_n^TX^k-(X^k\mathbf{1}_n)^T+(L^k)^T>0, && \text{Supply positive}\\
& X_{ij}^kW_{ij}^k\leqslant m^kX_{ij}^k, && \text{Sparsity constraint (distance limited)}\\
& X_{ij}^k\geqslant 0
\end{aligned} \qquad (1)$$

Computational complexity: polynomial of variable number (τn^2)
Spatial-temporal granularity

Experiment: RHC dispatch with Real-Time information

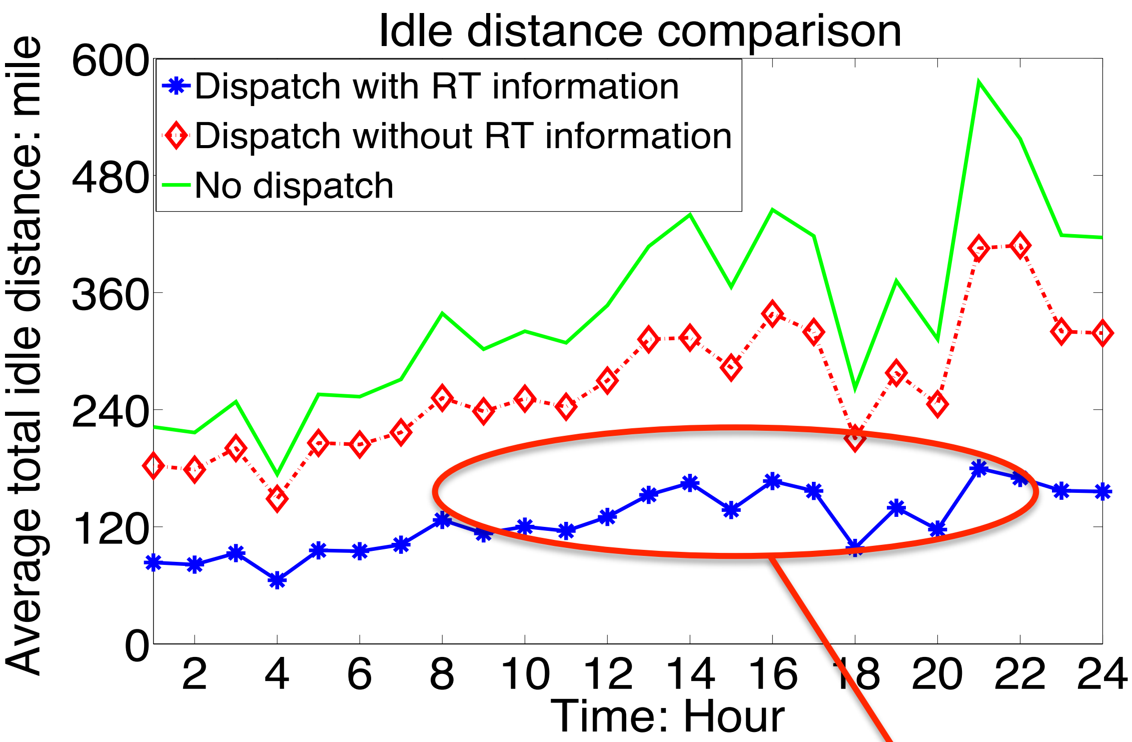

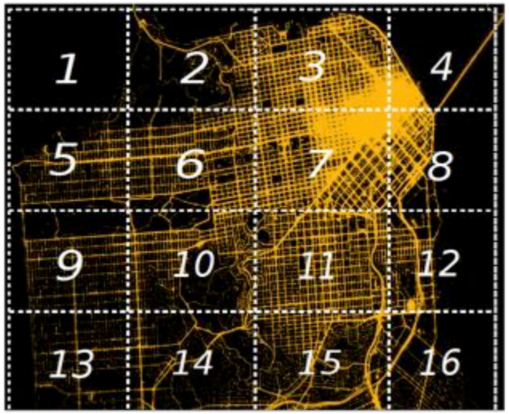

Collection Period	28 days
Number of Taxis	500
GPS Record Number	1,000,000

- System-level vehicle balancing and local shortest path dispatcher
- Average demand, GPS vehicle locations

Average idle distance ↓ 42%

Experiment: Dispatch with Real-Time Information

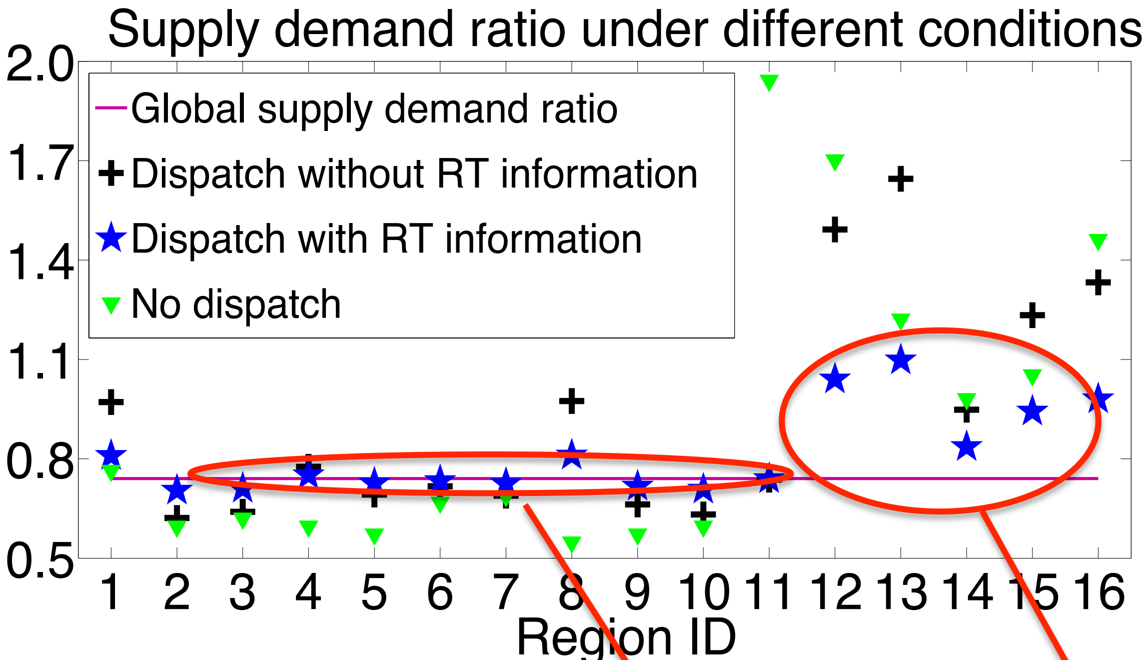

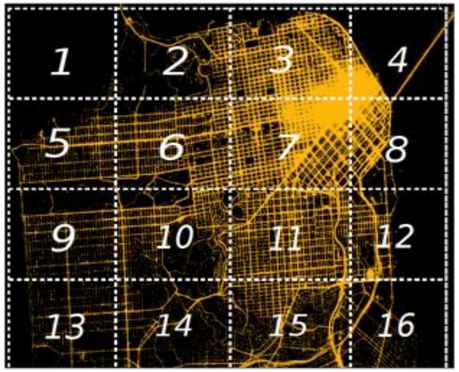

Collection Period	28 days
Number of Taxis	500
Record Number	1,000,000

- System-level vehicle balancing and local shortest path dispatcher
- Average demand, GPS vehicle locations

Average supply-demand ratio error ↓ 45%

Large error caused by prediction error

Contributions Outline: Data-Driven Control, Optimization of CPS

Uncertain demand: computationally tractable, probabilistic cost guarantee

1. Robust optimization for the worst-case resource allocation cost
2. Distributionally robust optimization for the expected cost

Data to Uncertain Demand or Demand Distributions

Spatial-temporal data
- Partition city map, cluster dataset
- Trip/trajectory→ Aggregated demand

$$\hat{r}_{k+1} = f_r(I_{[k-l,k]})$$

$$r_{k+1} = \hat{r}_{k+1} + \delta_{k+1}.$$

Vector time series
Poisson distribution

- Bootstrap (repeated experiments)
- Confidence region of H_0: mean, covariance, probability distribution
- Closed and convex uncertainty sets

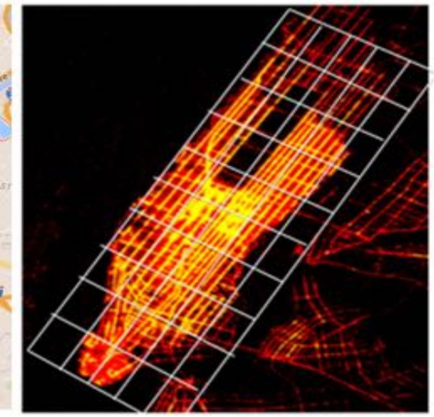

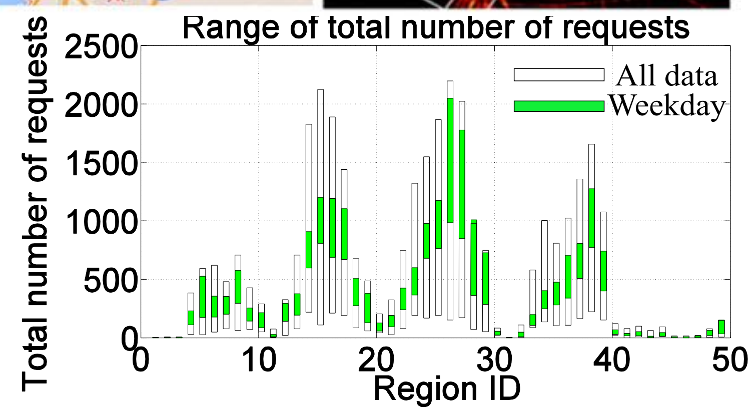

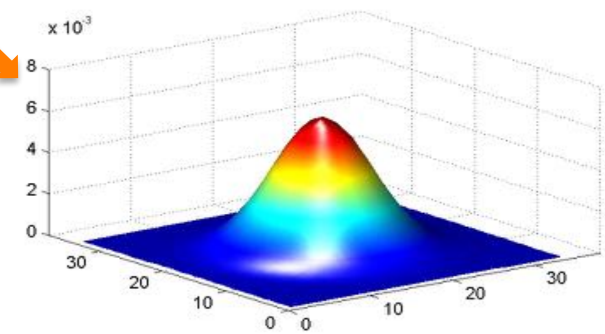

Uncertain Demand: Computationally Tractable Approximation

- Predicted demand —a closed and convex set

$$r_c = (r^1, r^2, \ldots, r^\tau) \in \Delta \in \mathbb{R}^{\tau n}$$ (demand at time k: r^k)

- Decision variables $X_{ij}^k \geqslant 0$: number of vacant taxis from region i to j

$S_i^k = f(X^{1:k})$ supply at region i time k, linear of previous decisions

- Objective not computationally intractable under uncertain demand

$$J_E = \sum_{k=1}^{\tau} \sum_{i=1}^{n} \left| \frac{S_i^k}{r_i^k} - \frac{\sum_{j=1}^{n} S_j^k}{\sum_{j=1}^{n} r_j^k} \right|$$

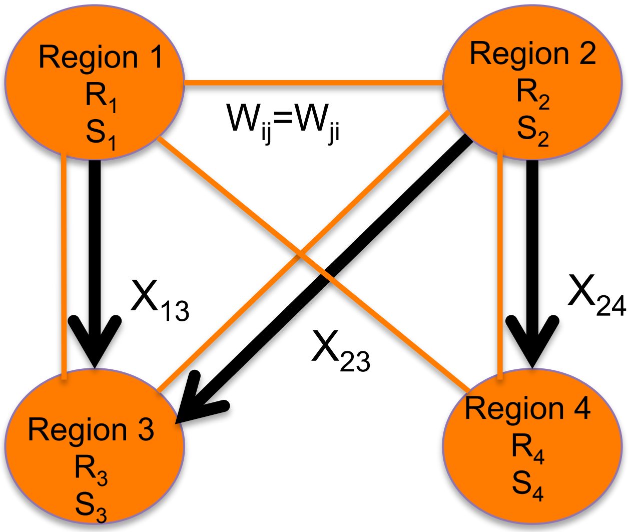

Uncertain Demand: Computationally Tractable Approximation

- Predicted demand —a closed and convex set

  $$r_c = (r^1, r^2, \dots, r^\tau) \in \Delta \in \mathbb{R}^{\tau n}$$ (demand at time k: r^k)

- Decision variables $X_{ij}^k \geqslant 0$: number of vacant vehicles from region i to j

  $S_i^k = f(X^{1:k})$ supply at region i time k, linear of previous decisions

*Theorem: computationally tractable approximation (concave of demand)

*Fei Miao et. al, CDC, 2015; Fei Miao et. al, under revision, IEEE TCST 2016

$$J_E = \sum_{k=1}^{\tau} \sum_{i=1}^{n} \left| \frac{S_i^k}{r_i^k} - \frac{\sum_{j=1}^{n} S_j^k}{\sum_{j=1}^{n} r_j^k} \right|$$

$$\boxed{\sum_{k=1}^{\tau} \sum_{i=1}^{n} \frac{r_i^k}{(S_i^k)^\alpha} \qquad \alpha > 0}$$

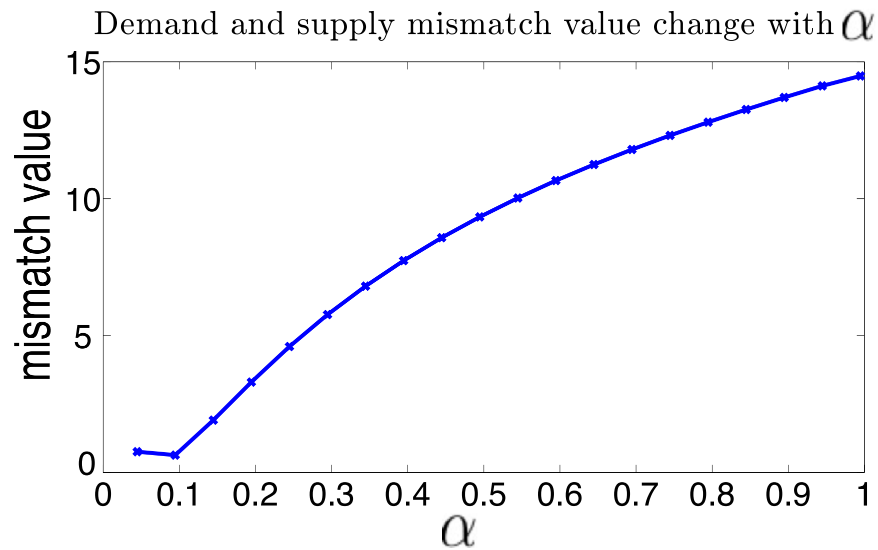

Probabilistic Cost Guarantee for Worst-Case by Robust Solutions

- Predicted demand —a closed and convex set

$$r_c = (r^1, r^2, \ldots, r^\tau) \in \Delta \in \mathbb{R}^{\tau n}$$ (demand at time k: r^k)

Probabilistic guarantee

$$\begin{aligned} \min_{X^{1:\tau}} \quad & J(X^{1:\tau}, r_c) \\ \text{s.t.} \quad & P_{r_c \sim \mathbb{P}^*(r_c)}(f(X^{1:\tau}, r_c) \leqslant 0) \geqslant 1 - \epsilon. \end{aligned}$$

→

Robust optimization problem

$$\begin{aligned} \min_{X^{1:\tau}} \boxed{\max_{r_c \sim \Delta}} \quad & J(X^{1:\tau}, r_c), \\ \text{s.t.} \quad & f(X^{1:\tau}, r_c) \leqslant 0, \end{aligned}$$

***Theorems: robust optimization → Equivalent convex optimization**

1. Approximated objective: concave of uncertainties, convex of variables
2. Closed and convex uncertainty set (related to ε, first/second order)

*Fei Miao et. al, CDC, 2015; Fei Miao et. al, under revision, IEEE TCST 2016

Minimize Expected Cost: Distributionally Robust Optimization (DRO)

- **Motivation:** Trade-off between the expected and the worst-case costs

- **Formulation**

Demand: $r_c = (r^1, r^2, \ldots, r^\tau) \in \mathbb{R}^{\tau n}, \quad r_c \sim F^*, \quad F^* \in \mathcal{F}$

Resource allocation decision (spatial-temporal): $X^{1:\tau} = \{X^1,\ X^2,\ \ldots\ X^\tau\}$

Stochastic programming

$$\begin{aligned} \min_{X^{1:\tau}} \quad & \mathbb{E}_{r_c \sim F^*}\left[J(X^{1:\tau}, r_c)\right] \\ \text{s.t.} \quad & X^{1:\tau} \in \mathcal{D}_c. \end{aligned}$$

$\rightarrow$

Distributionally robust optimization

$$\begin{aligned} \min_{X^{1:\tau}} \quad & \boxed{\max_{F \in \mathcal{F}}} \quad \mathbb{E}\left[J(X^{1:\tau}, r_c)\right] \\ \text{s.t.} \quad & X^{1:\tau} \in \mathcal{D}_c. \end{aligned}$$

***Theorem: DRO → Equivalent convex optimization**

1.Approximated objective: concave of uncertainties, convex of variables

2.Closed and convex set of probability distributions (first/second order)

*Fei Miao et. al, accepted, CPSWeek ICCPS 2017.

Applications of Data-Driven DRO Resource Allocation

Allocation cost | Metric of service quality (d/s ratio related)

$$\min_{X^{1:\tau}, S^{1:\tau}} . \max_{F \in \mathcal{F}} \mathbb{E}\left[\sum_{k=1}^{\tau}\left(\boxed{J_D(X^k)} + \beta \boxed{\sum_{i=1}^{n} \frac{r_i^k}{(S_i^k)^\alpha}}\right)\right] \quad (3)$$
$$\text{s.t.} \quad X^{1:\tau}, S^{1:\tau} \in \mathcal{D}_c$$

$S_i^k = f(X^{1:k})$ supply at region i time k, linear of previous decisions

Resource allocation under demand uncertainties

- Autonomous vehicle balancing for mobility-on-demand system
- Bicycle balancing: supply-demand ratio at each station is in a range
- Hierarchical carpooling framework: global balance, local carpool
- Real-time demand response with limited resource

Evaluations: Robust *VS* Non-Robust Solutions

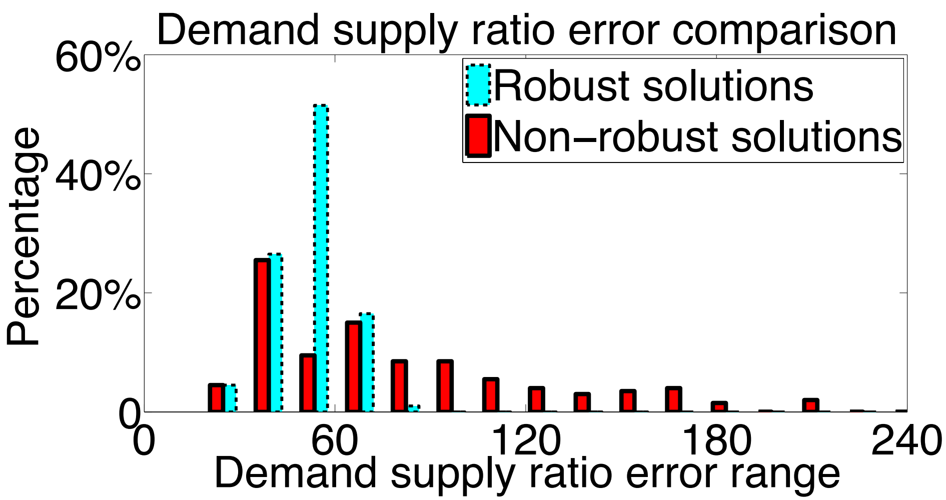

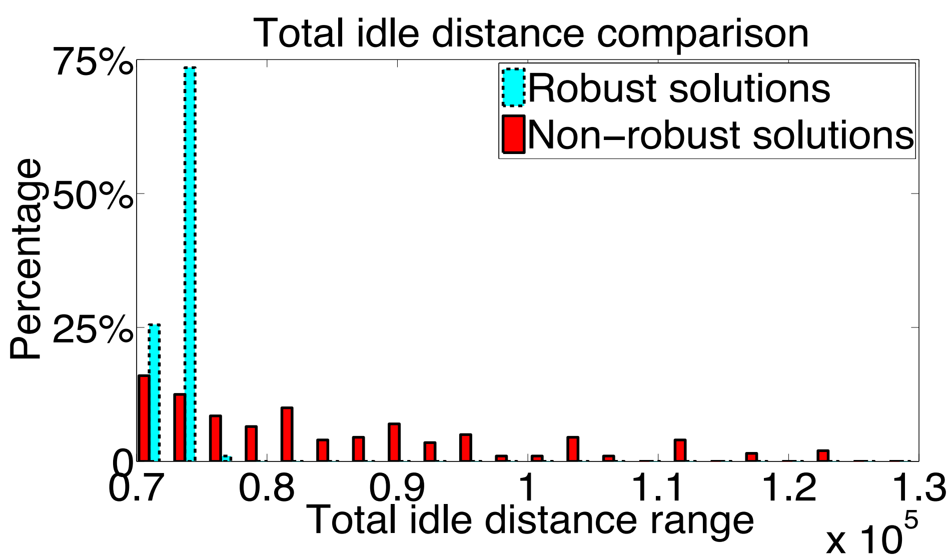

Collection Period	4 years
Data Size	100 GB
Trip number	700 million

Cross-validation

Second-order-cone uncertainties

Probabilistic guarantee 1-ε=0.75

The average demand supply ratio error reduced by 31.7%

The average total idle distance

Robust VS non-robust: ↓10.13%

Evaluation: Average Costs of (distributionally) Robust Solutions

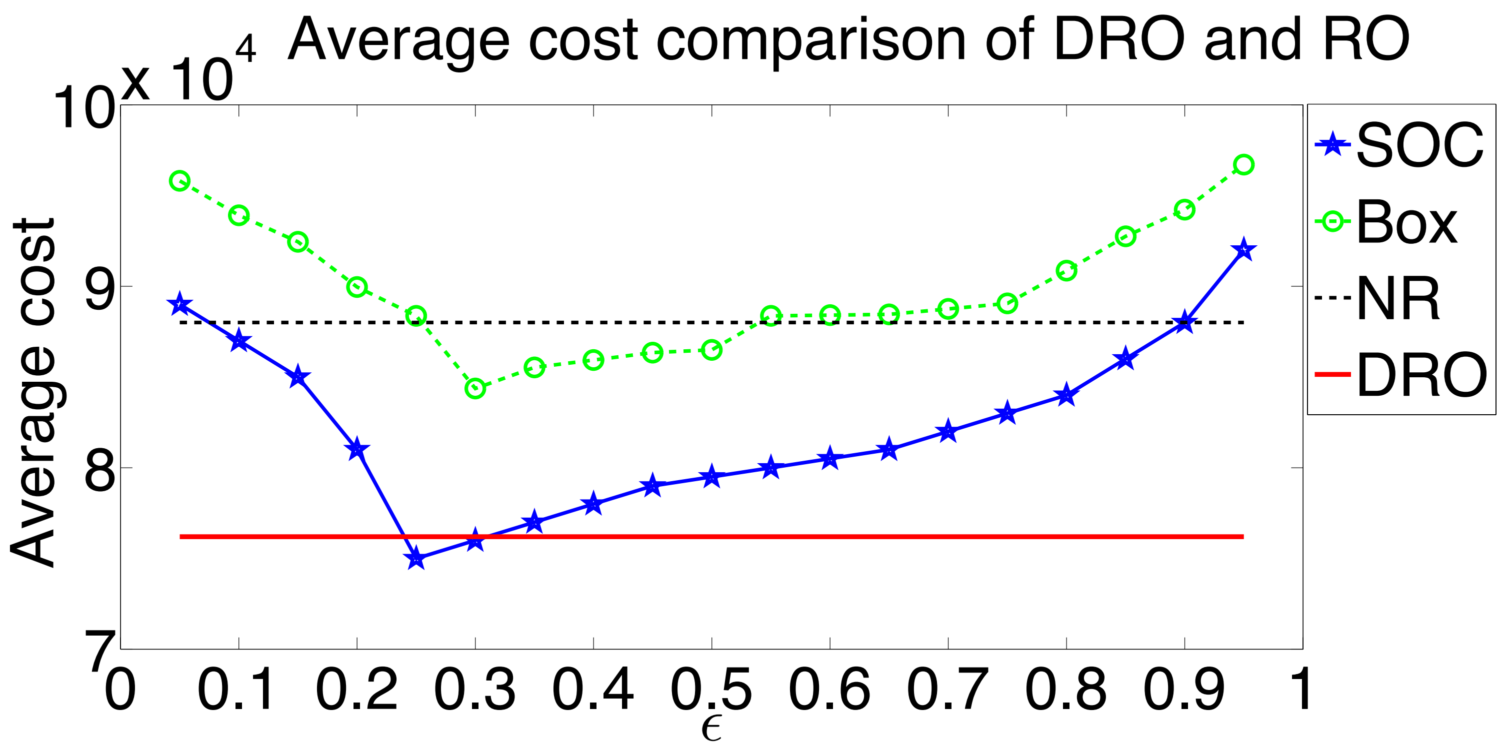

SOC: second-order-cone uncertainty set
Box: range of demand at each region
Compared with non-robust (NR) solutions (100GB NYC taxi data)
The average total idle driving distance is reduced by 10.05%

Dynamic Region Partition with DRO to Reduce Cost

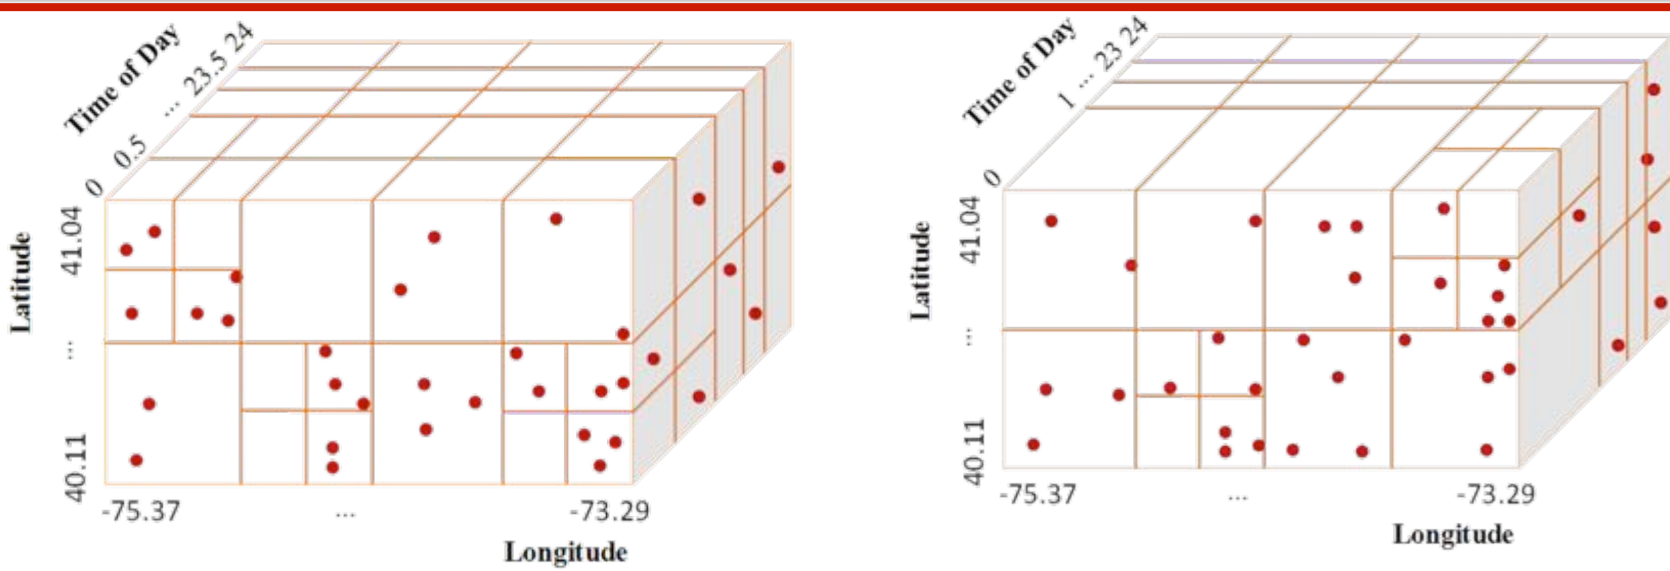

Region Division	Grid Idle Mile	Quad-Tree Idle Mile	Change Rate
t=1 hour	7.63×10^4	6.62×10^4	13.1%
T=30 minutes	6.84×10^4	5.47×10^4	20.0%

- [Quad-tree region partition, uncertainty set of demand probability distributions, distributionally robust vehicle balancing] ***VS*** [static grid, non-robust model]

→ total idle distance↓60 million miles, 8 million dollars gas consumption/year

- Compared with total idle distance of original data: ↓55%

Summary of Contributions in Data-Driven CPS

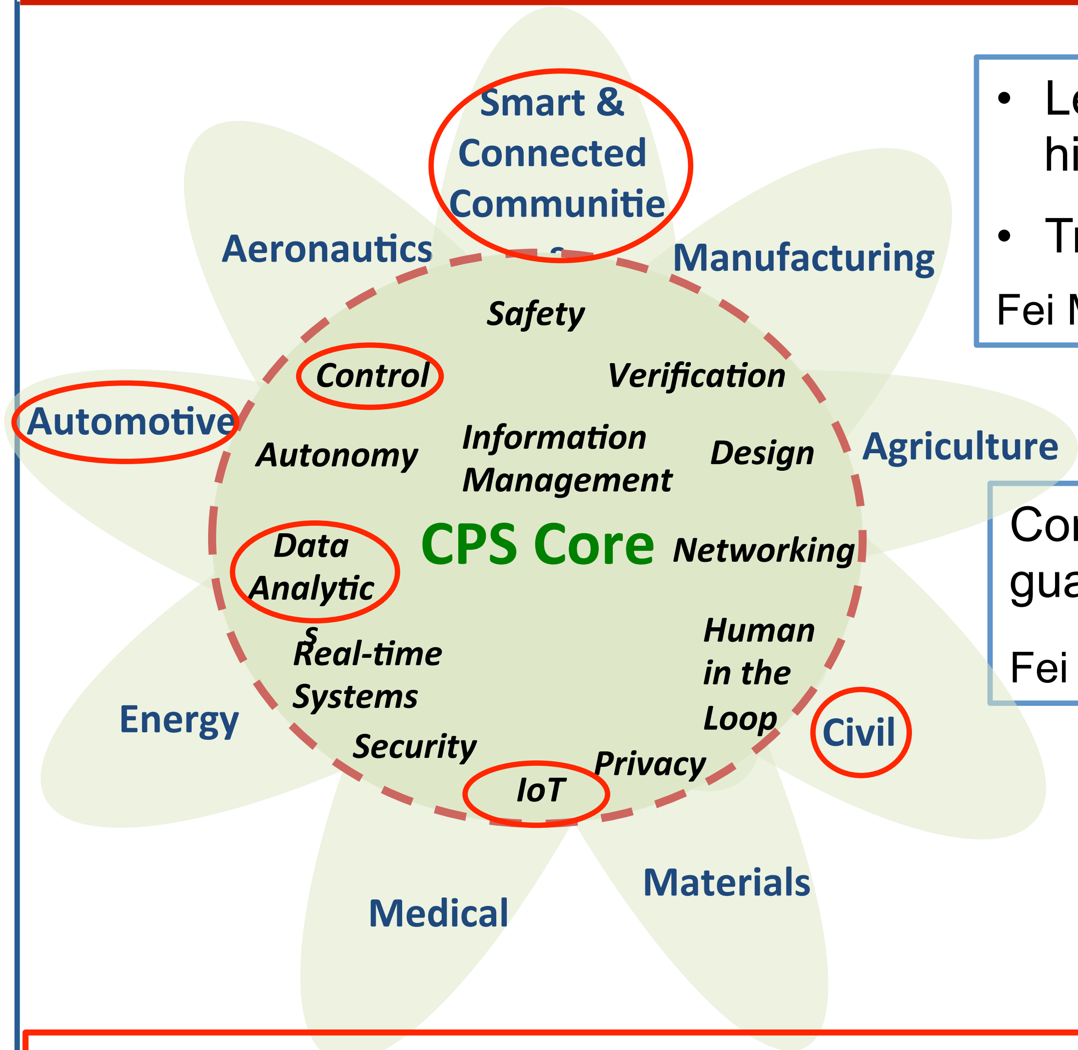

- Learn demand from data; real-time, hierarchical decision-making
- Transportation efficiency

Fei Miao et.al, ICCPS15 best paper finalist

Computationally tractable, probabilistic guarantee under model uncertainties

Fei Miao et.al, ICCPS17;CDC15,TCST16

Challenges with growing **complexity and dynamics: scale, uncertainty**

Future work of CPS Safety and Security

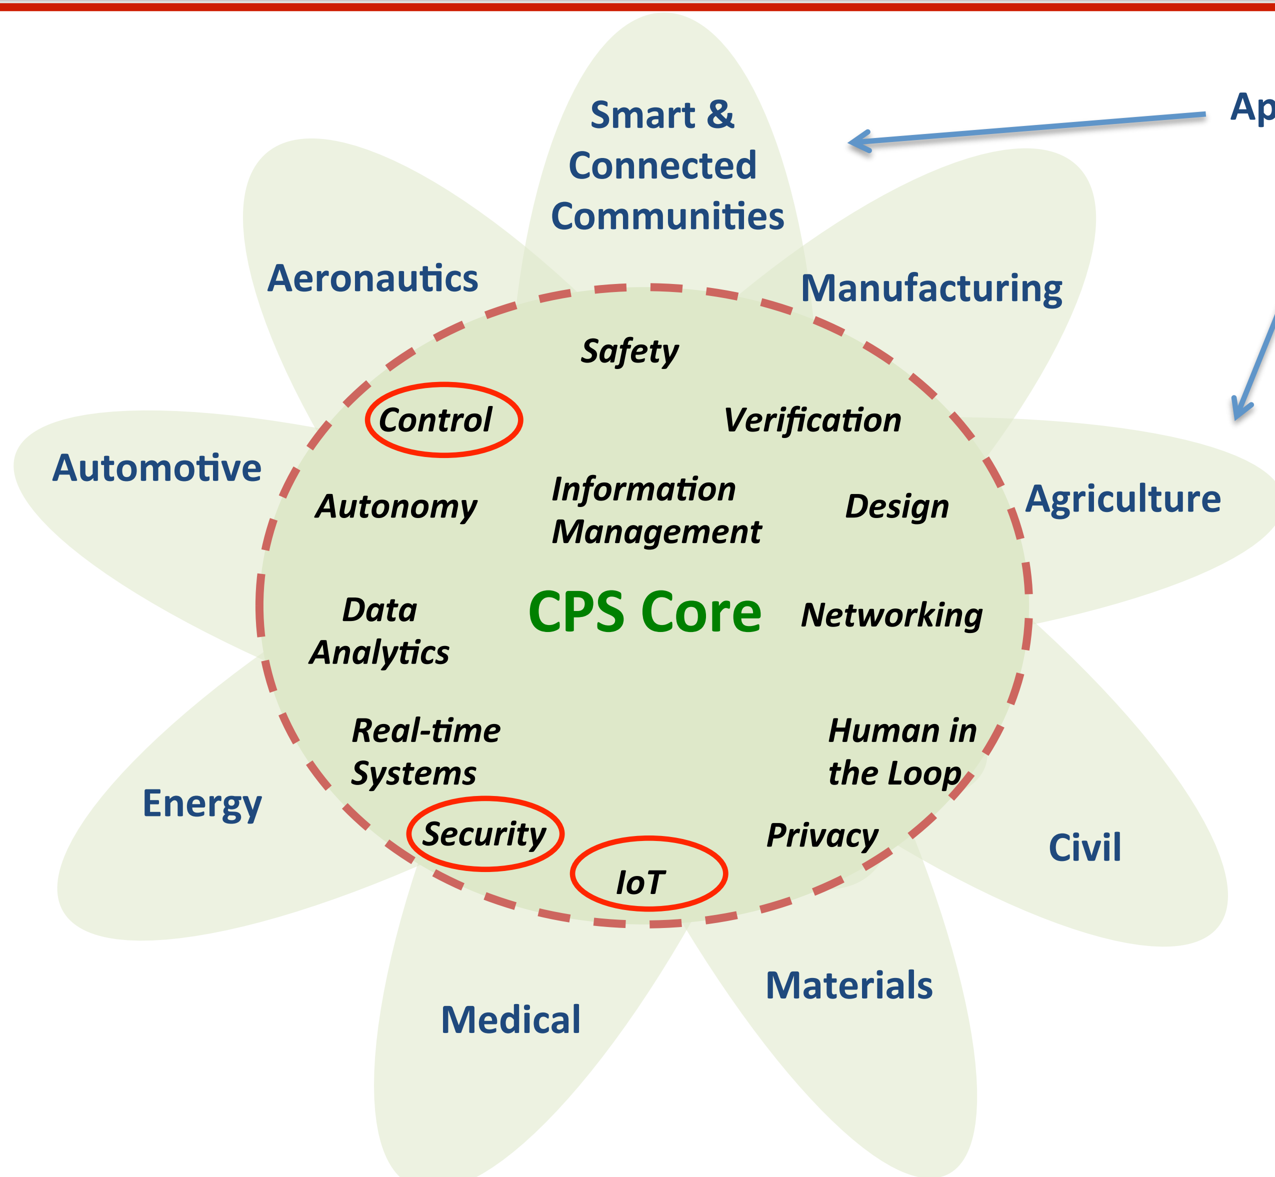

Challenges: Safety, Security and Resilience

- With integration of communication, computation and control, cyber attacks can cause disasters in the physical world
- Knowledge of physical system dynamics helpful for security, resilience


BBC
NEWS TECHNOLOGY
News Sport Weather Capital Culture Autos
Home US & Canada Latin America UK Africa Asia Europe Mid-East Business Health Sci/Environment
25 July 2013 Last updated at 19:04 ET
Car hackers use laptop to control standard car
By Zoe Kleinman
Technology reporter, BBC News

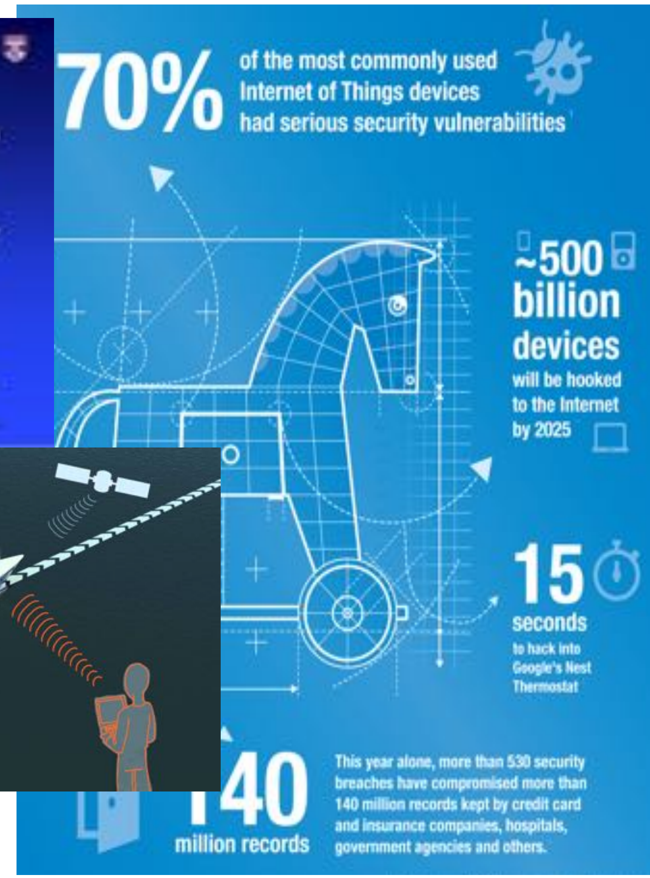

Connected vehicles

CPS Security: Coding for Stealthy Data Injection Attacks Detection

- **Problem**: Stealthy Data Injection Attacks

-Attacker is smarter with the system model knowledge: inject data to communication channel, drive system to unstable state and pass detectors.

-Communication cost for encrypted messages is too large

- **Goal:** A low cost technique to detect stealthy data injection attacks

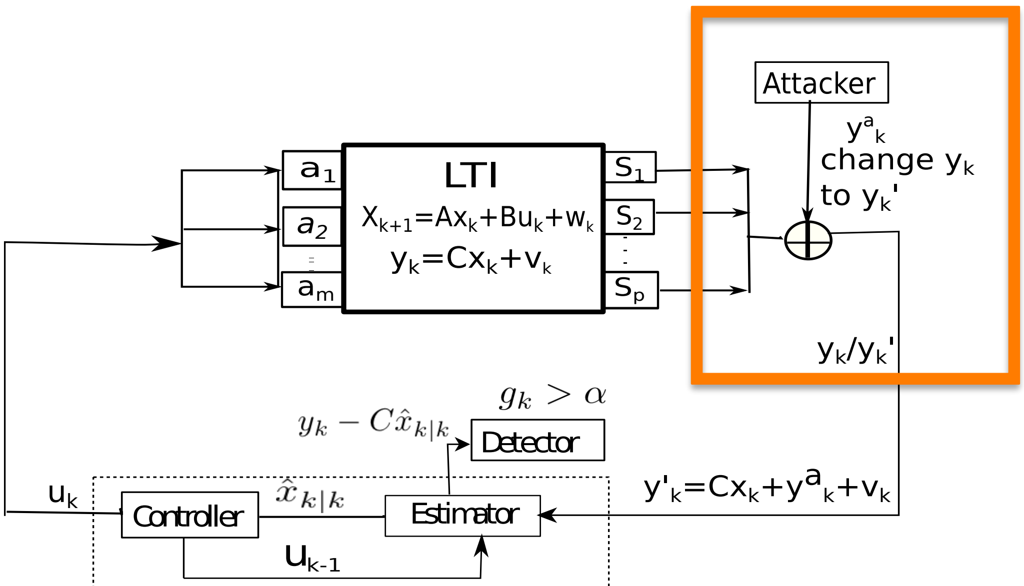

Coding Schemes for Stealthy Data Injection Attacks Detection

$$\mathbf{y}_k = Cx_k + v_k$$

$$Y_k = \Sigma(Cx_k + v_k)$$

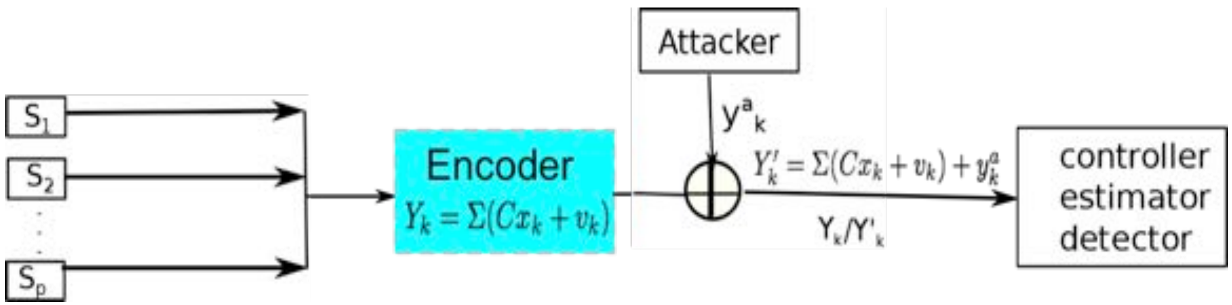

$$Y'_k = \Sigma(Cx_k + v_k) + y^a_k$$

Estimator decode

$$\tilde{y}'_k = \Sigma^{-1}Y'_k = y_k + \Sigma^{-1}y^a_k$$

Low cost: no extra bytes to communicate after coding

***Research Contributions**

-*Analyze* sufficient conditions of feasible, low cost coding

-*Design* an algorithm to calculate a feasible Σ in real-time

-Time-varying coding when the attacker can estimate Σ

*Fei Miao et.al, TCNS 2016. (funded by DARPA)

Dynamic Stochastic Game for Resilient Systems

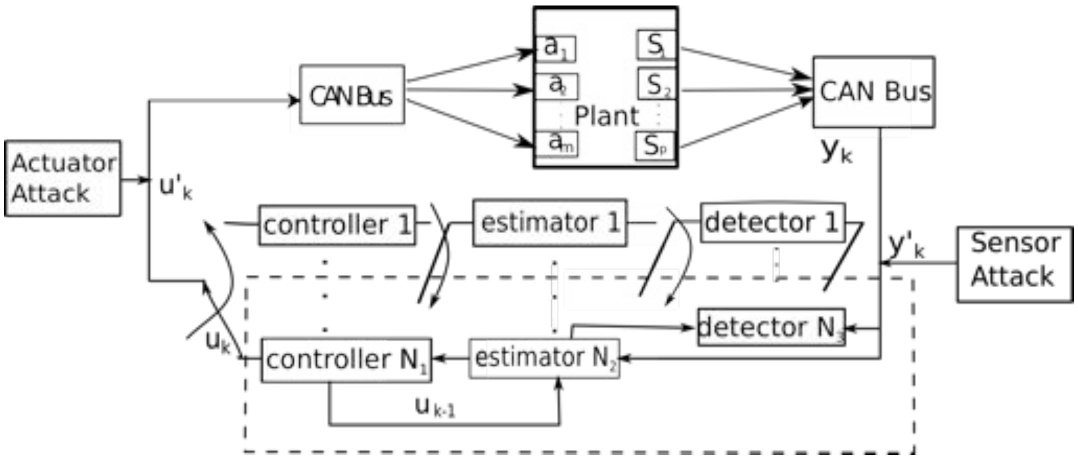

- **Problem:** When an attack happens? What type of attack? Not known!

Higher attack detection rate, higher investment cost in security in general

***Contributions:** dynamic stochastic game for an optimal switching policy between subsystems to balance security overhead and control cost

*Fei Miao et.al, CDC 2013,2014; journal version submitted to Automatica (funded by DARPA)

- **Agenda: safety, efficiency, security for CPSs with focus on smart cities, autonomous transportation systems**

- **Contributions: data-driven CPSs, CPSs/Smart Cities security**

- **Future work**

-Hierarchical decision making based on heterogeneous data information

-Design incentive mechanisms (e.g., dynamic pricing) of users and suppliers for social optimal behavior

-Safety assurance of coordinated control of connected autonomous vehicles

-Security and resiliency of smart cities infrastructure with physical dynamics knowledge, distributed sensor networks