Big Data Meets CPS

Shankar Sastry

University of California, Berkeley

Joint work with Roy Dong, Lillian Ratliff, Henrik Ohlsson, Galina Schwartz, Claire Tomlin, Alex Bayen (Berkeley), Saurabh Amin (MIT) and Alvaro Cardenas (UT Dallas)

Work Supported by NSF CPS Frontiers Grant FORCES

FORCES Review Year 2, June 2014

Outline

Sensor Webs to Action Webs to Resilient CPS

Action Webs & Networked Control Systems (NCS)

Resilient Societal Scale CPS

Mechanism Design for Societal CPS

Disaggregation of CPS Big Data

Incentive Design Via Disaggregation pf CPS Big Data

Security

Financial Attacks on Societal CPS Systems

CPS Vulnerabilities

1. Threat assessment
2. Attack diagnosis
3. Resilient control

CPS Big Data and Privacy

Privacy in Societal CPS Systems

Conclusions and Future Research

Outline

Sensor Webs to Action Webs to Resilient CPS

Action Webs & Networked Control Systems (NCS)

Resilient Societal Scale CPS

Mechanism Design for Societal CPS

Disaggregation of CPS Big Data

Incentive Design Via Disaggregation pf CPS Big Data

Security

Financial Attacks on Societal CPS Systems

CPS Vulnerabilities

1. Threat assessment
2. Attack diagnosis
3. Resilient control

CPS Big Data and Privacy

Privacy in Societal CPS Systems

Conclusions and Future Research

The swarm at the edge of the cloud

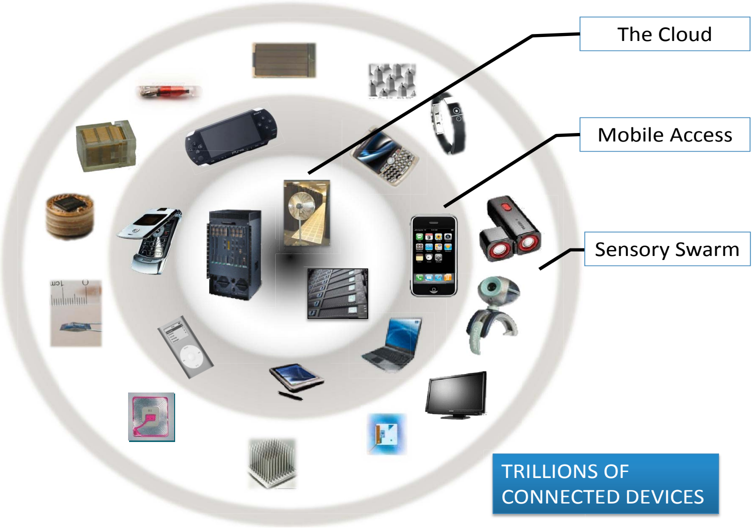

The Cloud

Mobile Access

Sensory Swarm

TRILLIONS OF CONNECTED DEVICES

Source: J. Rabaey [ASPDAC'08]

Wireless Sensor Webs Everywhere

Change detection: Thresholds, phase transitions, anomalies

- Security systems
- Health care
- Wildfire detection
- Fault diagnosis
- Tracking & surveillance

Action Webs in CPS Infrastructures

Supervisory Control & Data Acquisition (SCADA)

- Robust estimation
  - Noisy measurements
  - Lossy communication
- Real-time control
  - Safety
  - Performance

COTS IT for SCADA

- Cost ↓, Reliability ↑
- Digital and IP based: New vulnerabilities!
- Reliability ⇒ Security

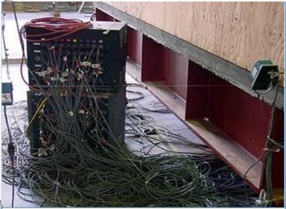

Wired networks are costly to maintain

Typical industrial infrastructure ~ $10B

Source: Emerson case study

Action Webs

Observe and infer for planning and modifying action

- Dealing with uncertainty
- Tasking sensors
- Programming the ensemble
- Multiple objectives
- Embedding humans

Example: Building energy management

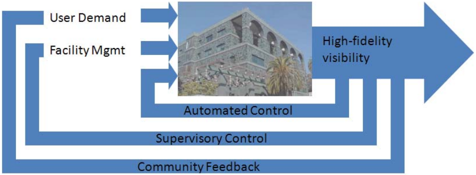

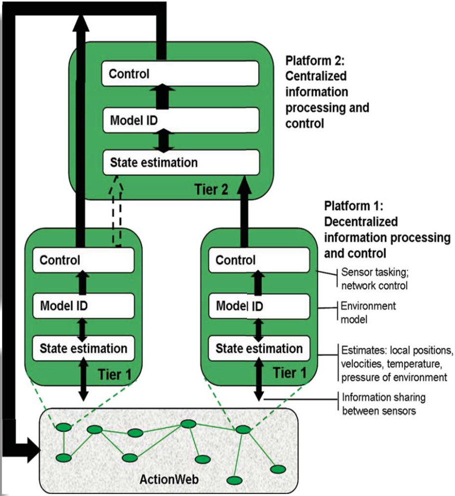
Control
Model ID
State estimation
Tier 2
Platform 2: Centralized information processing and control
Platform 1: Decentralized information processing and control
Control
Model ID
State estimation
Tier 1
Control
Model ID
State estimation
Tier 1
Sensor tasking; network control
Environment model
Estimates: local positions, velocities, temperature, pressure of environment
Information sharing between sensors
ActionWeb

Courtesy: Claire Tomlin

From Action Webs to Resilient CPS

Resilient/High Confidence Networked Control

- Fault-tolerant networked control
  - Limits on stability, safety, & optimality
  - Scalable model predictive control
- Security & Resilient Control
  - Availability, Integrity, & Confidentiality
  - Graceful degradation
- Economic Incentives
  - Incentive Design for investing in security
  - Interdependent Risk Assessment & Cyber Insurance

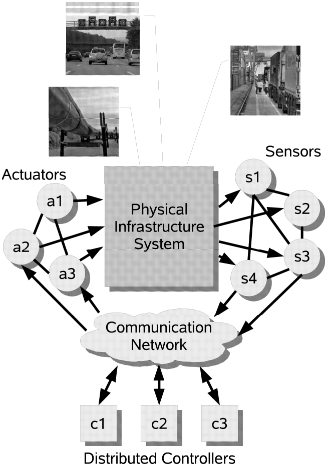

Societal Scale CPS

A complex collection of sensors, controllers, compute nodes, and actuators that work together to improve our daily lives

- From very small: Ubiquitous, Pervasive, Disappearing, Perceptive, Ambient
- To very large: Always Connectable, Reliable, Scalable, Adaptive, Flexible

Emerging Service Models

- Building energy management
- Automotive safety and control
- Management of metropolitan traffic flows
- Distributed health monitoring
- Smart Grid

Economical, Social and Environmental Drivers

Electricity Grid:

- Smart meters are being used for demand response currently. However, the potential of smart meters go far beyond D/R. The market for energy analytics in the smart grid is estimated to be worth $9.7 billion by 2020

Transportation Systems:

- It is estimated that more than 4.2 billion hours are lost sitting in traffic, resulting in 2.8 billion gallons of wasted fuel and costing more than $ 87 billion dollars annually. By utilizing CPS analytics in intelligent transportation systems (ITS) we can actively manage our transportation network to improve safety, efficiency, and multimodal connectivity.

Other Critical Infrastructures:

- Healthcare systems, Water systems, Natural gas and Oil, ...

Outline

Sensor Webs to Action Webs to Resilient CPS

- Action Webs & Networked Control Systems (NCS)
- Resilient Societal Scale CPS

Mechanism Design for Societal CPS

- Disaggregation of CPS Big Data
- Incentive Design Via Disaggregation pf CPS Big Data

Security

- Financial Attacks on Societal CPS Systems
- CPS Vulnerabilities

1. Threat assessment
2. Attack diagnosis
3. Resilient control

CPS Big Data and Privacy

- Privacy in Societal CPS Systems

Conclusions and Future Research

Need for Incentives in Societal CPS

- There is often a substantive gap between competitive Nash equilibria and the social planner's optimum (Hal Varian, et al).
- Due to information asymmetries and misaligned objectives, the actions taken by agents in S-CPS are not socially optimal.
- Incentives are the natural mechanism for aligning agents so that they behave in a socially optimal way.

In Energy CPS:

- Consumers are not well informed about their energy consumption patterns; utilities can use incentives to motivate consumers to use less energy.

Iin Transportation CPS:

- Drivers often travel at peak hours; incentives can be used to encourage drivers to shift their departure time for some reward resulting in overall reduced congestion.

Motivations for Incentive Design in Energy CPS

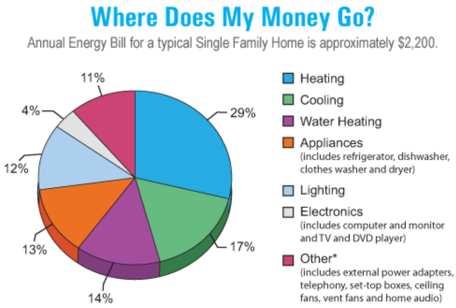

www.energystar.gov

- Studies have shown that providing device-level feedback on power consumption patterns to energy users can modify behavior and improve energy efficiency.
- Provide **incentives** in the form of **rebates and monetary rewards** focusing on devices that fall into largest consumption categories in order to reduce energy consumption.

Creyts, et al., Reducing U.S. greenhouse gas emissions: How much at what cost? U.S. Greenhouse Gas Abatement Mapping Initiative, 2007.
Laitner, et al., Examining the scale of the behaviour energy efficiency continuum. European Council for an Energy Efficient Economy, 2009.
Perez–Lombard, et al., A review on buildings energy consumption information. Energy and Buildings, 2008.

Incentive Design for Energy CPS Systems

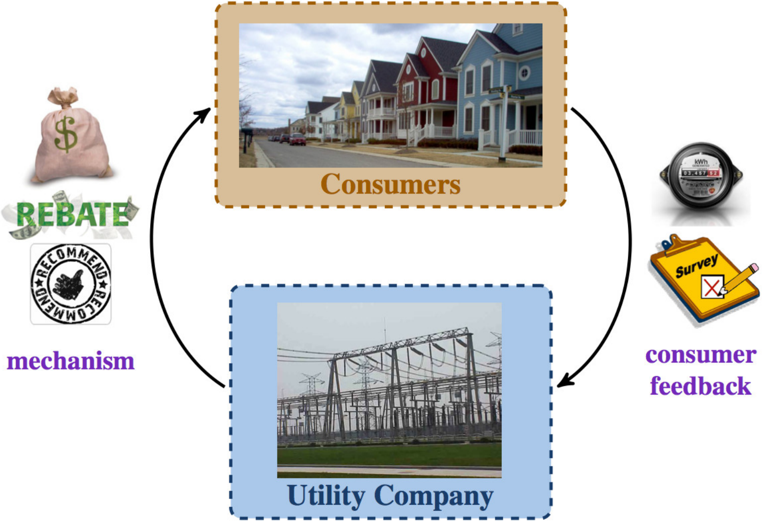

Consumers

Utility Company

mechanism

consumer feedback

Energy Disaggregation or Non Intrusvie Load Monitoring

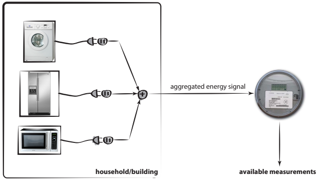

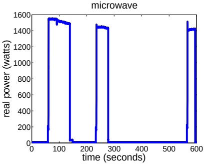

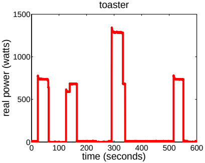

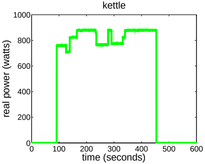

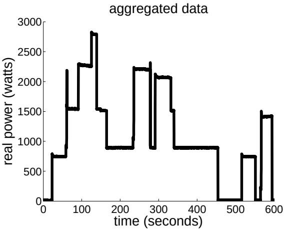

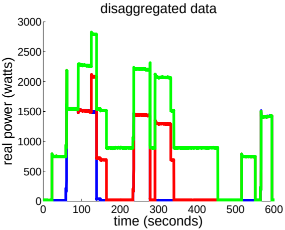

Approaches to Disaggregation

Hidden Markov Models

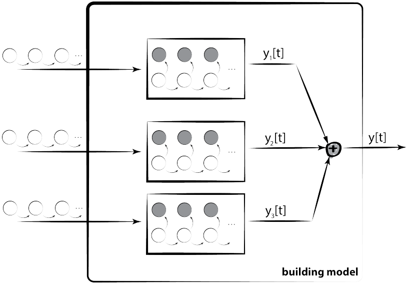

- Unsupervised
- Requires tuning of parameters.
- The states are constant wattage levels; usage patterns and device signatures are encoded in transition probabilities.

Sparse Coding

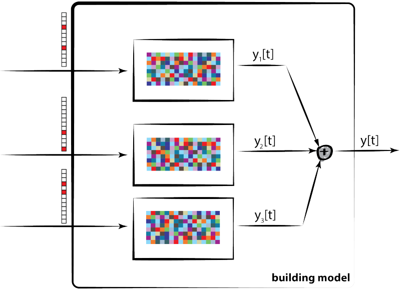

- Supervised
- Assume inputs are sparse.
- Reconstruct the aggregate signal by selecting as few signatures as possible from a library.

A New Systems Framework for Disaggregation

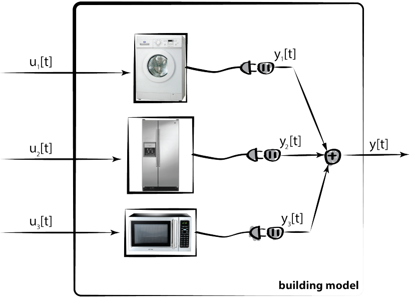

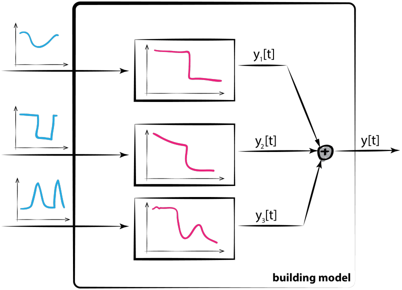

We learn *dynamical models* for the devices!:

- We have theoretical results guaranteeing recovery of the most likely device consumption signals.
- We also learn dynamics of devices, which is useful for other Smart Grid operations.

Disaggregation Summary

- We take a semi-supervised, systems approach to energy disaggregation by identifying dynamical models of the devices.
- We utilize system dynamics and priors on device usage in our approach to energy disaggregation to provide a way to formulate and regularize an otherwise naturally ill-posed problem.
- In the energy disaggregation problem, questions of consumer **privacy** arise naturally since a fundamental part of disaggregation is inference about consumer behavior.

Incentive Design via Energy Disaggregation

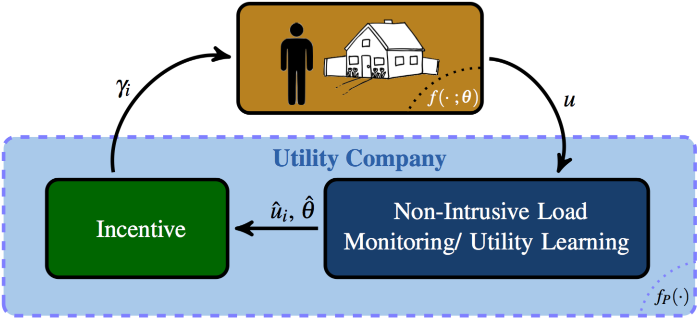
γ_i
$f(\cdot\,;\theta)$
u
Utility Company
Incentive
$\hat{u}_i,\ \hat{\theta}$
Non-Intrusive Load Monitoring/ Utility Learning
$f_P(\cdot)$

- Given an upper bound on the probability of distinguishing devices, the utility company can design incentives that induce the consumer to use the desired amount of energy for a device with an error bound derived from the probability of distinguishing devices.

Ratliff, Dong, Ohlsson, Sastry. Behavior Modification and Utility Learning via Energy Disaggregation. IFAC, 2014.

Incentive Design Mechanisms introduce New Vulnerabilities!

- In regulated markets, utility companies are incentivized to reduce the overall consumption of their consumer base.
- Demand response programs incentivize customers to shift their demand thereby alleviating inaccuracies in load forecasting. Device-level incentives can be designed via non-intrusive load monitoring.
- Introduces new vulnerabilities by allowing **adversarial** agents who may **spoof** their energy signal or otherwise **disrupt** the energy system.

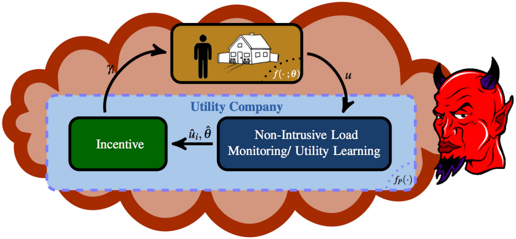
γ_i
$f(\cdot\,;\theta)$
u
Utility Company
Incentive
$\hat{u}_i, \hat{\theta}$
Non-Intrusive Load Monitoring/ Utility Learning
$f_P(\cdot)$

Outline

Sensor Webs to Action Webs to Resilient CPS

Action Webs & Networked Control Systems (NCS)

Resilient Societal Scale CPS

Mechanism Design for Societal CPS

Disaggregation of CPS Big Data

Incentive Design Via Disaggregation pf CPS Big Data

Security

Financial Attacks on Societal CPS Systems

CPS Vulnerabilities

1. Threat assessment
2. Attack diagnosis
3. Resilient control

CPS Big Data and Privacy

Privacy in Societal CPS Systems

Conclusions and Future Research

Security: Revenue Protection

- Non–technical losses are caused by actions external to the power system such as theft, non–payment by consumers, or errors in accounting.
- Both faults and theft can be the result of **adversarial** agents acting on the system, e.g. ***spoofing*** energy signals.

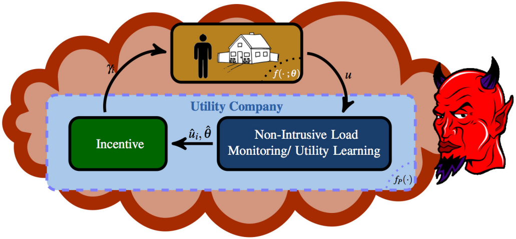
γ_i
$f(\cdot\,;\theta)$
u
Utility Company
Incentive
$\hat{u}_i, \hat{\theta}$
Non-Intrusive Load Monitoring/ Utility Learning
$f_P(\cdot)$

Reducing Technical and Non–Technical Losses in the Power Sector. World Bank Group Technical Report, 2009.

Security: Revenue Protection

- In conjunction with C3 Energy, we have developed algorithms for detecting non-technical loss in the electricity grid.
- We trained and tested our algorithms on data from a utility company with over 30 million customers.
- The data included time-series consumption data and meter events from AMIs, weather data, customer demographics, and work orders.
- We identified ∼50 features and selected those that were highly correlated with anomalous or tampering events.
- We utilized machine learning algorithms to develop a model for identifying non-technical loss.

C3 Energy

CPS Attacks

Maroochy Shire sewage plant *(2000)*

Los Angeles traffic control *(2008)*

Tehama Colusa canal system *(2007)*

Cal-ISO power system computers *(2007)*

NCS/CPS security concerns

Attackers

- Malicious insiders
- Computer hackers
  - Cyber criminals
  - Cyber warriors
  - Hacktivists
  - Rogue hackers
  - Corporate spies

Stuxnet worm

- Targets SCADA systems
- Four zero-day exploits, antivirus evasion techniques, p-2-p updates, network infection routines
- Reprograms *Programmable Logic Controller (PLC)* code

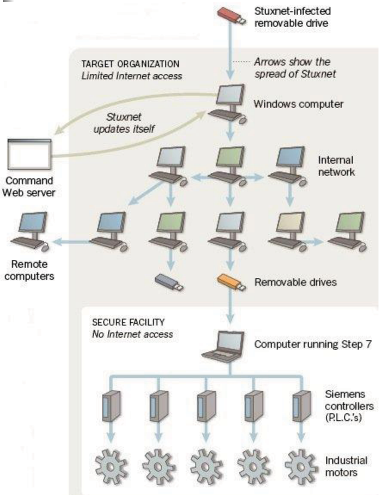

Source: Symantec, NYT

Resilient Control for NCS

1. Threat assessment
   - How to model attacker and his strategy?
   - Consequences to the physical infrastructure
2. Attack diagnosis
   - How to detect manipulations of sensor-control data?
   - Stealthy [undetected] attacks
3. Resilient control
   - Design of resilient control algorithms
   - Tradeoffs between performance and containment
   - Incentive mechanisms to improve NCS reliability & security

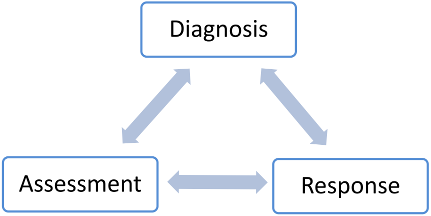

Threat assessment

- How to model attacker and his strategy?
- Consequences to the physical infrastructure

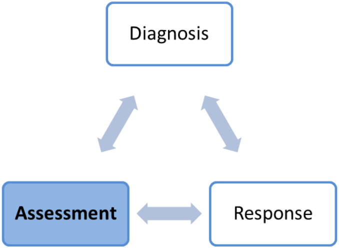

Field operational test on the Gignac canal network
[Amin, Litrico, Sastry, Bayen. HSCC'10]

Models of deception and denial-of-service (DoS) attacks
[Amin, Cárdenas, Sastry. HSCC'09]

Assessment for Tennessee Eastman process control system (TE-PCS)
[Cárdenas, Amin, Lin, Huang, Sastry. ASIACCS'11]

Gignac water canal network

SCADA components

- Level & velocity sensors
- PLCs & gate actuators
- Wireless communication
- Multiple stakeholders

Communication station

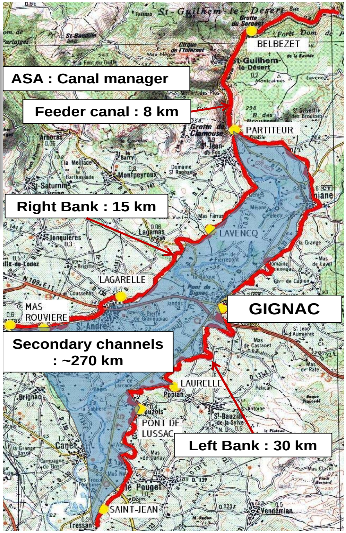

Map of Gignac canal

Presented by permission from Cemagref, France

Gignac canal network

Physical infrastructure

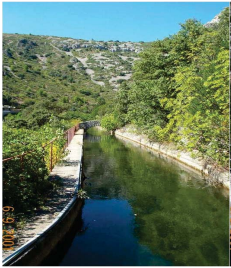

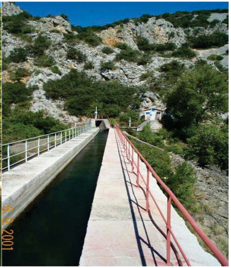

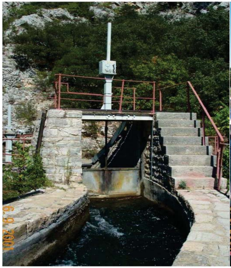

Cyber infrastructure

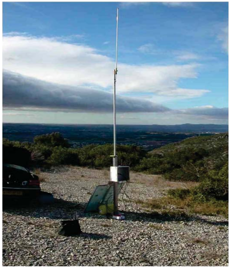

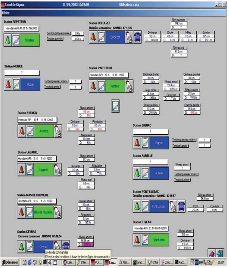

h

Reported attacks on water SCADA systems

Gignac canal system attacks

- Stealing water by compromising sensors
- Tampering PLCs
- Theft of solar panels

Other SCADA vulnerabilities

- Time between telemetry requests can be used for malicious traffic injection
- Encryption provides confidentiality but does not provide data integrity

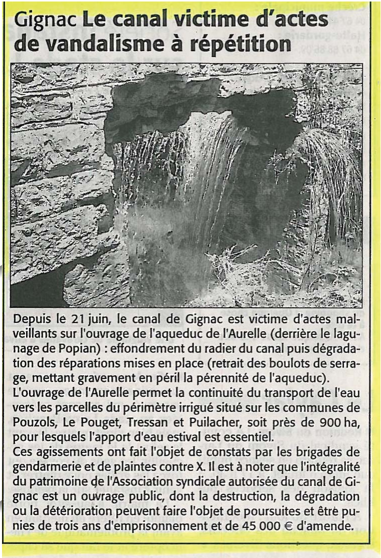

Gignac **Le canal victime d'actes de vandalisme à répétition**

Depuis le 21 juin, le canal de Gignac est victime d'actes malveillants sur l'ouvrage de l'aqueduc de l'Aurelle (derrière le lagunage de Popian) : effondrement du radier du canal puis dégradation des réparations mises en place (retrait des boulots de serrage, mettant gravement en péril la pérennité de l'aqueduc).
L'ouvrage de l'Aurelle permet la continuité du transport de l'eau vers les parcelles du périmètre irrigué situé sur les communes de Pouzols, Le Pouget, Tressan et Puilacher, soit près de 900 ha, pour lesquels l'apport d'eau estival est essentiel.
Ces agissements ont fait l'objet de constats par les brigades de gendarmerie et de plaintes contre X. Il est à noter que l'intégralité du patrimoine de l'Association syndicale autorisée du canal de Gignac est un ouvrage public, dont la destruction, la dégradation ou la détérioration peuvent faire l'objet de poursuites et être punies de trois ans d'emprisonnement et de 45 000 € d'amende.

Courtesy: C. Hugodot, Manager

Cyber-attack on the Avencq canal pool

Successful attack

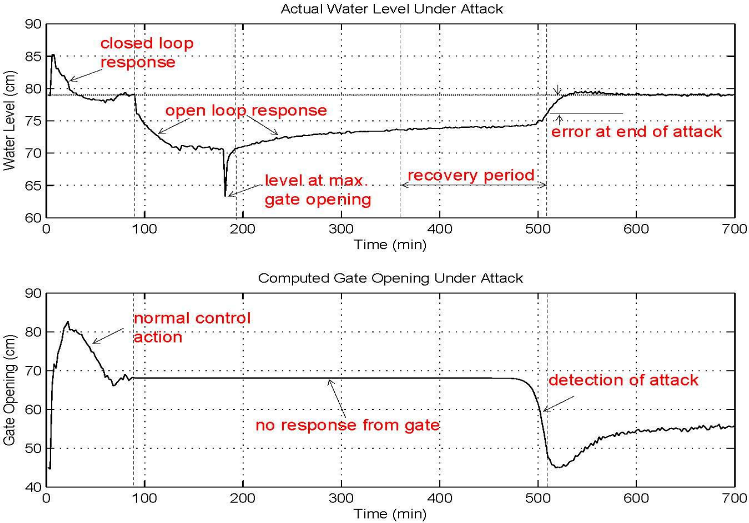

Taxonomy of Attacks on NCS

Cyber Attacks

SCADA Manager [IT Security] **A6**

- Unauthorized access, Viruses

Supervisory Control **A3-A5**

- Deception: set-point change, parameter substitution
- Denial-of-Service (DoS): network flooding, process disruption

Regulatory Layer **A1-A2**

- Deception: compromise of measurements & controls, spoofing, replay
- DoS: jamming, ↑ comm. latency

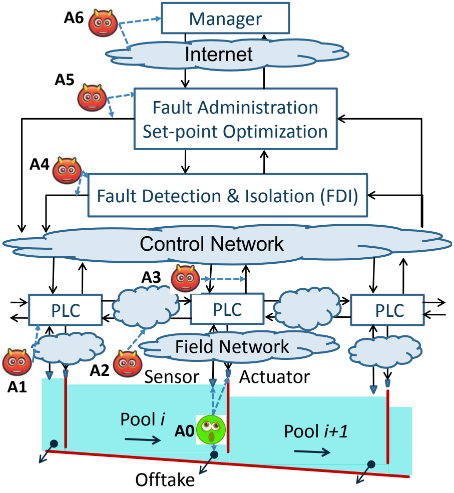

Physical Faults [Control th.] **A0**

- Sensor-actuator faults
- Unauthorized leaks

Attack diagnosis

- How to detect manipulations of sensor-control data?
- Stealthy [undetected] attacks

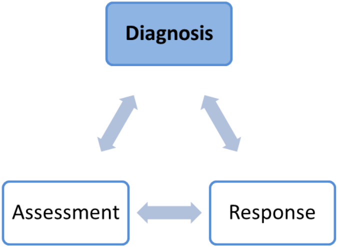

Observer-based diagnosis for Gignac SCADA system
[Amin, Litrico, Sastry, Bayen. IEEE TCST'11]

Non-parametric CUSUM statistic based diagnosis for TE-PCS
[Cárdenas, Amin, Sastry, et.al. ASIACCS'11]

Study of stealthy attacks on power system state estimators
[Teixeira, Amin, Sandberg, Johansson, Sastry. IEEE CDC'10]

Attacks on supervisory control layer

Supervisory Layer Attacks **A3**

- Deception: set-point change, parameter substitution
- Denial-of-Service (DoS): network flooding, process disruption

Physical Faults/Attacks **A0**

- Sensor-actuator faults
- Unauthorized withdrawals

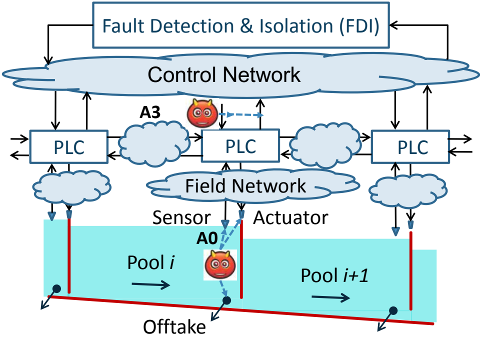

Design of a model-based diagnosis scheme

Recommendations on Security Diagnosis

Recommendations to the European Commission on Canal Automation & the Cemagref Research Institute

- Enhanced model (redundancy) improves detection
- Sensors located closer to the offtakes are critical
- Localized sensor attacks do not lead to global degradation
- Multiple pool sensor attacks can evade detection [stealth]

Attack diagnosis for [other] SCADA systems

Process control

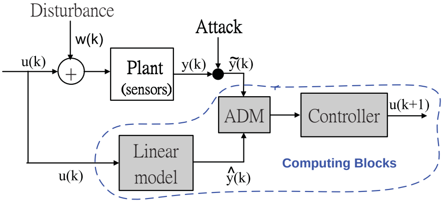

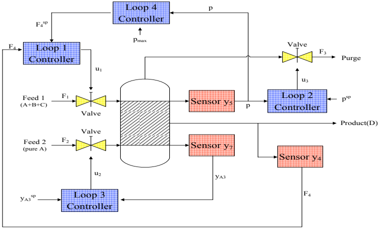

[Cárdenas, Amin, Lin, Huang, Sastry. ASIACCS'11]

Power transmission

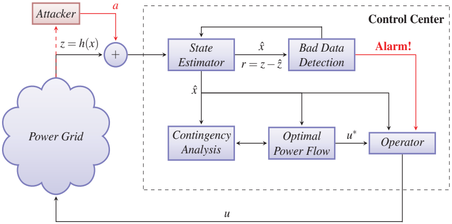

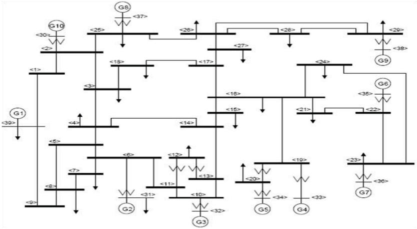

[Teixeira, Amin, Sandberg, Johansson, Sastry. IEEE CDC'10]

Resilient control

- Design of resilient control algorithms?
- Fundamental limitations & interdependent security

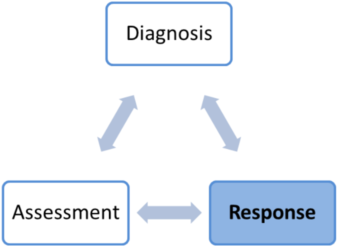

Stability of hyperbolic PDEs under switching boundary control
[Amin, Hante, Bayen. IEEE TAC'10]

Incentives to secure under network induced interdependent risks
[Amin, Schwartz, Sastry. GameSec'10]

Safety-preserving control for stochastic systems under comm. losses
[Amin, Cárdenas, Sastry. HSCC'09]

Attacks on regulatory control layer

Regulatory layer **A1-A2**

- Deception: compromise of measurements & controls
- DoS: jamming, ↑ latency

Physical faults or attacks **A0**

- Sensor-actuator faults
- Unauthorized withdrawals

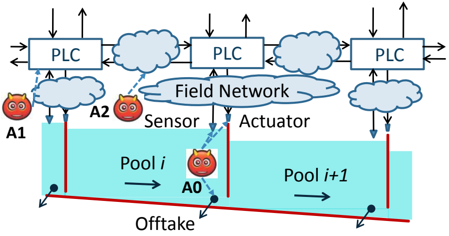

Switching attacks can lead to instability!

Outline

Sensor Webs to Action Webs to Resilient CPS

- Action Webs & Networked Control Systems (NCS)
- Resilient Societal Scale CPS

Mechanism Design for Societal CPS

- Disaggregation of CPS Big Data
- Incentive Design Via Disaggregation pf CPS Big Data

Security

- Financial Attacks on Societal CPS Systems
- CPS Vulnerabilities

  1. Threat assessment
  2. Attack diagnosis
  3. Resilient control

CPS Big Data and Privacy

- Privacy in Societal CPS Systems

Conclusions and Future Research

Privacy Issues

In energy CPS:

- Disaggregation can be used to infer a person's schedule.
- Disaggregation can be used to infer specific consumption of entertainment

In transportation CPS:

- Smart phone data can be used to supplement sensor data in intelligent transportation systems.
- Driver intent and transit patterns can be inferred from disaggregated traffic data or from GPS data.

Privacy Issues in Energy CPS

Data minimization principle (NISTIR 7628)

Limit the collection of data to only that necessary for Smart Grid operations, including planning and management, improving energy use and efficiency, account management, and billing.

But can we quantify these ideas?

- Quantify trade-off between amount of data and performance of Smart Grid ?
- Analyze the amount of private information which can be inferred from data.

The power consumption signal is not private in and of itself. It is what we can *infer* that is private.

- Household occupancy.
- Behavioral patterns.
- Which devices are present in a household, and their usage.

Privacy Formulation

If an AMI measures only the aggregate power consumption, what can we infer?
Recall the problem of *energy disaggregation*:

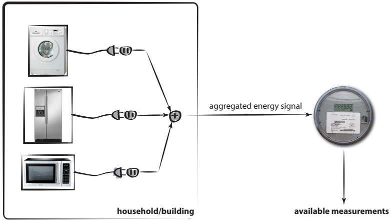

This leads to our notion of **privacy**. In this model, our **adversary**:

- Observes the AMI signals
- Has knowledge of what devices are in the house
- Knows the dynamics and signatures of these devices

What can he infer?

Privacy Metrics

- We can place an upper bound on the probability of successfully distinguishing devices.
- These bounds are properties of the disaggregation problem itself and hold for any algorithm.
- This upper bound can thus act as a **guarantee for privacy**.

- Inputs u_1 and u_2. For each input, the aggregate power consumption signal follows distributions F_1 and F_2, respectively.
- For simplicity, we assume these distributions to be Gaussian.
- F_1 and F_2 have means μ_1 and μ_2, and both distributions have the same covariance $\sigma^2 I$.

Dong, Ratliff, Ohlsson, Sastry. Fundamental Limit of Non-Intrusive Load Monitoring. HiCoNS, 2014.
Ratliff, Dong, Ohlsson, Cárdenas, Sastry. Privacy and Customer Segmentation in the Smart Grid. Submitted to CDC, 2014.

Privacy Theorems: Gaussian case

Theorem

If $\mathbf{u} = \mathbf{v_0}$ with probability $1/2$ and $\mathbf{u} = \mathbf{v_1}$ with probability $1/2$, and if $G(\mathbf{v_0}) \sim N(\mu_0, \Sigma), G(\mathbf{v_1}) \sim N(\mu_1, \Sigma)$ are independent, then for any NILM algorithm S and deciding function I:

$$\mathbb{P}((I \circ S)(G(\mathbf{u})) = \mathbf{u}) \leqslant \frac{1}{2}\left(1 - \operatorname{erf}\left(\frac{-\frac{1}{\|a\|_2}(a^{\mathscr{T}}\mu_0 + b)}{\sqrt{2\sigma^2}}\right)\right)$$

with

$$a^{\mathscr{T}} = (\mu_0 - \mu_1)^{\mathscr{T}}\Sigma^{-1}$$

$$b = \frac{1}{2}\left(\mu_1^{\mathscr{T}}\Sigma^{-1}\mu_1 - \mu_0^{\mathscr{T}}\Sigma^{-1}\mu_0\right)$$

$$\sigma^2 = \frac{1}{\|a\|_2^2}a^{\mathscr{T}}\Sigma a = \frac{(\mu_0 - \mu_1)^{\mathscr{T}}\Sigma^{-1}(\mu_0 - \mu_1)}{(\mu_0 - \mu_1)^{\mathscr{T}}\Sigma^{-2}(\mu_0 - \mu_1)}$$

Generalizations to M inputs is easy!

Privacy versus Performance Tradeoffs

Quantify the trade-off between the amount of data and the performance of Smart Grid operations.

- For a direct load control (DLC) application, consider the simple model:

$$x_{k+1} = x_k + u_k + \mu_k + d_k$$

  - $x_k \in \mathbb{R}$ is the power consumption of a unit (e.g. HVAC, sector of grid) at time k.
  - $u_k \in \mathbb{R}$ represents the DLC signal.
  - $\mu_k \in \mathbb{R}$ represents the affine term which generates our nominal demands.
  - d_k represents the disturbance.
- A privacy-aware sampling policy:
  - DLC policies must be employed to return the power consumption to the nominal demand.
  - Controller receives measurements every N time steps yet issues control commands at every time step.
  - $\mathscr{H}_\infty$: measure of how much the uncertainty in demand gets amplified.

Downsampling policies

Analyze the amount of private information which can be inferred from data.

We can use our results to analyze how privacy increases as N, the downsampling rate, increases.

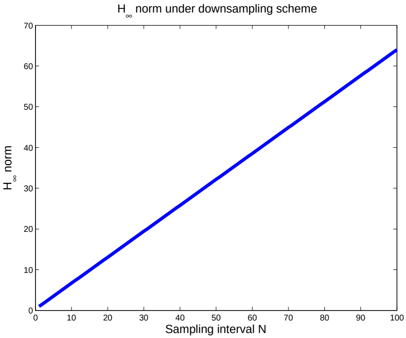

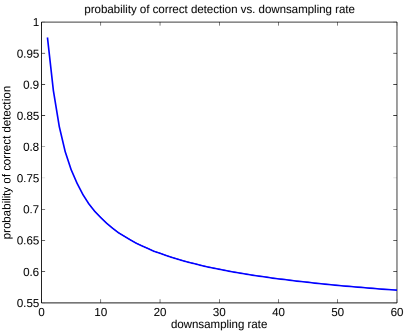

Ratliff, Dong, Ohlsson, Cárdenas, and Sastry 2014, *Under review*.

Privacy Contracts

The utlity company selects ***down-sampling*** as a privacy preserving metering policy. Then,

- The performance of direct load control is dependent on the sampling rate.
- Higher sampling rates result in improved performance of the direct load control scheme with diminishing returns.

Contract Design:

Utility company can design screening mechanisms to obtain the consumer's privacy preferences (unknown type) by offering contracts where privacy is the good and privacy-setting is the quality of the good.

Privacy Contracts: Two-Type Model

- The contracting device that the utility company can use is the privacy setting that they offer to the consumer.
- There are two privacy settings offered: x_L, x_H such that $x_L \leqslant x_H$, $x_L, x_H \in \mathbb{R}$.
- The consumer type θ is unknown to the utility company.
- We consider two types: $\theta \in \{\theta_L, \theta_H\}$ where θ represents how much the consumer values privacy.

Privacy Contracts: Two-Type Model

- Utility company announces price t for choosing privacy level x. The consumer's utility is equal to zero if he does not select a privacy setting, and it is

$$U(x,\theta) - t \geqslant 0 \qquad \text{(\textbf{Individual Rationality})}$$

- Assumption: U is increasing in (x,θ).
- **Incentive-compatibility**: all of the participants fare better when they truthfully reveal any private information asked for by the mechanism:

$$U(x_H,\theta_H) - t_H \geqslant U(x_L,\theta_H) - t_L$$

$$U(x_L,\theta_L) - t_L \geqslant U(x_H,\theta_L) - t_H$$

Privacy Contracts: Utility Company

Utility company's utility function:

$$v(x,t) = t - g(x)$$

where $g(x)$ is the unit cost resulting from the privacy setting x and is a strictly increasing, continuous function.

Screening Problem

$$\begin{aligned}\max_{\{(t_L,x_L),(t_H,x_H)\}} & (1-p)v(x_L,t_L) + pv(x_H,t_H) \\ \text{s.t.} & U(x_i,\theta_i) - t_i \geqslant 0,\ \ i = H, L \\ & U(x_H,\theta_H) - t_H \geqslant U(x_L,\theta_L) - t_L \\ & U(x_L,\theta_L) - t_L \geqslant U(x_H,\theta_L) - t_H\end{aligned}$$

where $p = \mathrm{Prb}(\theta = \theta_H) = 1 - \mathrm{Prb}(\theta = \theta_L) \in (0,1)$ (prior on distribution of types in the population)

Simplification of the Contract Design Problem

- Depending on the form of $U(x, \theta)$ and $g(x)$ this problem can be difficult to solve.
- Assumption: $U(x, \theta_H) - U(x, \theta_L)$ is increasing in x (marginal gain from raising the value of the privacy setting).
- The individual rationality and incentive compatibility constraints reduce to

$$\begin{aligned} t_H - t_L &= U(x_H, \theta_H) - U(x_L, \theta_H) \\ t_L &= U(x_L, \theta_L) \end{aligned}$$

Reduced screening problem

$$\begin{cases} \max_{x_H}\{U(x_H, \theta_H) - g(x_H)\} \\ \max_{x_L}\{-p(U(x_L, \theta_H) - U(x_L, \theta_L)) + (1-p)(U(x_L, \theta_L) - g(x_L))\} \end{cases}$$

Ratliff, Dong, Ohlsson, Cárdenas, Sastry. Privacy and Customer Segmentation in the Smart Grid. Submitted to CDC, 2014.

Privacy Contracts for Direct Load Control

Recall DLC example: as you decrease the sampling rate the performance degrades, i.e. the $\mathscr{H}_\infty$ norm increases, and it degrades in a linear way.

- Let $g(x) = \zeta x$, $0 < \zeta < \infty$.
- Utility company's utility function $v(x,t) = t - g(x)$.
- $p = P(\theta = \theta_H)$.
- Consumer's utility $U(x,\theta) = \frac{1}{2}(\bar{x}^2 - (x-\bar{x})^2)\theta$
- Optimal quality:

$$(x_H^*, x_L^*) = \left(\bar{x} - \frac{\zeta}{\theta_H}, \left[\bar{x} + \frac{(1-x)\zeta}{(p\theta_H - \theta_L)}\right]_+\right)$$

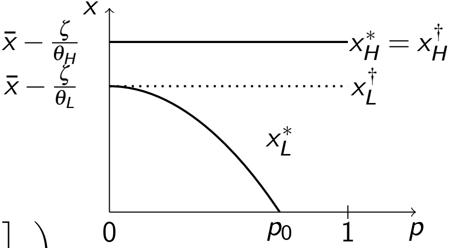

Ratliff, Dong, Ohlsson, Cárdenas, Sastry. Privacy and Customer Segmentation in the Smart Grid. Submitted to CDC, 2014.

Privacy Contracts in Energy CPS

We view consumers as utility maximizers and privacy as a good!

- Based on their valuation of privacy as a good, consumers can select the quality of the service contract with the utility company.
- Electricity service is offered as a product line differentiated according to privacy where consumers can self-select the level of privacy that fits their needs and wallet.
- Revenue Protection: Utility company has a right to find out if a consumer is hiding behind privacy to steal electricity.

Outline

Sensor Webs to Action Webs to Resilient CPS

Action Webs & Networked Control Systems (NCS)

Resilient Societal Scale CPS

Mechanism Design for Societal CPS

Disaggregation of CPS Big Data

Incentive Design Via Disaggregation pf CPS Big Data

Security

Financial Attacks on Societal CPS Systems

CPS Vulnerabilities

1. Threat assessment
2. Attack diagnosis
3. Resilient control

CPS Big Data and Privacy

Privacy in Societal CPS Systems

Conclusions and Future Research

Big Data in Resilient CPS

Resilient CPS Systems

- Assessment, detection & response
- Networked and fault-tolerant control
- Disaggregation of Big Data in Societal CPS systems: Unsupervised and Supervisedd
- Fundamental Limits of Performance

Mechanism Design

- Incentive Theory for Cost Effective operations of Societal CPS Systems
- Utility Based Privacy Metrics
- Contract Mechanisms to Allow for Privacy Opt-In

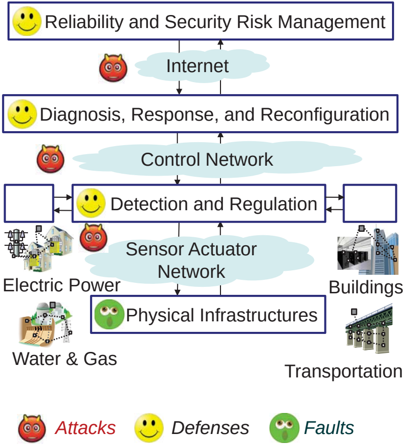

Thank you for your attention. Questions?

Shankar Sastry
sastry@coe.berkeley.edu