


Short and Long-Term Research Challenges for Protecting Critical Infrastructure Systems

UC Berkeley
Oct, 2013

Alvaro A. Cárdenas
Department of Computer Science
University of Texas at Dallas

From Sensor Nets to Cyber-Physical Systems (CPS)

- Control
- Computation
- Communication
- Interdisciplinary Research!
- Example: Smart Grid



Computer-Enabled Attacks & Threats

■ Attacks


■ Maroochy Shire 2000



■ HVAC 2012



■ Stuxnet 2010



■ Smart Meters 2012



FBI: Smart Meter Hacks Likely to Spread

39 views
A series of hacks perpetrated against so-called "smart meter" installations over the past several years may have cost a single U.S. electric utility hundreds of millions of dollars annually, the FBI said in a cyber intelligence bulletin obtained by

■ Threats

Obama Adm Demonstrates
In Feb. 2012 attack to power
Grid

DHS and INL study impact of
cyber-attacks on generator

US Video Shows Hacker Hit on Power Grid

US Video Shows Potential Destruction Caused by Hackers Seizing Control of Electrical Grid



In this image from video released by the Department of Homeland Security, smoke pours from an expensive electrical turbine during a March 4, 2007, demonstration by the Idaho National Laboratory, which was simulating a hacker attack against the U.S. electrical grid. (AP Photo/Dept. of Homeland Security)

By TED BRIDIS and EILEEN SULLIVAN
Associated Press Writers
WASHINGTON Sep 27, 2007 (AP)

A government video shows the potential destruction caused by hackers seizing control of a crucial part of the U.S. electrical grid: an industrial turbine spinning wildly out of control until it becomes a smoking hulk and power shuts down.

Font Size
A A A
E-mail
Print
Share

Your Opinion
Comment & Contribute
WHAT OTHERS ARE SAYING 44 Comments
I agree with the person(s) about h
old t...
locoyoc02 Sep-27
'seems for every smart program
there is...
OFEARTHLYGOOD Sep-27
scare tactics,Someone already
pointed out...
tmrabu Sep-27


COMMENT
Post Video
TALK
Connect with Newsmakers

Don't Forget Physical Attacks

BBC NEWS

You are in: World: **Americas**
 Tuesday, 18 January, 2000, 00:10 GMT

Colombia rebels blast power pylons



Blackout: Rebels attacked at least 17 power cables

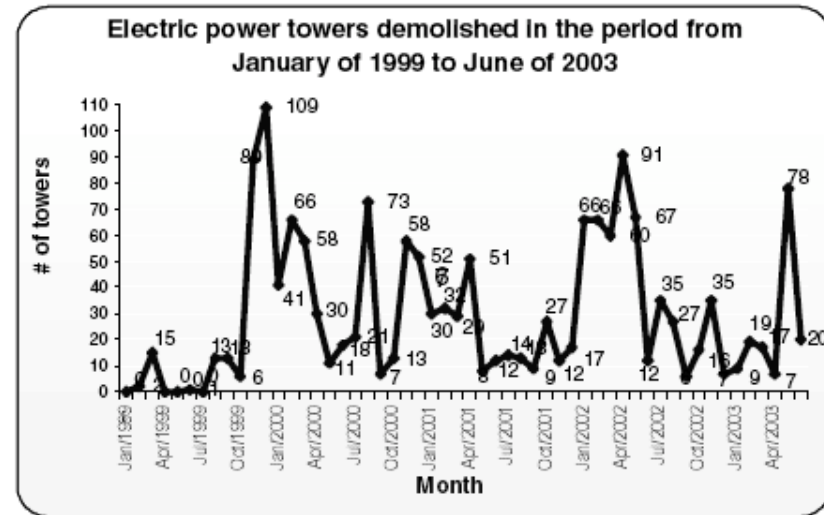
Power supplies to large areas of northern Colombia have been cut off after a series of rebel attacks on the country's electricity pylons.

Front Page World

Africa
 Americas
 Asia-Pacific
 Europe
 Middle East
 South Asia

From Our Own Correspondent

Letter From America
 UK
 UK Politics
 Business
 Sci/Tech
 Health
 Education
 Sport
 Entertainment
 Talking Point
 In Depth
 AudioVideo



Vulnerabilities can be Exploited

TECHNOLOGY | APRIL 8, 2009

Electricity Grid in U.S. Penetrated By Spies

Article

Video

Comments (146)



Email



Printer Friendly

Share:



Yahoo Buzz



Text Size



By SIOBHAN GORMAN




Associated Press

Robert Moran monitors an electric grid in Dallas. Such infrastructure grids across the country are vulnerable to cyberattacks.

WASHINGTON -- Cyberspies have penetrated the U.S. electrical grid and left behind software programs that could be used to disrupt the system, according to current and former national-security officials.

Cybermagedon?



NSA "Cyber Terrorism" Warning



Anonymous can take out the entire power grid.
Should be subjected to battlefield detention laws

Reality Check

- While the cyber-war rhetoric is a bit alarmist, there is a problem
- Cyber-Physical Systems are Vulnerable
 - By design
 - By lack of secure software development
 - As an attractive target in cyber-conflict
 - By lack of investment in security

Three Research Challenges to Improve CPS Security

- **Short Term**

- Incentives
- Software reliability
- Solve basic vulnerabilities

- Medium Term

- Leverage big data for situational awareness

- Long Term Research

- Attack-resilient estimation and control

Security is a Hard Business Case

- *“Making a strong business case for cybersecurity investment is complicated by the difficulty of quantifying risk in an environment of rapidly changing, unpredictable threats with consequences that are hard to demonstrate”*
 - DoE

As a Result Systems are Vulnerable with Basic Security Gaffes

- Unauthenticated remote connection to devices
- Unencrypted communications
- Hardcoded backdoor from manufacturer
- Hardcoded keys in devices
- Devices have several easily exploitable vulnerabilities
 - (e.g., Project Basecamp from DigitalBond)
 - Vendors not patching (mostly legacy devices)

Matter of Incentives

- Governments are responsible for Homeland Security, and critical infrastructure security
 - Utilities are not (outside their budget/scope?)
 - Problem:
 - Interdependencies (e.g., cascading failures)
 - It doesn't matter if one utility sets an example because this is a weakest security game
 - Nations have much more to lose from an attack than utilities
- What are the best ways to incentivize all players (vendors, asset owners, consumers, etc.) to implement best-security practices in the protection of Critical Infrastructures?

Legislation as an Incentive

- **Cybersecurity Act (S.3414)**
 - Currently trying to pass votes in US Senate (has failed twice)
 - Trimmed down regulation needs after opposition from republicans and some industry
- **SECURE IT Act (MIA?)**
 - Fun fact: bill uses the term “cyber-physical systems”
 - “collaborative research and development activities for cyber-physical systems with participants from universities federal laboratories and industry. Cyber-physical systems are systems found in infrastructure, healthcare, transportation, energy, and manufacturing where the systems’ s information technology and physical elements are tightly integrated.”



Incentives for Asset Owners: ROI Case Studies

- Game Theory in electricity theft
- Revenue = billed electricity + recovered theft

$$R(\theta, Q) = \sum T(q_B) + \sum \rho(e, q_U, f_1) F^r(q_U),$$

- Goal: find optimal investment in protection

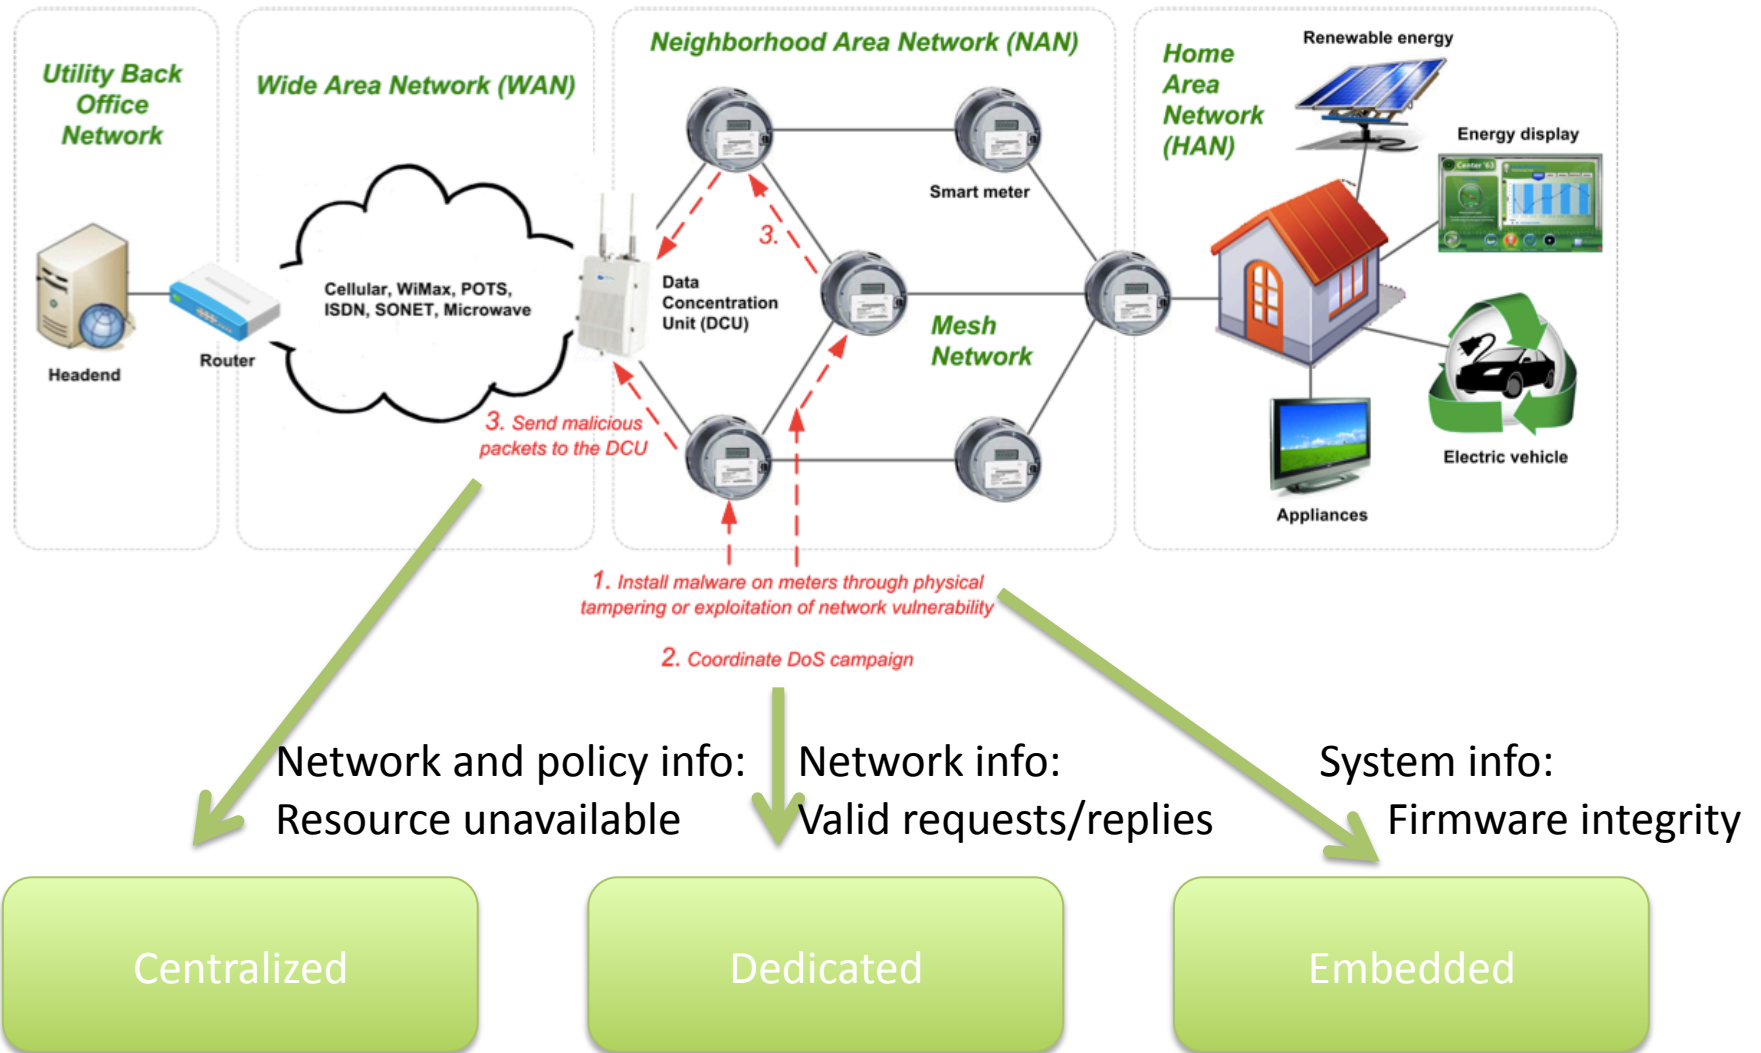
$$\max_{e \geq 0} R(\theta, Q) - C(Y) - \psi(e).$$

R = Revenue

C = Operational Cost

ψ = Security Investment

Alternatives for Investing in Intrusion Detection for AMI systems



[To appear in IEEE Transactions on Smart Grid 2014?]

Incentives for Vendors

- Asset owners need to request vendors secure coding practices, hardened systems, and quick response when new vulnerabilities and attack vectors are identified
- American Law Institute (ALI)
 - Principles of the Law of Software Contracts (2009)
 - Vendors liable for knowingly shipping buggy software
 - Implied warranty of no material hidden defects (non-disclaimable)
 - Software for CIP can be first use case

Future Work: Security Economics of CIP

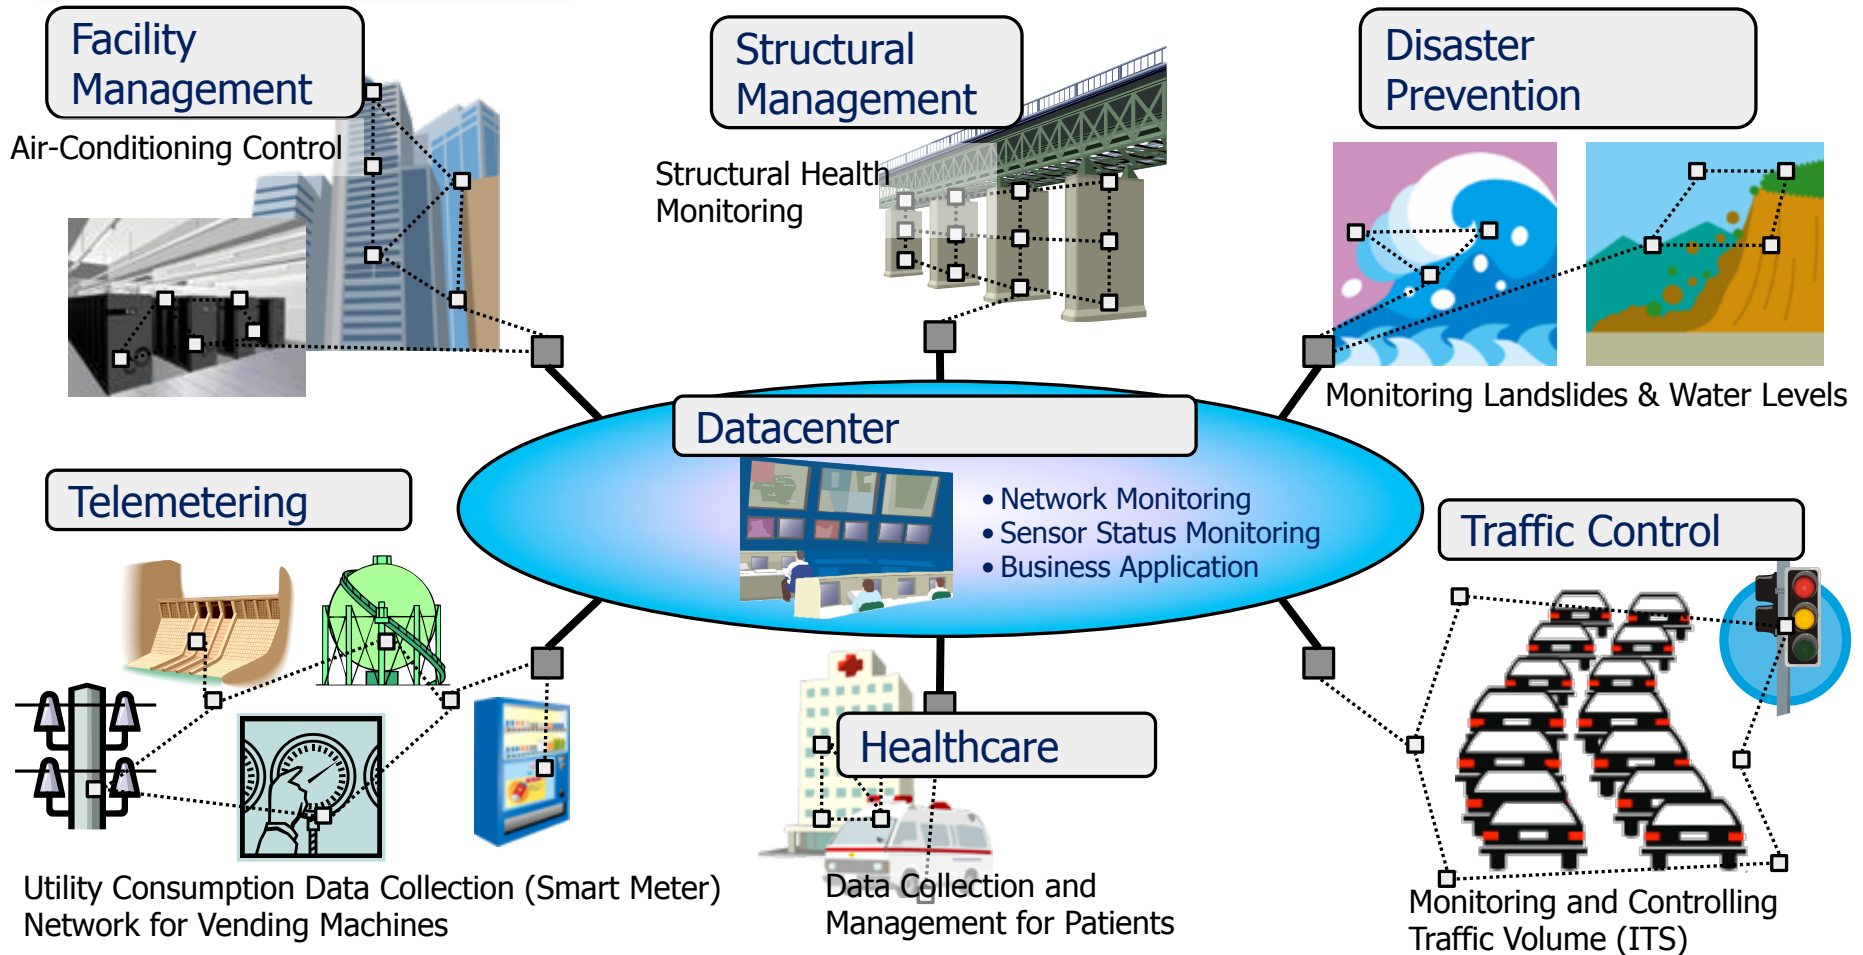
- Regulation
 - Federal (e.g., FERC) vs. State (e.g., PUC)
 - States need to take action first?
- Standards
- Case Law increasing responsibility and liability of vendors and asset owners
 - ALI: Principles of the Law of Software Contracts (2009)
- Procurement Language
- Insurance
- ROI
- Attacks

Three Research Challenges to Improve CPS Security

- Short Term
 - Incentives
 - Software reliability
 - Solve basic vulnerabilities
- **Medium Term**
 - Leverage sensor data for situational awareness
- Long Term Research
 - Attack-resilient estimation and control

Sensor Networks and Internet of Things (IoT)

Smart Infrastructures



Standards: Wireless HART (IEC), ISA SP 100.11a, IETF 6LoWPAN, ROLL, CoRE, Eman, LWIP, IRTF IoT, W3C EIX, IEEE 802.15.4 (g), 802.15.5, etc.

Business Case for “Data Analytics” is Easier than Security Business Cases

- Situational Awareness is part of the business case for modernizing our infrastructures
 - To understand the health of the system
 - Transmission grid, distribution grid, routing protocol in AMI, etc.
 - Wide Area Protection, Monitoring and Control
- Goal: leverage this data to improve cyber-security situational awareness
 - We get: Redundancy, Diversity
 - Data Analytics to identify suspicious behavior

Big Data Analytics in Smart Grid



EVENT

HOME APPLE BROADBAND CLEANTECH CLOUD COLLABORATION MEDIA



A note to our readers

We have updated our Privacy Policy and Terms of Service. Review continuing to use the site, you are agreeing to our updated Privacy Service.

This notice should appear only the first time you visit this site.

When big IT goes after big data on the smart grid

By Adam Lesser | Mar. 20, 2012, 10:49am PT | No Comments

Tweet 136 | Share 37 | Like 7 | +1 6

This article originally appeared on GigaOM Pro, our premium research service (subscription required).

With many utilities facing the task of storing petabytes of smart meter data for as long as seven years in order to satisfy regulatory requirements, the ability to house and leverage the massive load of



Big Data Offers Big Value for Utilities

03 Apr 2012 | United States

Share this


Smart meters produce data – it takes work to make the data ‘smart.’

What happened

Adam Lesser of Gigaom wrote about the difficulties faced by utilities when dealing with “big data” and the opportunities that this offers to IT companies.

According to Lesser, utilities face petabytes of data that needs to be stored for up to seven years to comply with regulation. Not only that, these utilities also need to “mine” this data and be able to pull out useful information, in a usable format, to allow them to save the time and money promised when deploying smart meters. In other words, make data ‘smart.’ This poses a “significant IT challenge,” one that is new to utilities.

In his report, “Smart Grid Billing Outlook 2012-2016,” author Danny Dicks says, “While smart meter deployments have been growing steadily over the last 3-4 years, utilities’ IT system priorities have been focused on preparing for how to deal with large volumes of smart meter data. This year we expect to see the emphasis change towards making use of that data – to develop innovative tariffs and new services ... All this will require changes to traditional billing systems and CISs.”



No end of possibilities for the fearless, forward thinking and imaginative.



Big Data Management for Energy and Smart Grid - Creating the Real-Time Utility Enterprise

Thursday, April 5, 2012 from 5:00 PM to 9:00 PM (PT)
Mountain View, United States



Ticket Information				
TYPE	REMAINING	END	QUANTITY	
General	Sold Out	Ended	Free	Sold Out

[Enter promotional code](#)

Share this! | Email | Share | Tweet | Like | 7 people like this. Be the first of your friends.

Event Details

Producer

Main Sponsor

When & Where


Building 152
Moffett Field, Hwy 101
Mountain View, 94035

Thursday, April 5, 2012 from 5:00 PM to 9:00 PM (PT)

[Add to my calendar](#)

Advanced Metering Infrastructure (AMI)

- Replacing old mechanical electricity meters with new digital meters
- Enables **frequent**, periodic 2-way communication between utilities and homes





[Iwao, et.al. IEEE SmartGridComm, 2010]

[Céspedes, Cárdenas, IEEE ISGT 2012]



[Herberg, Cárdenas, et.al. IETF-draft-dff-cardenas 2012]

Electricity Consumption

Weekly




Daily



Electricity Theft

Figure 4.

Revenue loss due to non-technical loss of electricity



Source: Investment and Financial Flows To Address Climate Change. United Nations
 Source: IEA, 2007; ENERDATA, 2007; Smith, 2004.

Attacks will happen:
 Devices are deployed
 for 20~30 years




ABOUT THIS BLOG

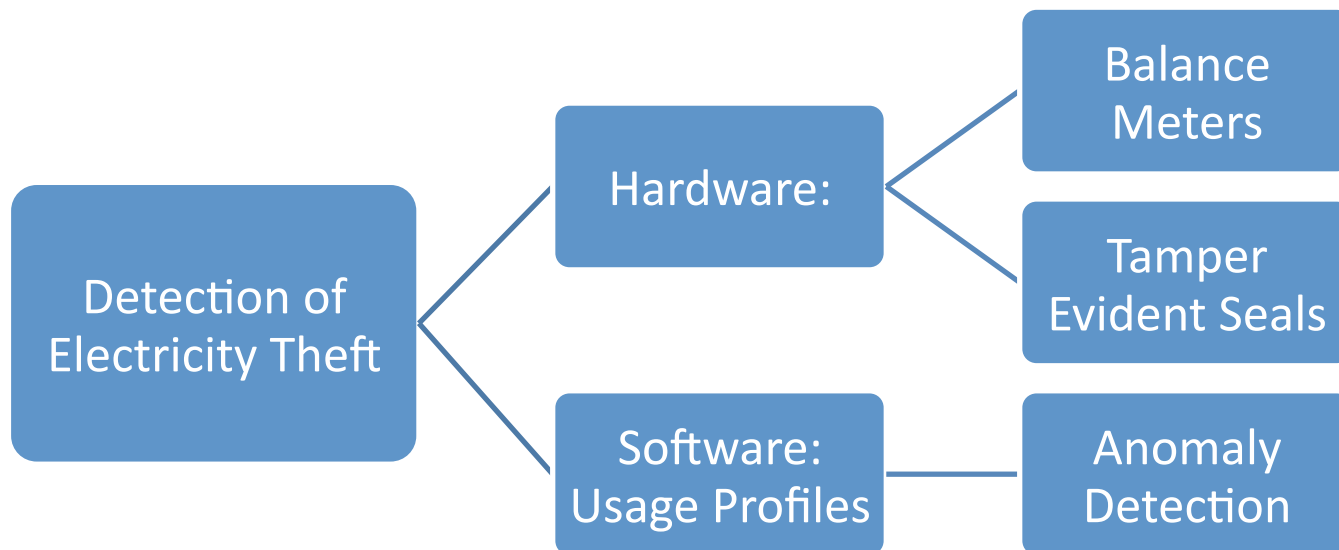
FBI: Smart Meter Hacks Likely to Spread

39 tweets retweet

A series of hacks perpetrated against so-called "smart meter" installations over the past several years may have cost a single U.S. electric utility hundreds of millions of dollars annually, the **FBI** said in a cyber intelligence bulletin obtained by





Anomaly Detection of AMI Data Can Complement other Detection Mechanisms



A tangle of wire atop this electricity pole in New Delhi, India in 2002 was testament to the capital city's power theft problems. Since then in North Delhi, automation has helped slash electrical losses.

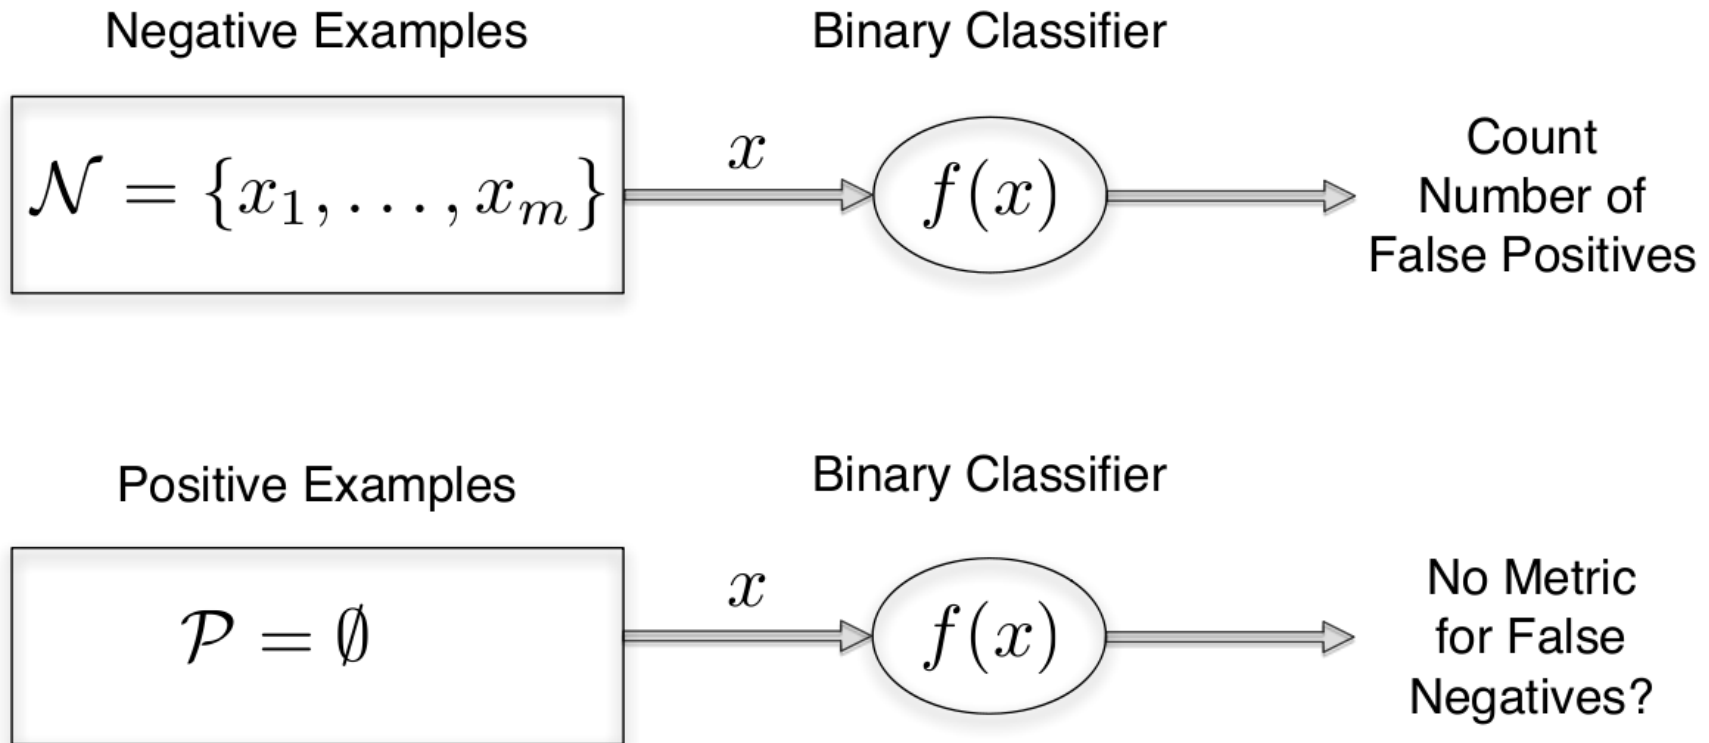
Evaluation

- Most Machine Learning Algorithms Assume a pool of Negative Examples and a Pool of Positive examples to evaluate the tradeoff between false alarms vs. detection rate:



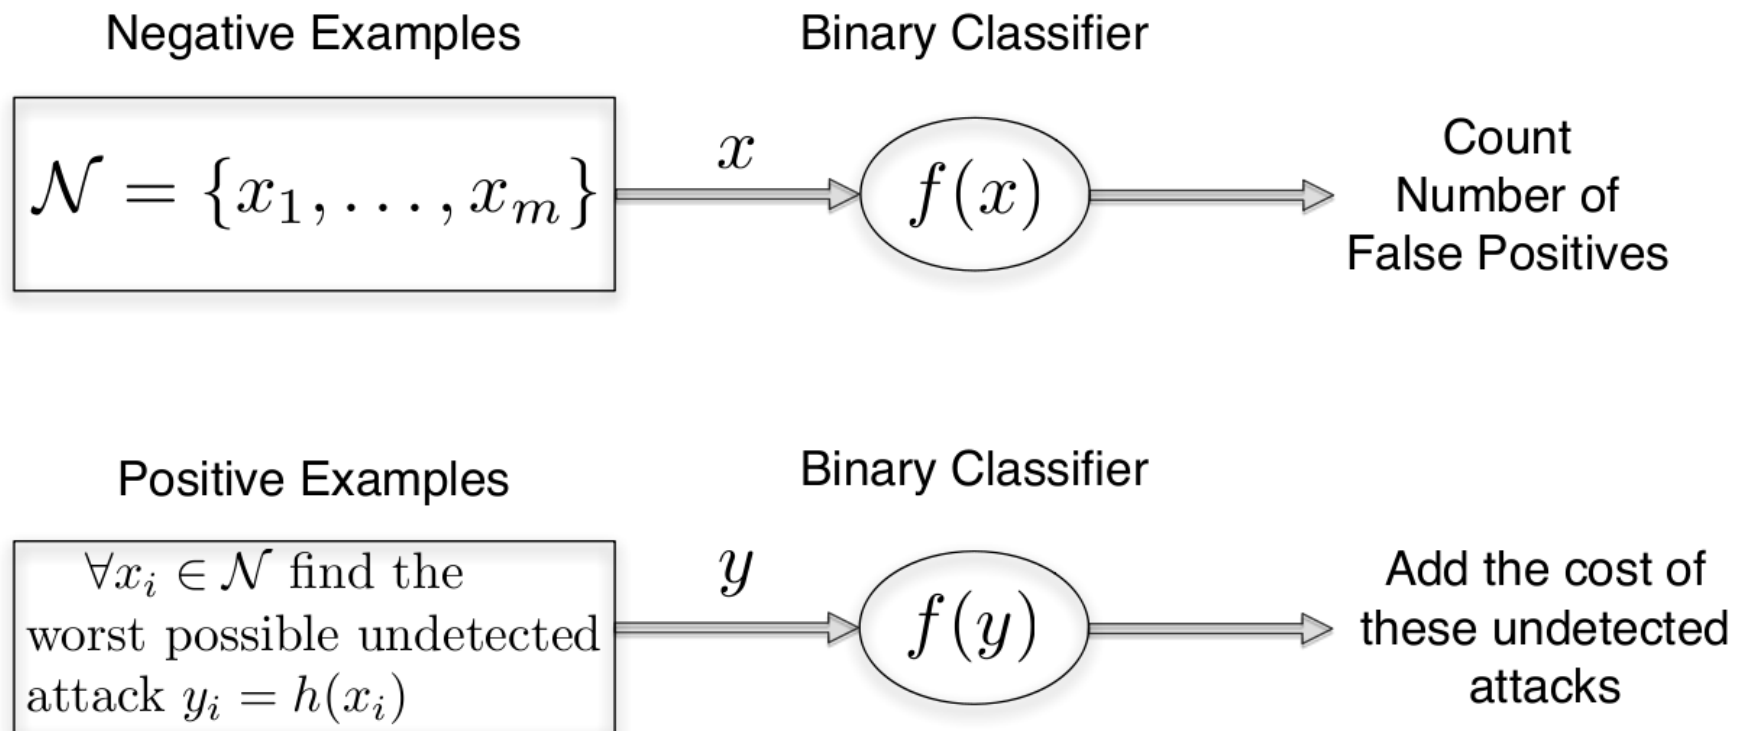
Problem: We Do Not Have Positive Examples

- Because meters were just deployed, we do not have examples of “attacks”



Our Proposal:

- Find the worst possible undetected attack for each classifier, and then find the cost (kWh Lost) of these worst-case undetected attacks




[Mashima, Cárdenas, Evaluating Electricity Theft Detectors. RAID, 2012]


Adversary Model



$f(t)$
Real Consumption
 Y_1, \dots, Y_n



$a(t)$
Fake Meter Readings
 $\hat{Y}_1, \dots, \hat{Y}_n$

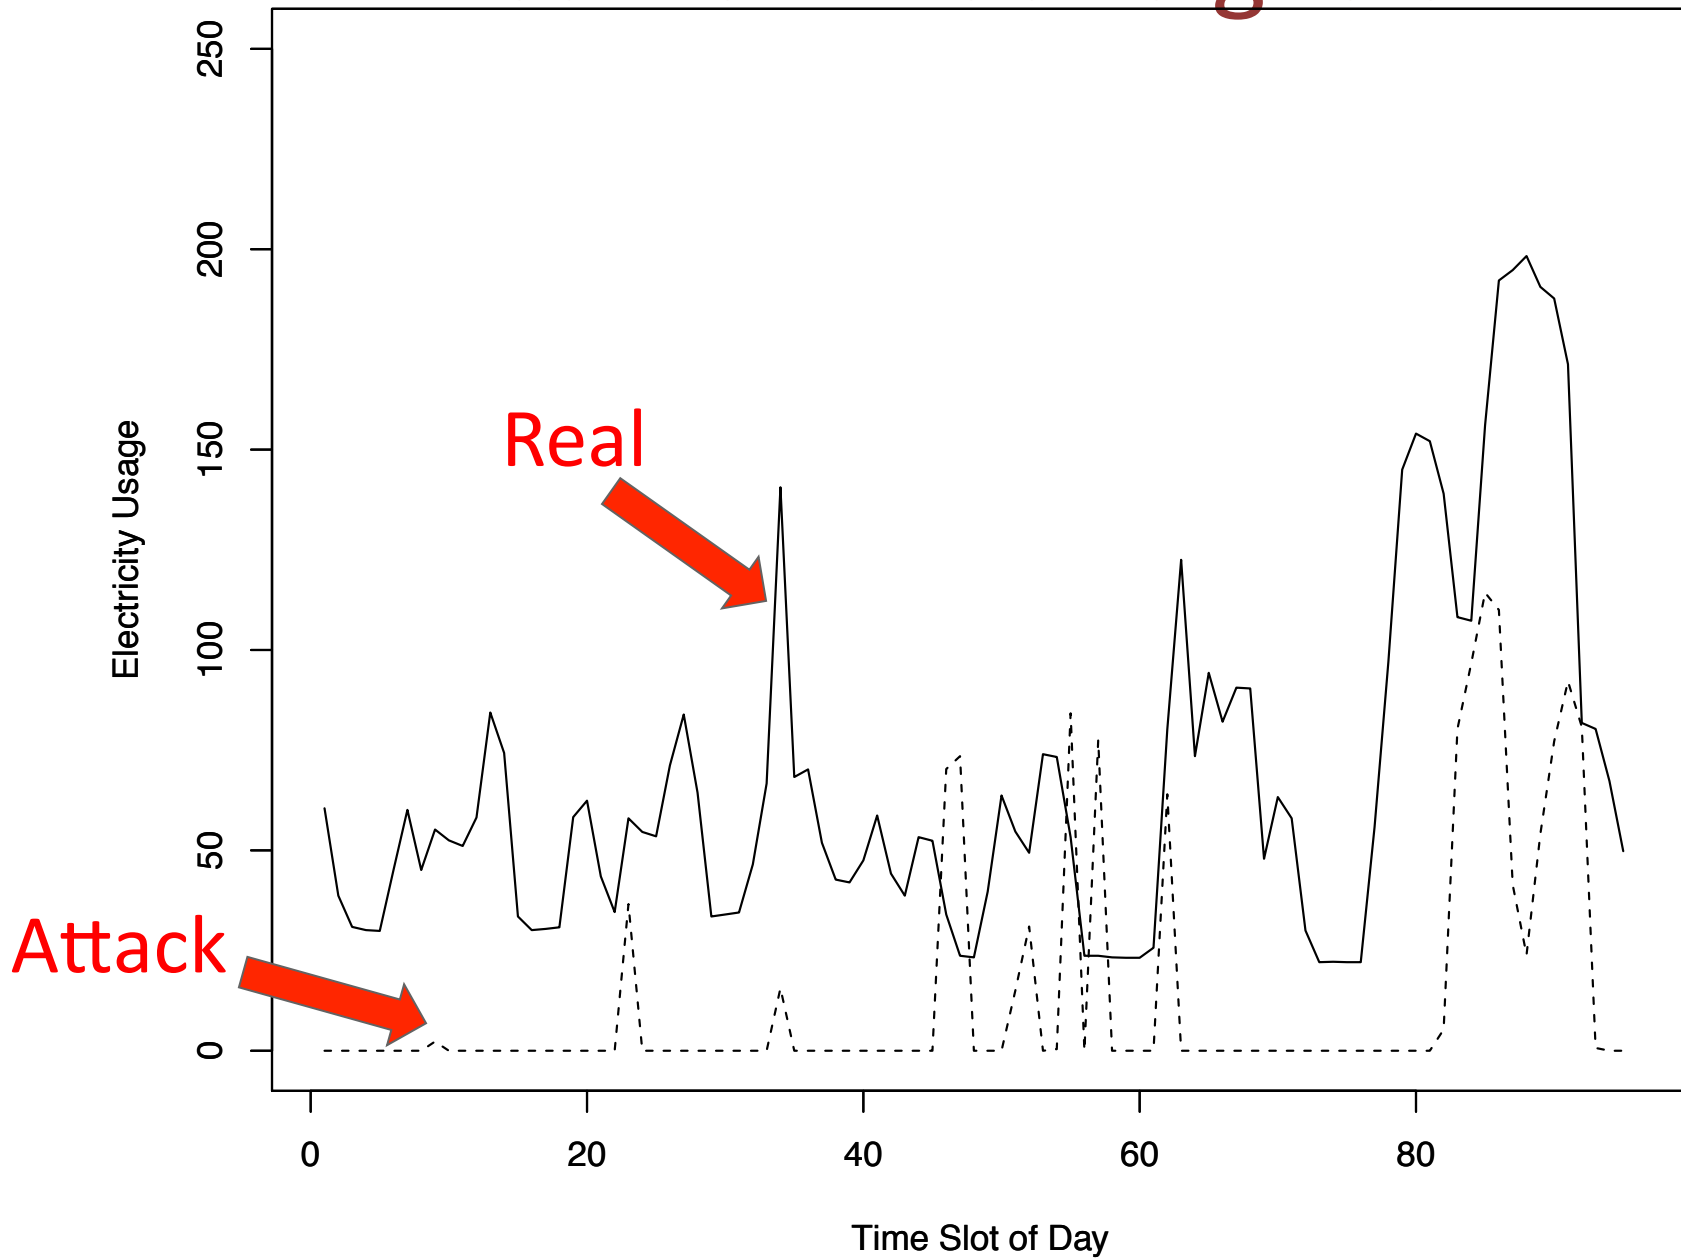


1st Goal of attacker: Minimize Energy Bill: $\min_{\hat{Y}_1, \dots, \hat{Y}_n} \sum_{i=1}^n \hat{Y}_i$

2nd Goal of Attacker: Minimization subject to not being detected by classifier "C":

$$C(\hat{Y}_1, \dots, \hat{Y}_n) = \text{normal}$$


Real vs. Attack Signals



New Tradeoff Curve: No Detection Rates

Y-axis: Cost of Undetected Attacks (can be extended to other fields)

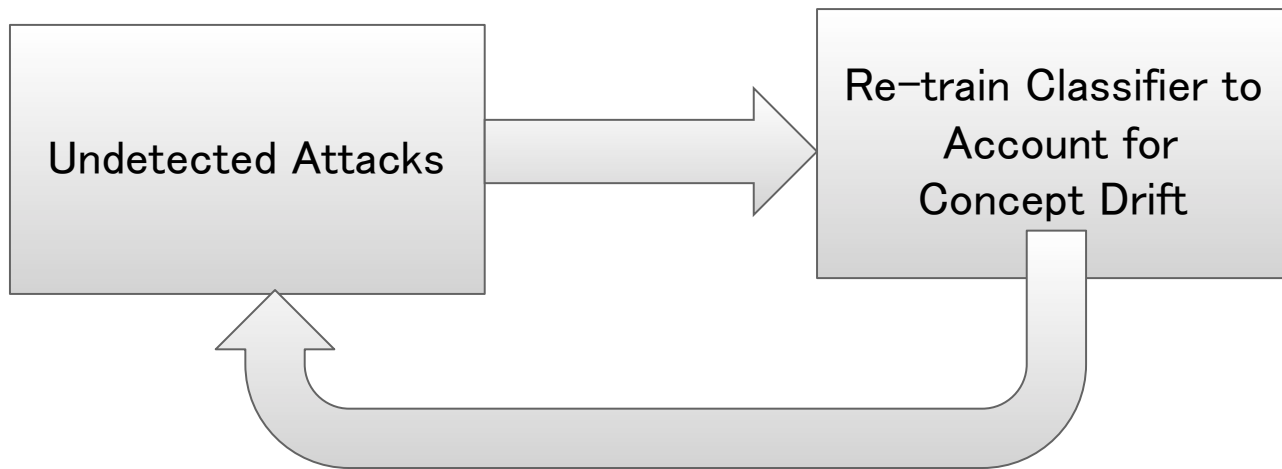
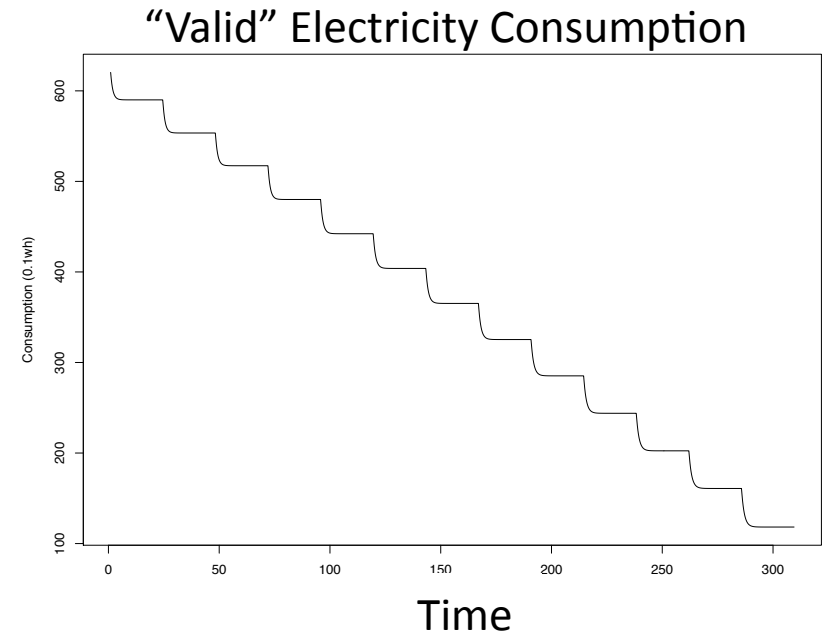
X-axis: False Positive Rate



Asymptotic Effects of Poisoning Attacks

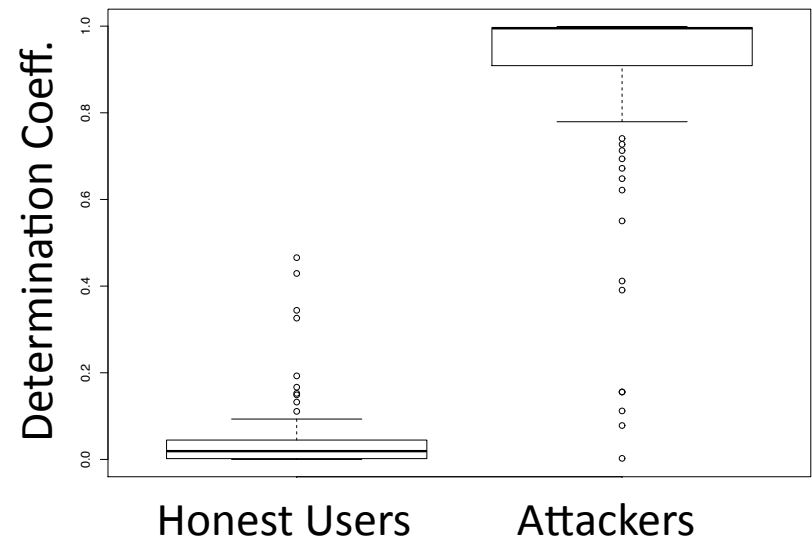
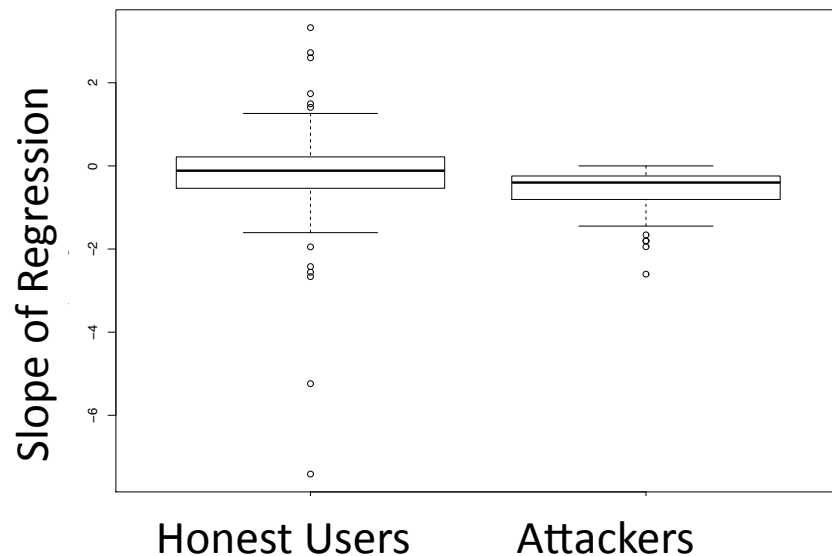
■ Concept Drift

- Electricity consumption is a non-stationary distribution
- We have to “retrain” models
- Attacker can use undetected attacks to poison training data





Detecting Poisoning Attacks



- Identify concept drift trends that could benefit an attacker
 - i.e., Lower electricity consumption over time.
- Countermeasure: linear regression of trend
 - Slope of regression was not good discriminant
 - Determination coefficients worked!



Ongoing Work: Detecting Other Anomalies



Ongoing Work: Cross-Correlation, Weather



[Mashima, Cárdenas, Evaluating Electricity Theft Detectors. RAID, 2012]

Three Research Challenges to Improve CPS Security

- Short Term
 - Incentives
 - Software reliability
 - Solve basic vulnerabilities
- Medium Term
 - Leverage Big Data for Situational Awareness
- **Long Term Research**
 - Attack-Resilient estimation and control

What is New and Fundamentally Different?

- So security is important; but
 - are there new research problems?
 - or can CPS security be solved with:
 - IT security best practices?
 - Control systems best practices?

Previous Work in Security: What can Help in Securing CPS?

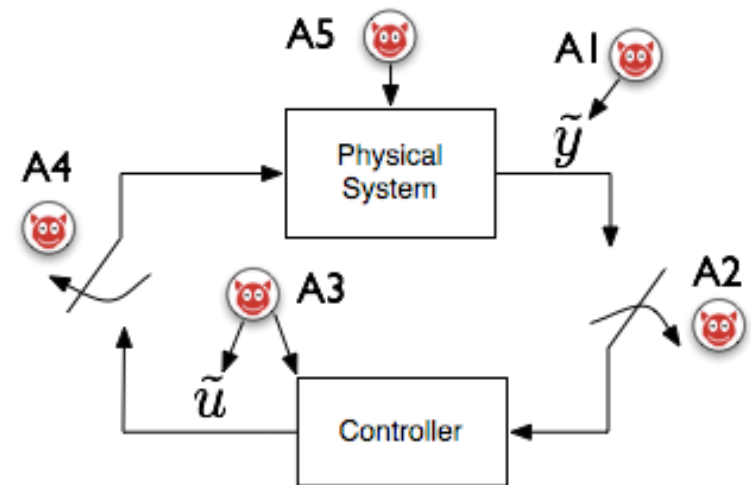
- **Prevention**
 - Authentication, Access Control, Message Integrity, Software Security, Sensor Networks, Trusted Computing, White Listing
- **Detection**
 - Intrusion detection, anomaly detection, forensics
- **Resiliency**
 - Separation of duty, least privilege principle
- Incentives for vendors and asset owners to implement security best practices

Previous Work in Security: What is Missing for Secure CPS?

- APT attacks will succeed, even with security best practices
- Can we improve security by modeling cyber-interaction with the physical world?
 - How can the attacker manipulate the physical world? (better threat analysis)
 - Design attack-resilient control and estimation algorithms

■ Attacks to Regulatory Control

- A1 and A3 are deception attacks: the integrity of the signal is compromised
- A2 and A4 are DoS attacks
- A5 is a physical attack to the plant

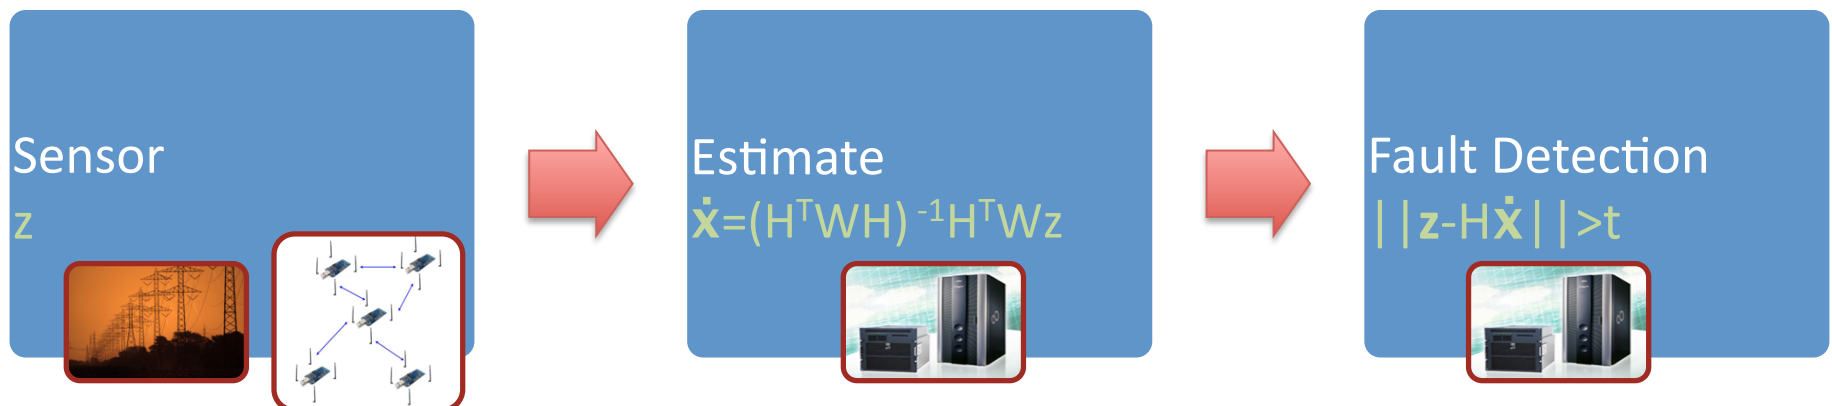


Previous Work in Control: What Can Help in Securing CPS?

- **Networked control**
 - Deals with control over lossy networks
 - Packet drops, network failures, etc. (similar to DoS)
- **Robust control**
 - Deals with uncertainties in the model and noise
 - Control algorithms resilient to worst-case disturbances
- **Fault-tolerant control**
 - Detects and isolates faulty components
- **Safety systems**
 - Takes over control when system is in danger

Previous Work in Control: What is Missing for Securing CPS?

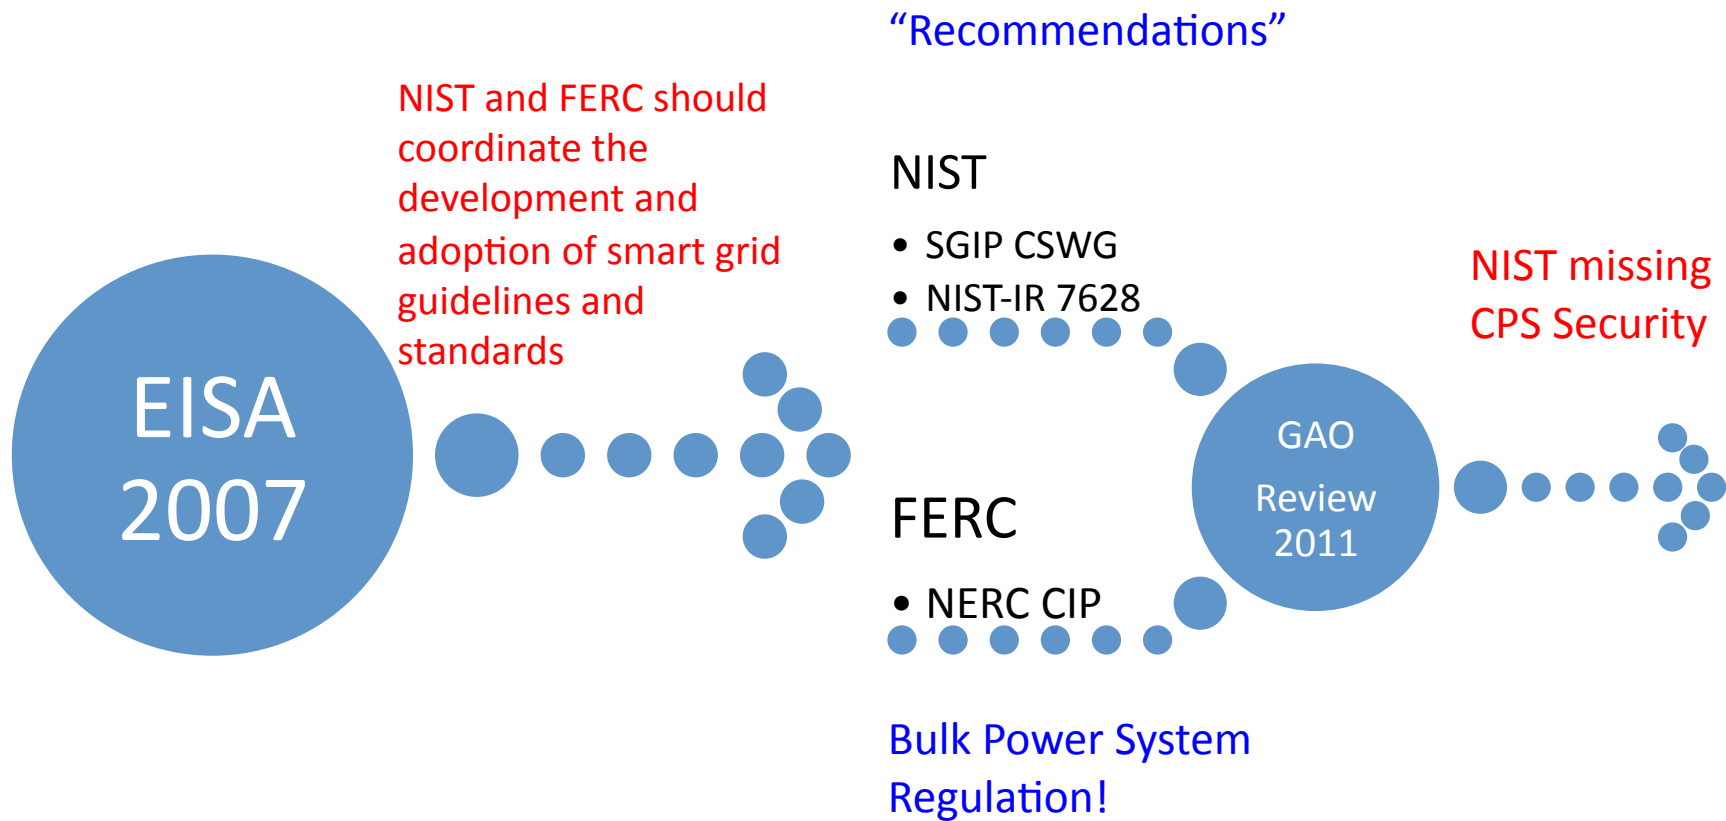
- Attacks are different than failures!
 - Attacks will evade fault-to
 - Non-correlated, non-independent, etc.
- **Example:**
 - **Fault-Detection Algorithms do not Work Against Attackers**
 - Liu, Ning, Reiter. CCS 09
 - Proof of concept attacks z such that $||z-H\hat{X}|| < t$



Control Theory + Computer Security Analysis = Resilient CPS

Improving Resiliency Against APT!

GAO Agrees: We Need new Research for CPS Security



New CPS Research Directions

■ Threat assessment:

- How to model attacker and his “control” strategy
- Consequences to the physical system

■ Attack-**resilient** control algorithms

- CPS systems that degrade gracefully under attacks

■ Attack-**detection** by using models of the physical system

- Study stealthy attacks (undetected attacks)

■ Privacy

- Privacy-aware CPS algorithms

Papers articulating these ideas:

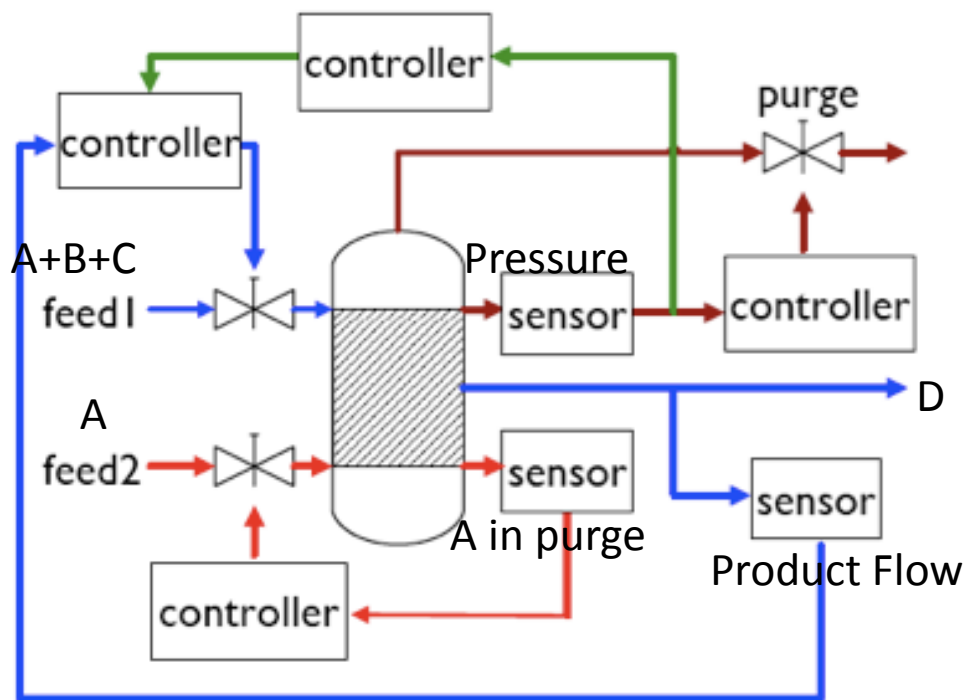
[Cárdenas, Amin, Sastry, HotSec 2008]

[Cárdenas, Amin, Sastry, ICDCS CPS Workshop 2008]

Requirements for Secure Control

■ Traditional Security Requirements: CIA (Confidentiality, Integrity, Availability)

■ What are the requirements of secure control?



• Safety Constraint:

– Pressure < 3000kPa

• Operational Goal:

– Minimize Cost:

• Proportional to the quantity of A and C in purge,

• Inversely proportional to the quantity of the final product D

$$\text{Cost} = \frac{F_3}{F_4} (2.206y_{A3} + 6.177y_{C3})$$


Risk Assessment

- If attacker compromises one (or more) sensor or actuators,
 - What attack strategy (false signals) can attacker use to disrupt our secure control requirements:
 - Violate Safety?
 - Maximize Operational Cost?
- At the end of this analysis we can identify high-priority sensor and actuators (the ones that require more security/trust)


[Journal of Critical Infrastructure Protection 2009]

Not all Compromises affect Safety


Production




Pressure






A in Purge




Feed of A



Resilient by Redundancy:




Purge Valve




Safety can be Compromised at Different Time Scales


Prioritize protection of control signal for A+B+C feed




It takes 20 hours to violate Safety by compromising the pressure sensor signal (prevention vs. detection&response)




DoS Attacks: No Impact when the System is at Steady State




However: A previous “innocuous” integrity attack becomes significant with the help of DoS attacks



Attacks to the Operational Cost Involve Devices that do not Matter in Safety

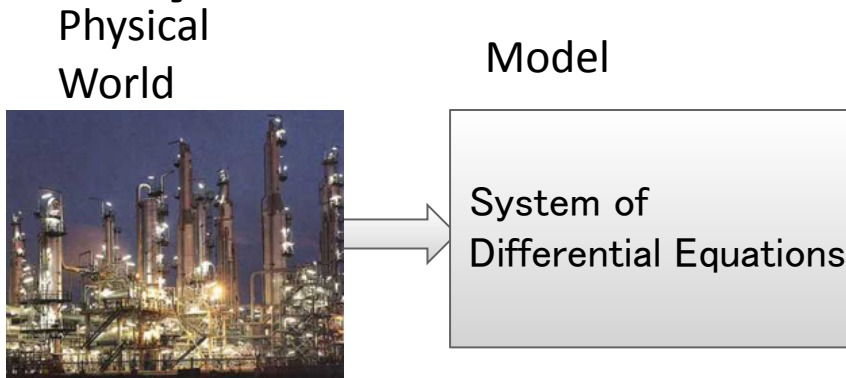


Attack increases safety but lowers profits



New Attack-Detection Mechanisms by Incorporating “Physical Constraints” of the System


- **1st Step: Model the Physical World**



- **3rd Step: Response to Attacks**


- **No attack**

- Use real plant signal



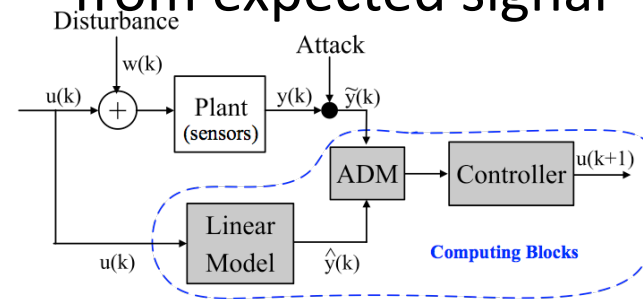
- **IDS detects attack**

- Switch to linear model
 - Detection time
 - False alarm



- **2nd Step: Detect Attacks**

- Compare received signal from expected signal



- **4th Step: Security Analysis**

- Missed Detections

- Study stealthy attacks

- False Positives

- Ensure safety of automated response

[Cárdenas, et.al. AsiaCCS, 2011]


Attacker Strategy: Stealthy Attacks

- Attacker
 - Knows our detection model and its parameters
 - Wants to be undetected for n time steps
 - Wants to maximize the pressure in the tank

- Surge attack
$$\tilde{y}_K = \begin{cases} y^{min} & \text{if } S_{k+1} \leq \tau \\ \hat{y}_K - |\tau + b - S_k| & \text{if } S_{k+1} > \tau \end{cases}$$


- Bias attack
$$\tilde{y}_k = \hat{y}_k - (\tau/n + b)$$

- Geometric attack
$$\tilde{y}_k = \hat{y}_k - \beta\alpha^{n-k}$$



Impact of Undetected Attacks

- Even geometric attacks cannot drive the system to an unsafe state
- If an attacker wants to remain undetected, she cannot damage the system



Control Resilient to DoS Attacks

For constrained linear systems

$$x_{k+1} = Ax_k + Bu_k^a + w_k, \quad k = 1, \dots, N-1$$

$$x_k^a = \gamma_k x_k, u_k^a = \nu_k u_k, \quad (\gamma_k, \nu_k) \in \{0, 1\}^2$$

find **causal feedback policies** $u_k = \mu_k(x_0^a, \dots, x_k^a)$, that
minimize $J(x_0, \mathbf{u}, \mathbf{w}) = \sum_{k=1}^N x_k^\top Q^{xx} x_k + \sum_{k=1}^{N-1} \nu_k u_k^\top Q^{uu} u_k$,
subject to **power constraints**

$$\begin{pmatrix} x_k^a \\ u_k^a \end{pmatrix}^\top \begin{pmatrix} H_i^{xx} & 0 \\ 0 & H_i^{uu} \end{pmatrix} \begin{pmatrix} x_k^a \\ u_k^a \end{pmatrix} \leq \beta_i, \quad i = 1, \dots, L_1,$$

and **safety constraints**

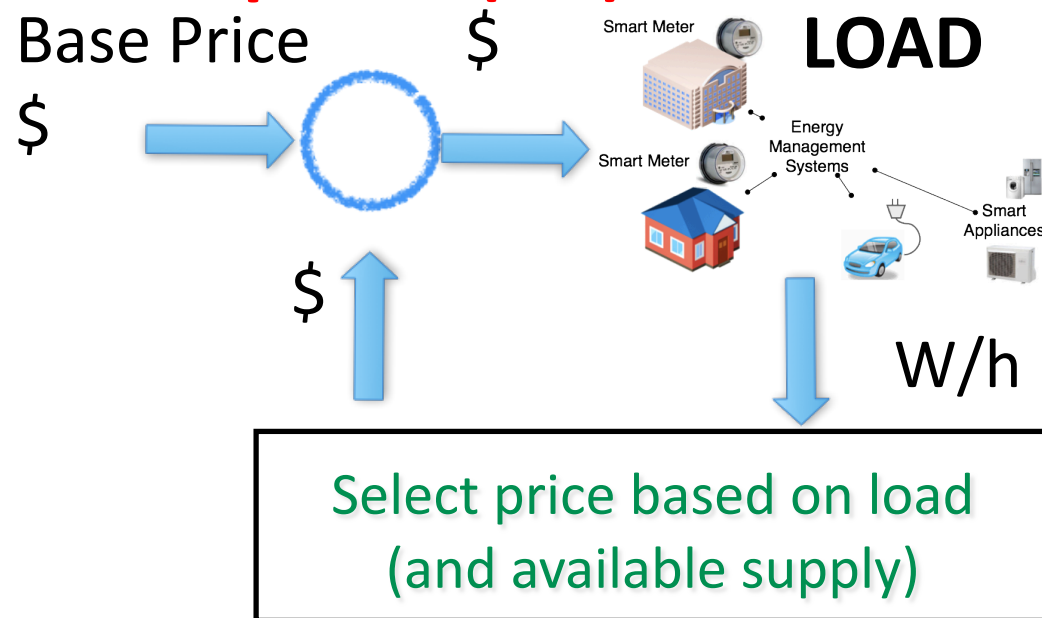
$$\begin{pmatrix} x_k^a \\ u_k^a \end{pmatrix} \in \mathcal{T}_j, \quad j = 1, \dots, L_2,$$

for all **disturbances** $\mathbf{w} \in \mathbf{W}_\alpha$ OR $\mathbf{w} \sim \mathcal{N}(0, W)$ and a given set of
 $(\gamma_0^{N-1}, \nu_0^{N-1}) \in \mathcal{A}_{pq}$ **attack signatures**.

[Amin, Cárdenas, Sastry. HSCC / CPSWeek 2009]

Privacy-Preserving Control

- Data Minimization Principle
 - How much data do we really need to collect for accurate estimation/control?
 - Quantity: sampling
 - Quality: quantization
- **Demand Response (DR)**



Conclusions

- First
 - Address basic security problems
 - No need for “research” in CS aspects of security but on “security economics”
- Second
 - Improve situational awareness
- Third
 - Design for resiliency
 - Leverage control systems expertise in security analysis