FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Energy Analytics

Henrik Ohlsson and Lillian J. Ratliff
joint with Roy Dong and S. Shankar Sastry
UC Berkeley

Berkeley UNIVERSITY OF CALIFORNIA

Massachusetts Institute of Technology

NSF

MICHIGAN

VANDERBILT UNIVERSITY

Background

- Most of the components in the electrical grid are > 50 years old.
- Recent modernization of the grid (installment of AMIs) comes with advantages to control and monitoring. However, in the wrong hands, the data might pose a privacy treat.
- The electrical grid is today a cyber-physical system.
- Utilities are currently utilizing smart meters for meter-to-cash. The potential of smart meters go far beyond this basic usage and the utilities are looking for a justification for their investments. The market for energy *analytics** in the smart grid is estimated to be worth \$9.7 billion by 2020.

*Analytics is the discovery and communication of meaningful patterns in data.

A Flora of Interesting Problems

Interesting and challenging problems for academia: complex systems, enormous amounts of data, multi-disciplinary, practical importance,... Also recent interest in industry (C3 Energy, Opower, GE, IBM, Bidgely,...)

- Revenue protection
- Customer segmentation
- Risk management
- Disaggregation
- Load forecasting
- Voltage optimization
- Outage management
- Energy efficiency/incentive design
- Security and privacy
- ...

Security

Revenue Attacks in the Smart Grid

- Non–Technical losses are caused by actions external to the power system such as theft, non–payment by consumers, or errors in accounting.
- The World Bank recently reported that in some countries (such as India) as little as 50 percent of the generated electricity is paid for.

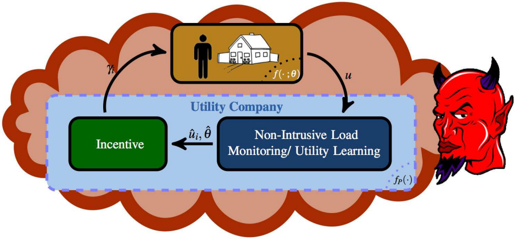

Revenue Protection

Revenue protection is the process of identifying non-technical losses in the grid and correcting for those losses.

- **Goal:** Implement a **revenue protection** scheme so that the smart grid is **resilient** to revenue attacks by **adversarial** agents.
- First step is building an efficient theft/anomaly detection algorithm which takes in historical consumption data (monthly/quarter hourly), event flags, work orders and meta information and returns a label (theft/normal). The algorithm is trained using historical investigation results.

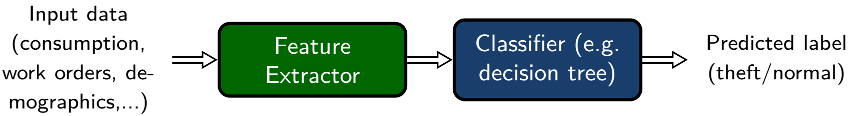

Input Data

- AMI data including time series consumption data (hourly, quarter hourly, etc.) and internal flags ('malfunction', 'tamper', etc.)
- Weather data
- Customer information (i.e. is the customer a producer?, billing information, rate category, region, etc.)
- Work order information such as dates, times, and descriptions of maintenance.

Features

- Max consumption drop over various intervals of time (month, day, year).
- Total variation over the past month/week/day.
- Production during the night.
- Region in which the customer resides.
- Work order based features such as scheduled visit followed by a large drop in consumption.
- Non-zero consumption after closing account.
- Missing payments.
- ...

Classification

- Classifier: Logistic regression, Adaboost, Random Forests,...

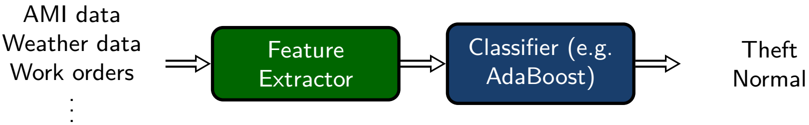

Revenue Protection Results

- In conjunction with C3 Energy and several large utilities, 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 millions of customers.
- The data included time-series consumption data and meter events from AMIs, weather data, customer demographics, and work orders.
- We identified $\sim$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.

On Going Research in Revenue Protection and Beyond

Technology transfer to industry: Our research led to the development and implementation of revenue protection machine learning algorithms for theft and anomaly detection that have been deployed at several utility companies.

Advancement of research agenda through industrial interaction: Through working with real data from utility companies, we have discovered new and fundamental research problems that are practically relevant.

Revenue Protection in the Electrical Grid, Henrik Ohlsson, Lillian Ratliff, et al., 2014, in preparation.

Energy Efficiency

Improved Feedback

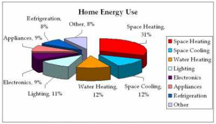

Consumers

consumer feedback

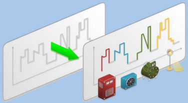

Disaggregation

Energy Disaggregation

- Yearly consumption data, weather & profile information
- Infer most likely profile if information is missing
- Collaboration with C3 Energy and several utility companies

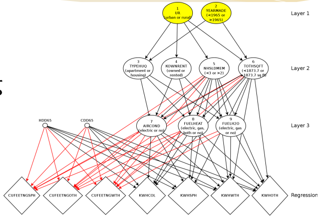

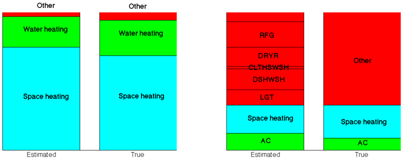

Ratliff, Dong, Ohlsson, Sastry, Incentive Design and Utility Learning via Energy Disaggregation. IFAC 2014.
Dong, Ratliff, *et al.*, Fundamental Limits of Non-Intrusive Load Monitoring. HICONS, 2014.
Dong, Ratliff, Ohlsson, Sastry, Energy Disaggregation via Adaptive Filtering. Allerton, 2013.
Dong, Ratliff, Ohlsson, Sastry, A Dynamical Systems Approach to Energy Disaggregation. CDC, 2013.

Incentive Design in Energy Systems

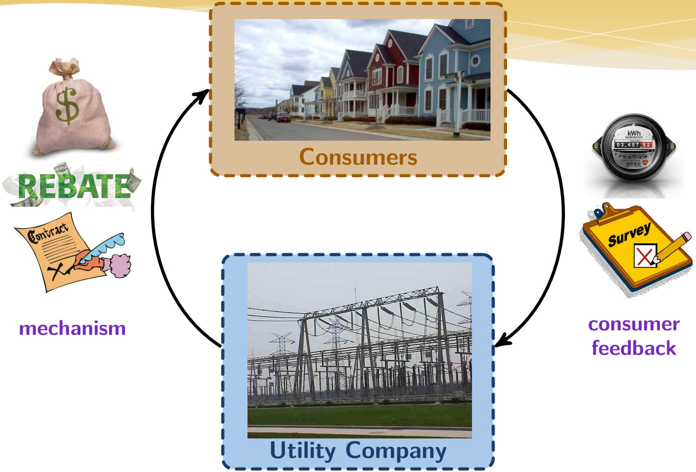

Motivations for Incentive Design in Energy Systems

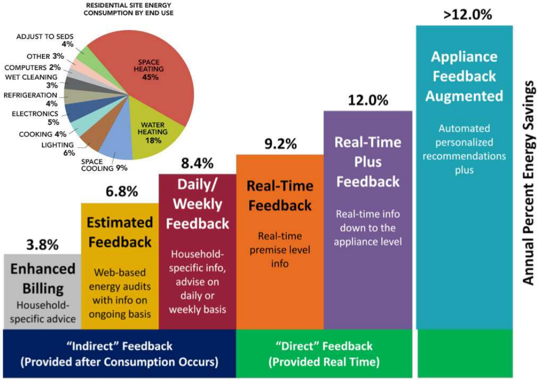

- Studies have shown that providing device-level feedback on power consumption patterns to energy users can modify behavior and improve energy efficiency.
- Provide **incentives** in the form of **rebates and monetary rewards** focusing on devices that fall into largest consumption categories in order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

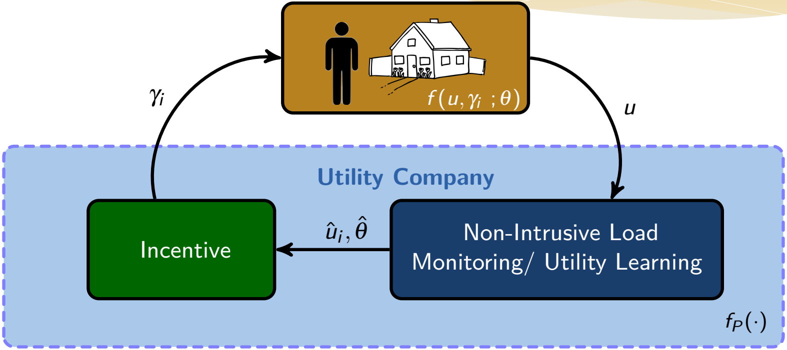

Ratliff, Dong, Ohlsson, Sastry, Incentive Design and Utility Learning via Energy Disaggregation. IFAC 2014.
Dong, Ratliff, *et al.*, Fundamental Limits of Non-Intrusive Load Monitoring. HICONS, 2014.
Dong, Ratliff, Ohlsson, Sastry, Energy Disaggregation via Adaptive Filtering. Allerton, 2013.
Dong, Ratliff, Ohlsson, Sastry, A Dynamical Systems Approach to Energy Disaggregation. CDC, 2013.

Knowledge Transfer – Consumer Feedback for Change

- Using this theoretical framework, in conjunction with C3 Energy and a utility company, we will monitor the energy consumption of a small group of homes and provide feedback to these users about their energy consumption patters.
- The utility company is in a regulated market in which the utility commission has incentivized them to reduce the overall consumption of its consumer base.
- Through this pilot program, the utility company aims to learn which types of feedback mechanisms will encourage lasting behavior changes in its consumers.

Incentives Introduce New Vulnerabilities

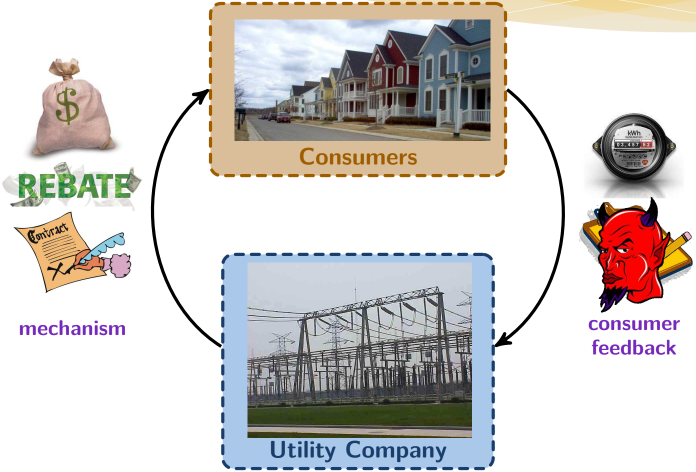

Incentives Introduce New Vulnerabilities

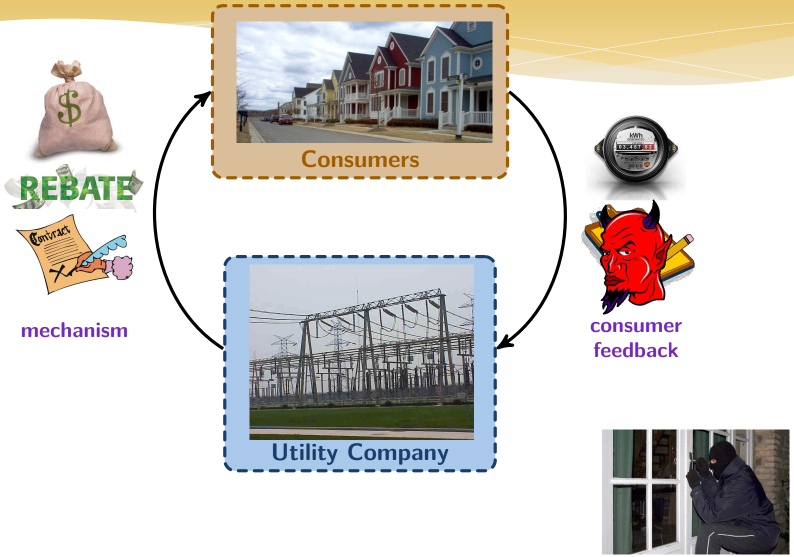

Incentives Introduce New Vulnerabilities

mechanism

Consumers

consumer feedback

Utility Company


The Post and Courier
Search Articles
Home News Sports Business Multimedia Entertainment Features Special Sections Blogs
Home > News > Local/State News
Marijuana bust shines light on utilities
Glenn Smith
Posted: Sunday, January 29, 2012 12:01 a.m., Updated: Friday, March 23, 2012 6:08 p.m.


privacy
conflict
public
reinvent
advocacy
golfing
social
journalism
annotated
minimization
development
online
communications
education
security
ubiquitous
international
government
rights
digital
realtime
narration
consumer
reputation
reform
management
strategy
internet
process
media
act
workplace
biometrics
video
law
Web
nerd
technology
policy
cdt
cyberethics
soccer
big
voter

Privacy Contracts: Two-Type Model

- It is not the signal itself that is private but what can be inferred that is private.
- There is a tradeoff between efficiency of grid operations and privacy-aware sampling polices.
- We have a bound on a successful privacy breach by an adversary.

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.

- There are two privacy settings offered: x_L, x_H such that $x_L \leqslant x_H$, $x_L, x_H \in \mathbb{R}$.
- We consider two types of consumers: $\theta \in \{\theta_L, \theta_H\}$ where θ represents how much the consumer values high-privacy over low-privacy.
- The consumer type is unknown to the utility company.

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Individual Rationality and Incentive Compatibility

- The utility company is to design a pair of contracts: $\{(t_L, x_L), (t_H, x_H)\}$.
- The consumer's utility is equal to zero if he does not select a privacy setting (opt-out), and it is

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

  if he selects the contract (t, x).
- Assumption: U is strictly increasing in (x, θ).
- **Incentive-compatible**: all of the participants fare best 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 cost:
$$v(x,t) = t - g(x)$$
where $g(x)$ is the unit cost resulting from the privacy setting x.
- $g(x)$ 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 $p = \Pr(\theta = \theta_H) = 1 - \Pr(\theta = \theta_L) \in (0,1)$ (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 (i.e. the marginal gain from raising the value of the privacy setting is greater for type θ_H).
- Then, 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: An Example

DLC example: as you decrease the sampling rate the performance degrades.

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

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

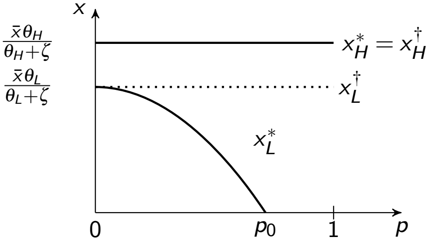

Ratliff, Dong, Ohlsson, Cárdenas, Sastry. Privacy and Customer Segmentation in the Smart Grid. Submitted to CDC, 2014.

Privacy in the Smart Grid Summary

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.

- Using the contract theory framework, we can design **insurance contracts** to be offered to either the consumer or the utility company.
- **Revenue Protection**: Utility company has a right to find out if a consumer is hiding behind privacy to steal electricity suggesting we should be studying these problems in parallel to develop a more **resilient smart grid**.
- **Multidimensional Screening**: We are currently exploring hidden convexities in the multidimensional screening problem in which consumers have higher dimensional types which live on a continuum.
- **Quantifying the Utility-Privacy Tradeoff:** We can quantify the tradeoff between privacy aware AMI sampling policies and smart grid operations.

Ratliff, Dong, Ohlsson, Cárdenas, Sastry. Privacy and Customer Segmentation in the Smart Grid. Submitted to CDC, 2014.
Dong, Cárdenas, Ratliff, Ohlsson, Sastry. Quantifying the Utility-Privacy Tradeoff in the Smart Grid. Submitted to IEEE Transactions on Smart Grid, 2014.

Customer Segmentation

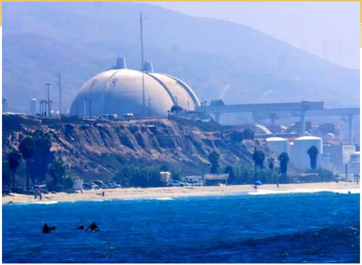

Source: inhabitat.com

- Contract theory provides an active way to do **customer segmentation**.
- Known consumer preferences can be used for targeting and incentive design.
- We are working with Lawrence Berkeley National Lab and Southern California Edison to do customer segmentation for targeting customers for demand response programs.

SOUTHERN CALIFORNIA EDISON®

An *EDISON INTERNATIONAL®* Company

FORCES FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

BERKELEY LAB Lawrence Berkeley National Laboratory

Workshop on Energy Analytics at the 2014 CDC!?

Conclusion

- Energy analytics is a very interesting and promising field
- Proposed a workshop on Energy Analytics at the Control and Decision Conference (CDC) to be held in Los Angeles, Dec 2014.
- Working close to companies enables us to focus on relevant problems and access data.
- The fundamental problems at the core of risk management, load forecasting, voltage optimization, outage management, etc. can be thought of as arising from **attacks** by adversarial agents.
- Disaggregation, customer segmentation, and incentive design all raise questions about **privacy**.

As a result, **energy analytics and game theory** must be used in conjunction to create a **resilient smart grid** so that we can recovery from **adversarial attacks** as well as **faults**.

- These ideas and framework extend to S-CPS.