FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Building cyber-enabled resilience in infrastructure networks

Saurabh Amin
MIT

All Hands Meeting, June 9-10, 2016

Berkeley UNIVERSITY OF CALIFORNIA
Massachusetts Institute of Technology
NSF
MICHIGAN
VANDERBILT UNIVERSITY

Main achievements

- We identified a unique set of domain-level RC+EI questions and approached them using different but complementary models and algorithms
- We derived new structural insights and design guidelines to improve efficiency and resilience by applying RC+EI theory to multiple domains
- Old decision-theoretic models (but new look):
  - Routing games, intrusion detection and network inspection
  - Water/ power/ traffic network control, stochastic hybrid systems (both data- and model-based)
- New economic models:
  - Interdependent security, cyber insurance, optimal dynamic policies for clean tech and infrastructure expansion
  - Applications of games and mechanisms: asymmetric info, multi-stage / sequential
- **Disclaimer:** I will present a summary of our team's progress, with some bias toward the topics in which I am involved. As a consequence, I will miss highlighting some important results.

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Summary of progress

1) **CPS resilience and security evaluation**
   - Testbed: evaluation platform for detection and mitigation
   - Security games: attacker-defender games over distribution networks
   - Cyber-physical security risks: interdependent security & insurance
2) **Incentives and Mechanism design**
   - Clean energy & renewables: optimal policy design under dynamic CPS constraints and endogenous evolution of technology
   - Data markets: pricing and regulation
   - Mobility services: routing, learning, and effect of information
3) **Resilient diagnostics and control algorithms**
   - Stochastic hybrid systems: learning operational models and control
   - Network diagnostics and estimation: fast approx. algorithms
   - Demand management: under uncertain supply

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Part I: CPS resilience and security evaluation: RCPS tesbed

* Threat models
  * Malicious (embedded) appln.
  * Spoofed sensor measurements
  * DDoS & deception attacks
  * Physical network faults
  * Cyber network reliability issues

* Domains
  * Fractionated spacecraft
  * *Vehicle control systems*
  * Road transportation systems
  * *Railway control systems*
  * Electric power T&D networks

* Research challenges
  * Security assessment
  * Resilient monitoring & control
  * Detection and mitigation of malicious applications

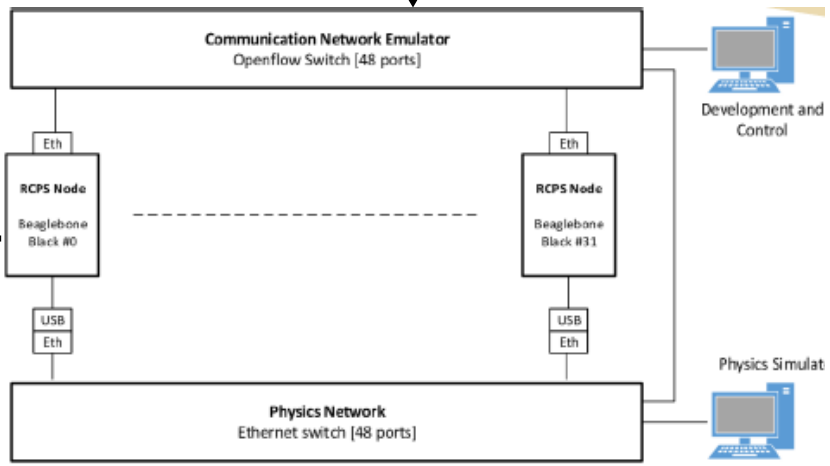

Role of Prof. Ilic's interaction variables?

(Karsai et al.; support, in part, by DARPA, AFRL, NIST, DoD)

Threat models + Domain models map to Network security games

Water distribution: network sensing under disruptions due to faults or malicious attacks

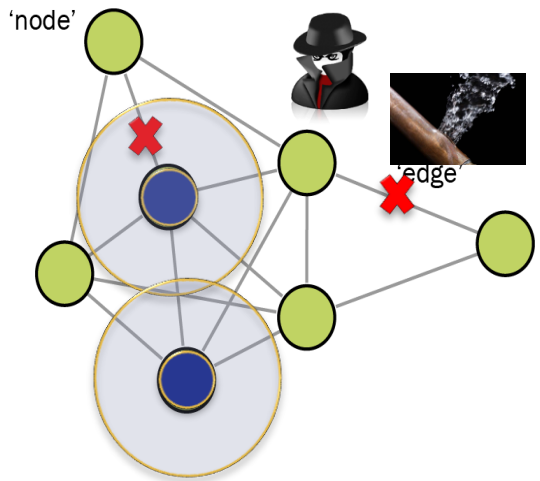

(Sela, Abbas, Dahan, Amin, Koutsoukos,…)

Transportation networks: optimal routing under link disruptions due to incidents or malicious attacks

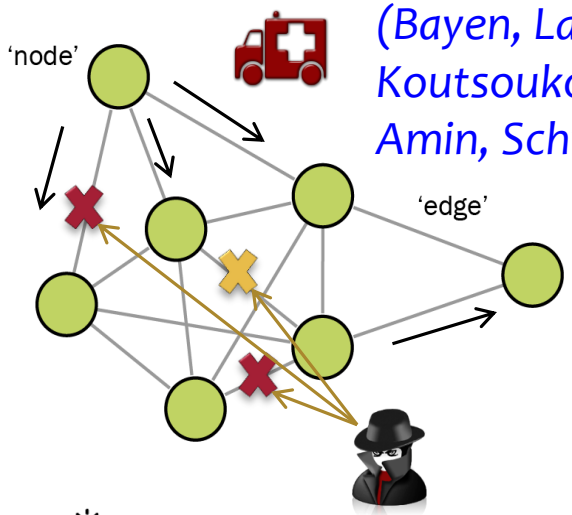

(Bayen, Laszka, Koutsoukos, Amin, Schwartz)

Electricity distribution: control in the face of random or strategic disturbances

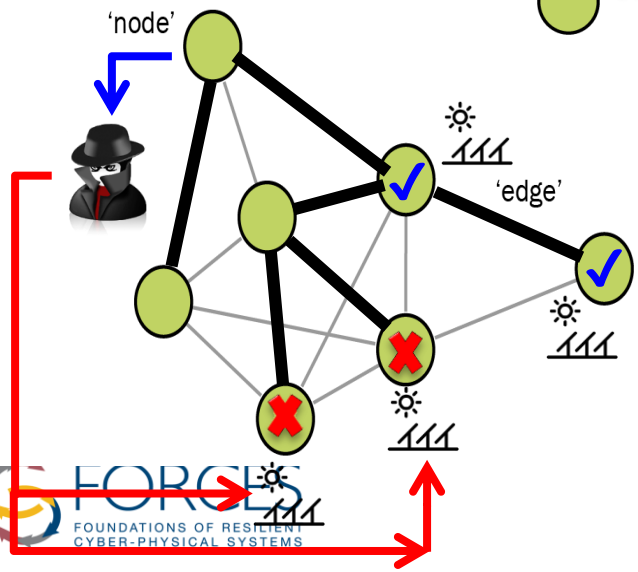

(Amin, Hiskens, Karsai,…)

Water network sensing: faults

Objective: For a given network, find minimum number of sensors and their placement:

* Detection: when a detectable event occurs, at least one sensor detects it.
* Location identification: For any pair of events, at least one sensor gives different output for them

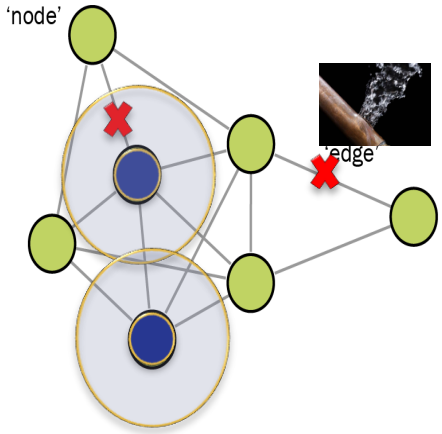

(Sela, Abbas, Koutsoukos, Amin)

Approach:

* Detection problem as Minimum Set Cover (MSC) problem, and Location Identification problem as a Minimum Test Cover (MTC) problem
* Augmented greedy algorithm which provides significant computational improvement while maintaining same approximation ratio as the classical greedy algorithm to solve MSC.

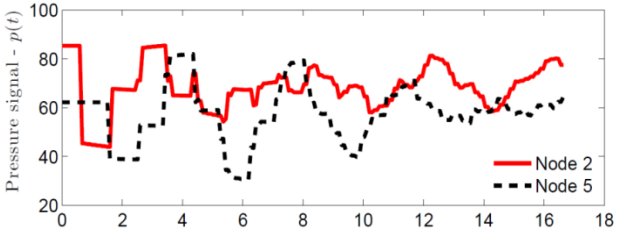

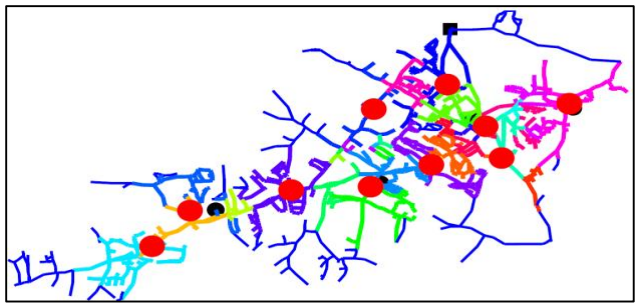

Extensions: Resource bounded monitoring

* Simultaneous placement and (sleep) scheduling
* Min. time to detect

(Abbas, Laszka, Koutsoukos)

Water network sensing: faults

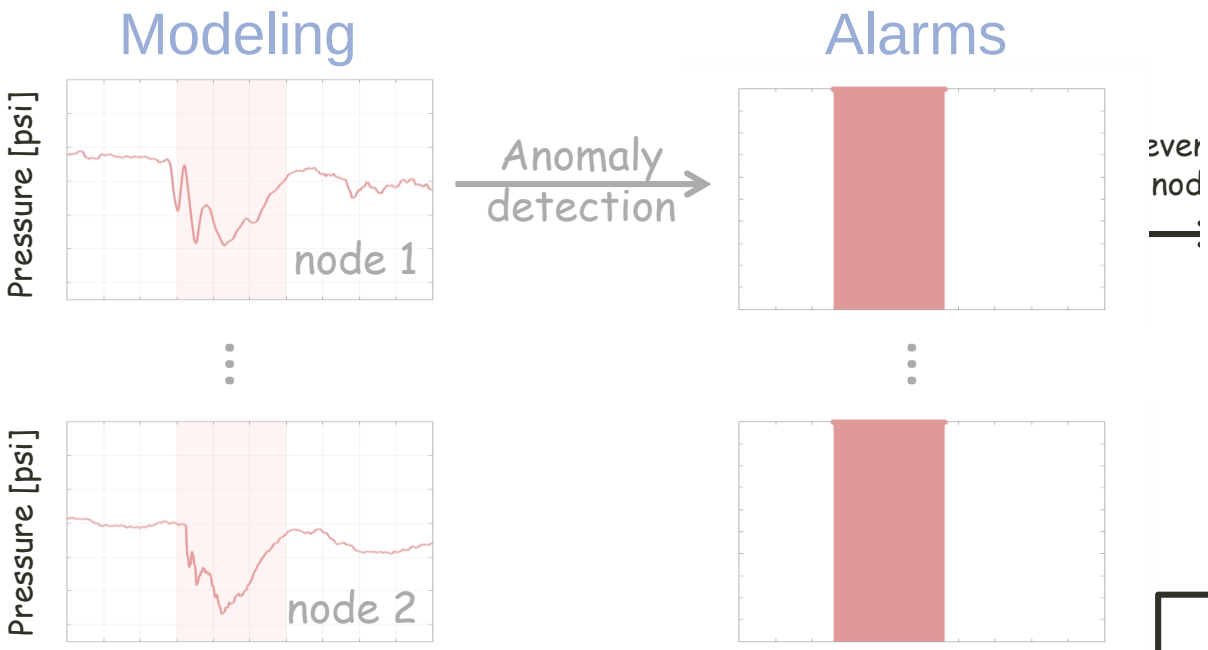

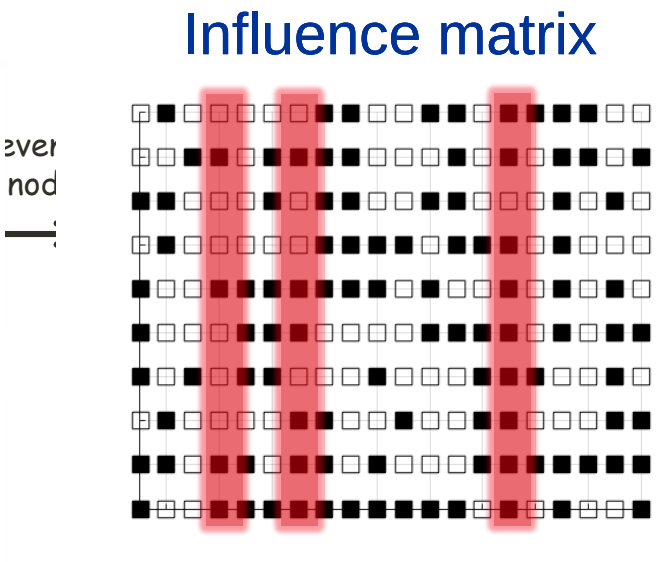

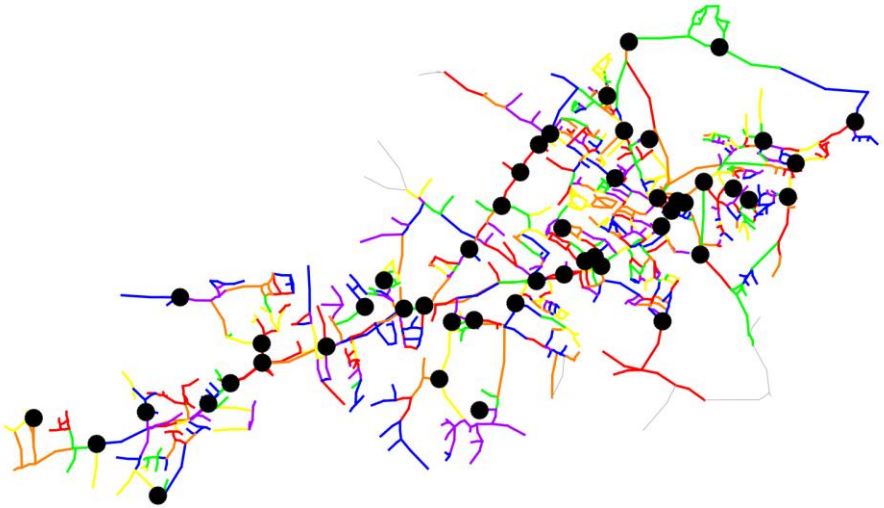

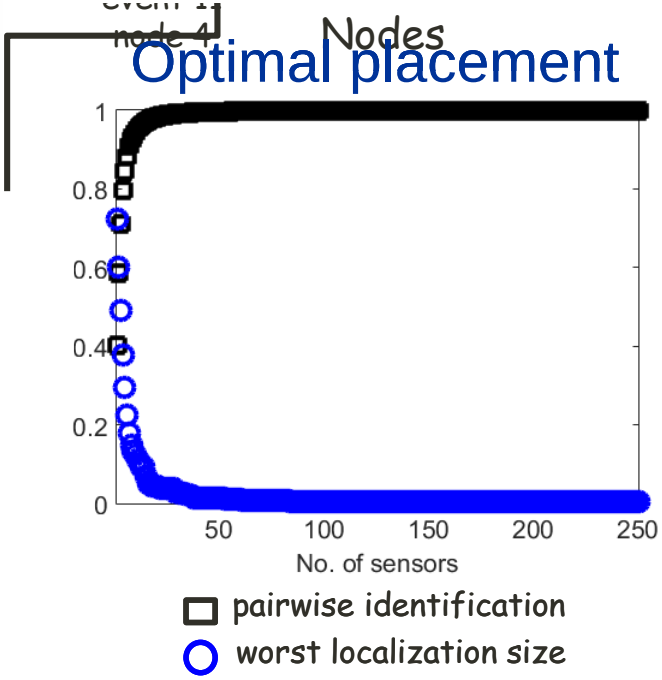

Data source: SMART project

Water network sensing: attacks

Objective: For a given network that attacks (link disruptions or contaminant injections), how to place and operate minimum number of sensors to achieve a target detection rate.

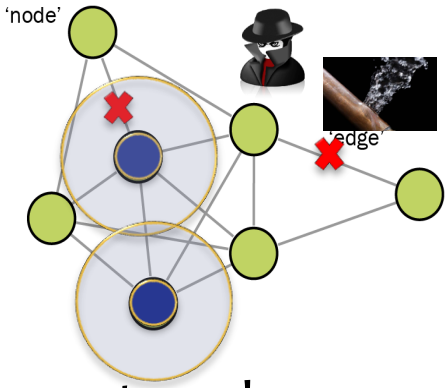

Model: We formulate a simultaneous game over the network:

- Attacker: simultaneously disrupts multiple edges
- Defender: strategically chooses a sensor placement

-- both players are resource-constrained

Approach:

- Zero-sum game over network with a general range sensing model
- Characterize support of equilibrium strategies in terms of minimum set cover and (extended) maximum matching problem
- Computable upper & lower bound on number of sensors to achieve target detection rate
- Prescribe randomized sensing strategies based on mixed Nash eq.

Water network sensing: attacks

Attacker

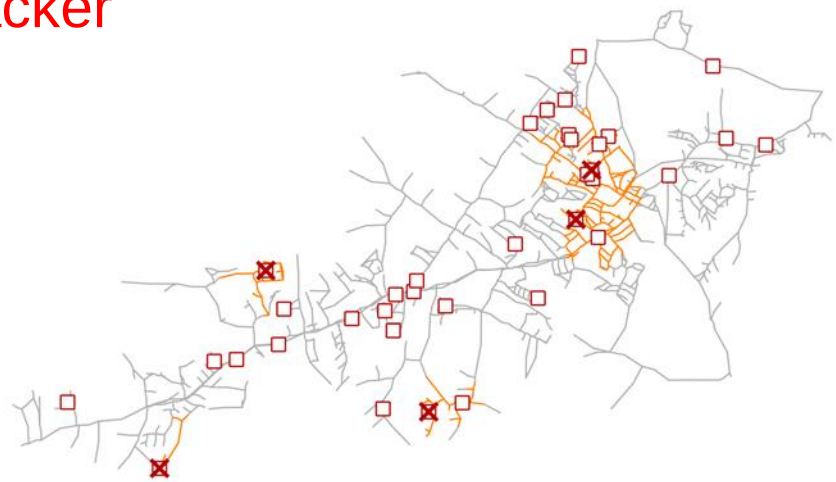

- EMM: maximum set of links that are covered by any node at most once

Payoffs

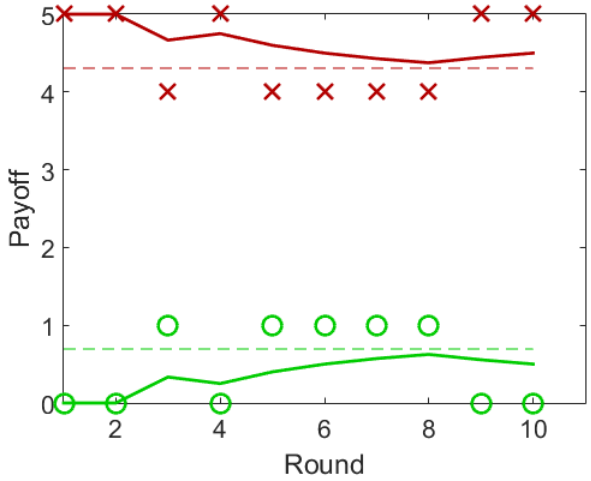

Defender

- MSC: minimum set of nodes that cover all edges

(Dahan, Sela, Amin)

Network intrusion (attack) detection

* Game-theoretic model for finding detection thresholds for intrusion detection systems in the face of strategic attacks
* Defender chooses strategy anticipating that attacker will play BR.

Defender:
Select false-negative probability f_s for each system.

Attacker:
Select a subset A of systems to attack.

$$\mathcal{L}(\boldsymbol{f}, A) = \mathcal{D}(A) \prod_{s \in A} f_s + \sum_{s \in S} C_s \cdot FP_s(f_s),$$

$$\mathcal{P}(\boldsymbol{f}, A) = \mathcal{D}(A) \prod_{s \in A} f_s$$

$$\arg\min_{\substack{\mathbf{0} \le \boldsymbol{f} \le \mathbf{1} \\ A \in \texttt{Best\_Response}(\boldsymbol{f})}} \mathcal{L}(\boldsymbol{f}, A)$$

$$\arg\max_{A \subseteq S} \mathcal{P}(\boldsymbol{f}, A)$$

* **Contributions:** 1) greedy heuristic; 2) extension to dynamic attacks

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

(Abbas, Laszka, Koutsoukos)

Control of electricity DNs: adversarial failures

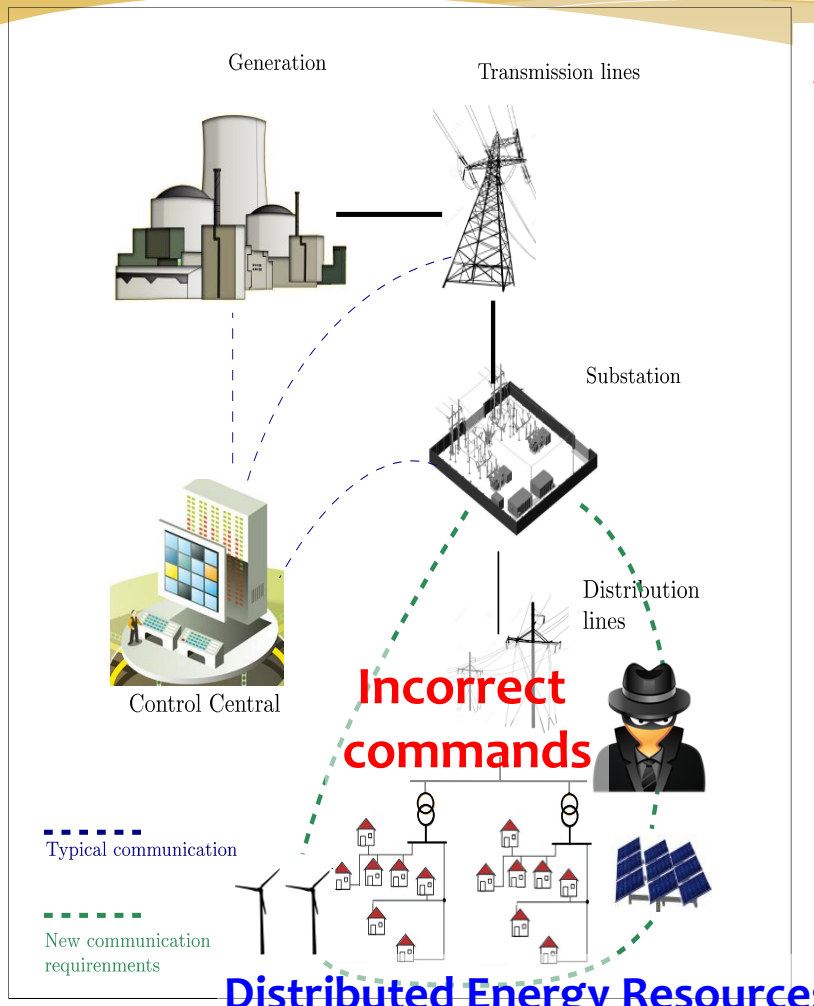

Adversary:

- Hack substation communications
- Introduce incorrect set-points and disrupt DERs and/or protection devices
- Create supply-demand mismatch
- Cause frequency and voltage violations
- Induce network failures (cascades)

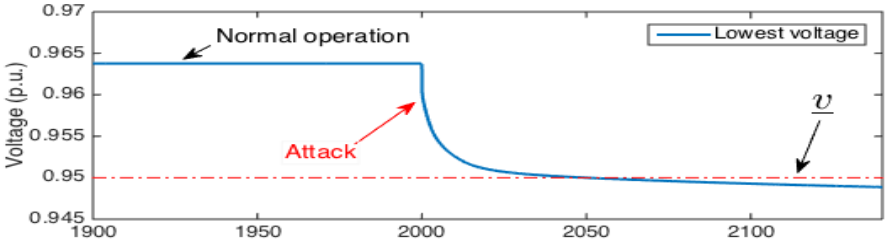

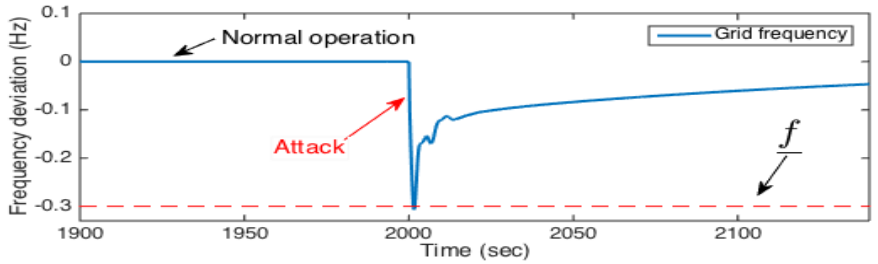

Resilient control of electricity distribution nets

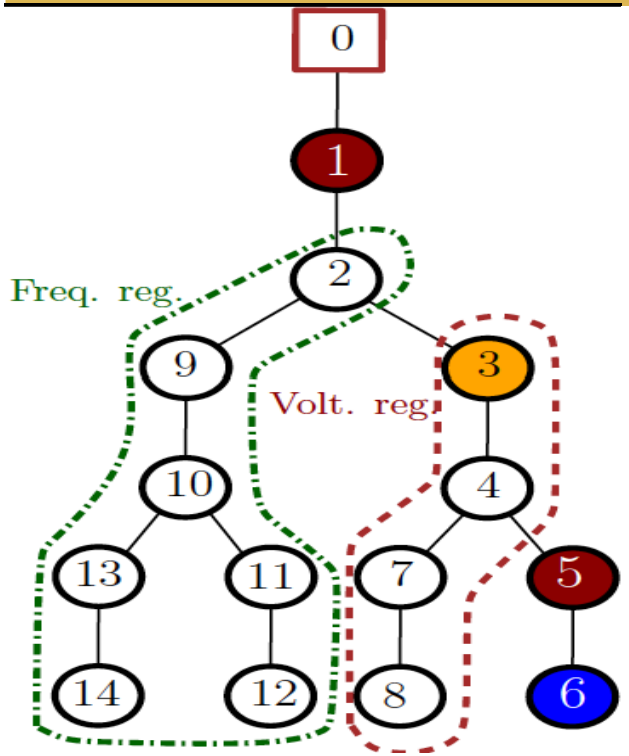

Approach: Network interdiction problem

- Detect attack and find worst affected nodes
- Launch distributed energy resources and/or reconfigure network
- Control voltage and frequency violations (distributed control) & limit network failures
  - *Extension: dynamic OPF solutions by Hiskens et al. ?*

Attacker-Defender interaction

- Attacker compromises DERs at 1, 5, 6
- Critical node 3 partitions network:
  - Subnet 1: control frequency
  - Subnet 2: regulate voltage.
- Decentralized control: new set-points

(Shelar, Amin)

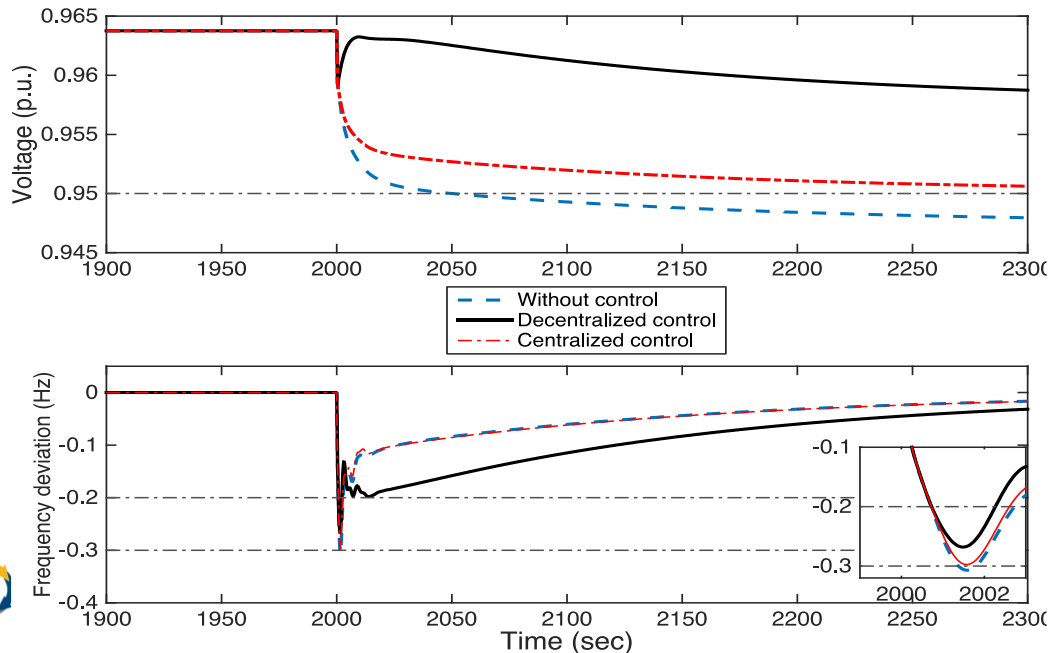

Network routing under adversarial disruptions

Strategic game on a flow network:

- **Network operator** (Defender): routes feasible flow through network to maximize her value of effective flow but faces transportation costs
- **Attacker**: disrupts one or more edges to maximize her value of lost flow but also faces cost of disrupting edges.

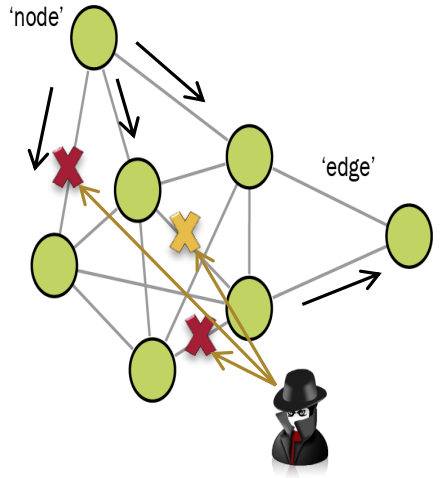

Features of the model:

- Strategic routing to maximize effective flow net transportation cost
- No rerouting of flow after attack
- Structural insights on players' eq. strategies
- Identification of vulnerable edges
- Optimal investment in network security

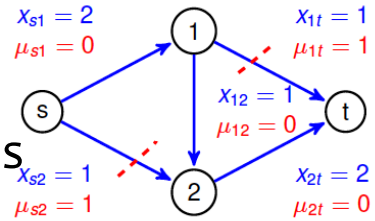

Initial flow and attack.

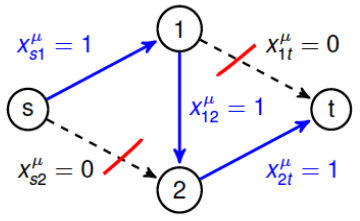

Resulting effective flow

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

(Dahan, Amin, also Laszka, Schwartz)

A metric of network vulnerability

Which edges are vulnerable to strategic disruptions?

- In equilibrium, any edge that is attacked with a positive probability must be saturated (i.e. the capacity bounds are met) by **every min-cost max-flow**. We call such an edge **vulnerable**. We view fraction of such edges as a vulnerability metric.
- Interestingly, one can find equilibria in which the attacker targets edges that are not part of a min-cut set (which is the notion of vulnerability in non-strategic situat

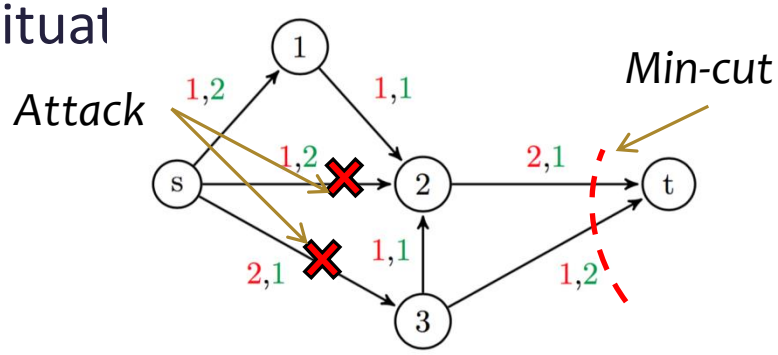

- Indeed, every edge that is part of a min-cut set is vulnerable because the edges that are part of a min-cut set are saturated by every min-cost max- flow.

Interdependent security model of CPS risks

- **Problem:** Estimate CPS risks while accounting for cyber-physical nature and strategic nature of attacks
- **Contribution:** Interdependent security modeling framework
  - Can model both discrete and continuous security choices
  - Strategic attackers: Cost of attack and cost of being caught versus benefit of successful attack
  - Strategic defenders: Cost versus benefit of security investments
  - Estimation of Nash cost for *i)* different topologies and *ii)* degree of interdependence based on cyber-physical interactions
  - How to steer system to superior equilibrium?
    - Quantification of policy impact
    - Comparison across different policies

(Schwartz, Sastry, also Laszka, Amin)

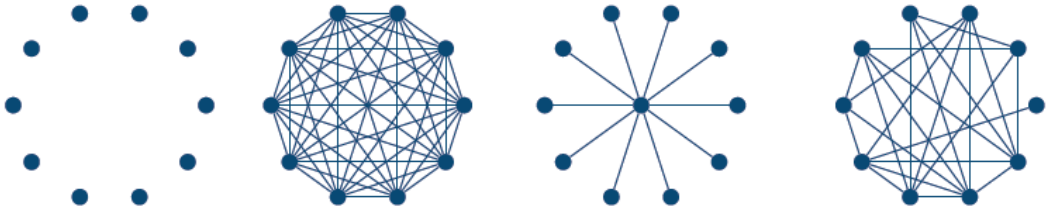

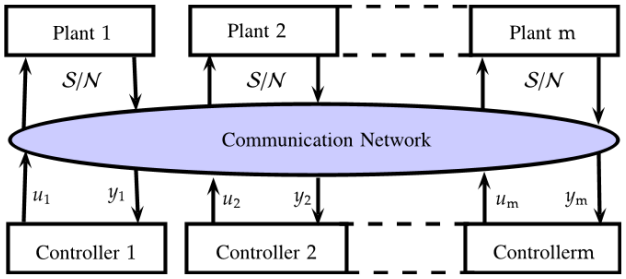

Part II

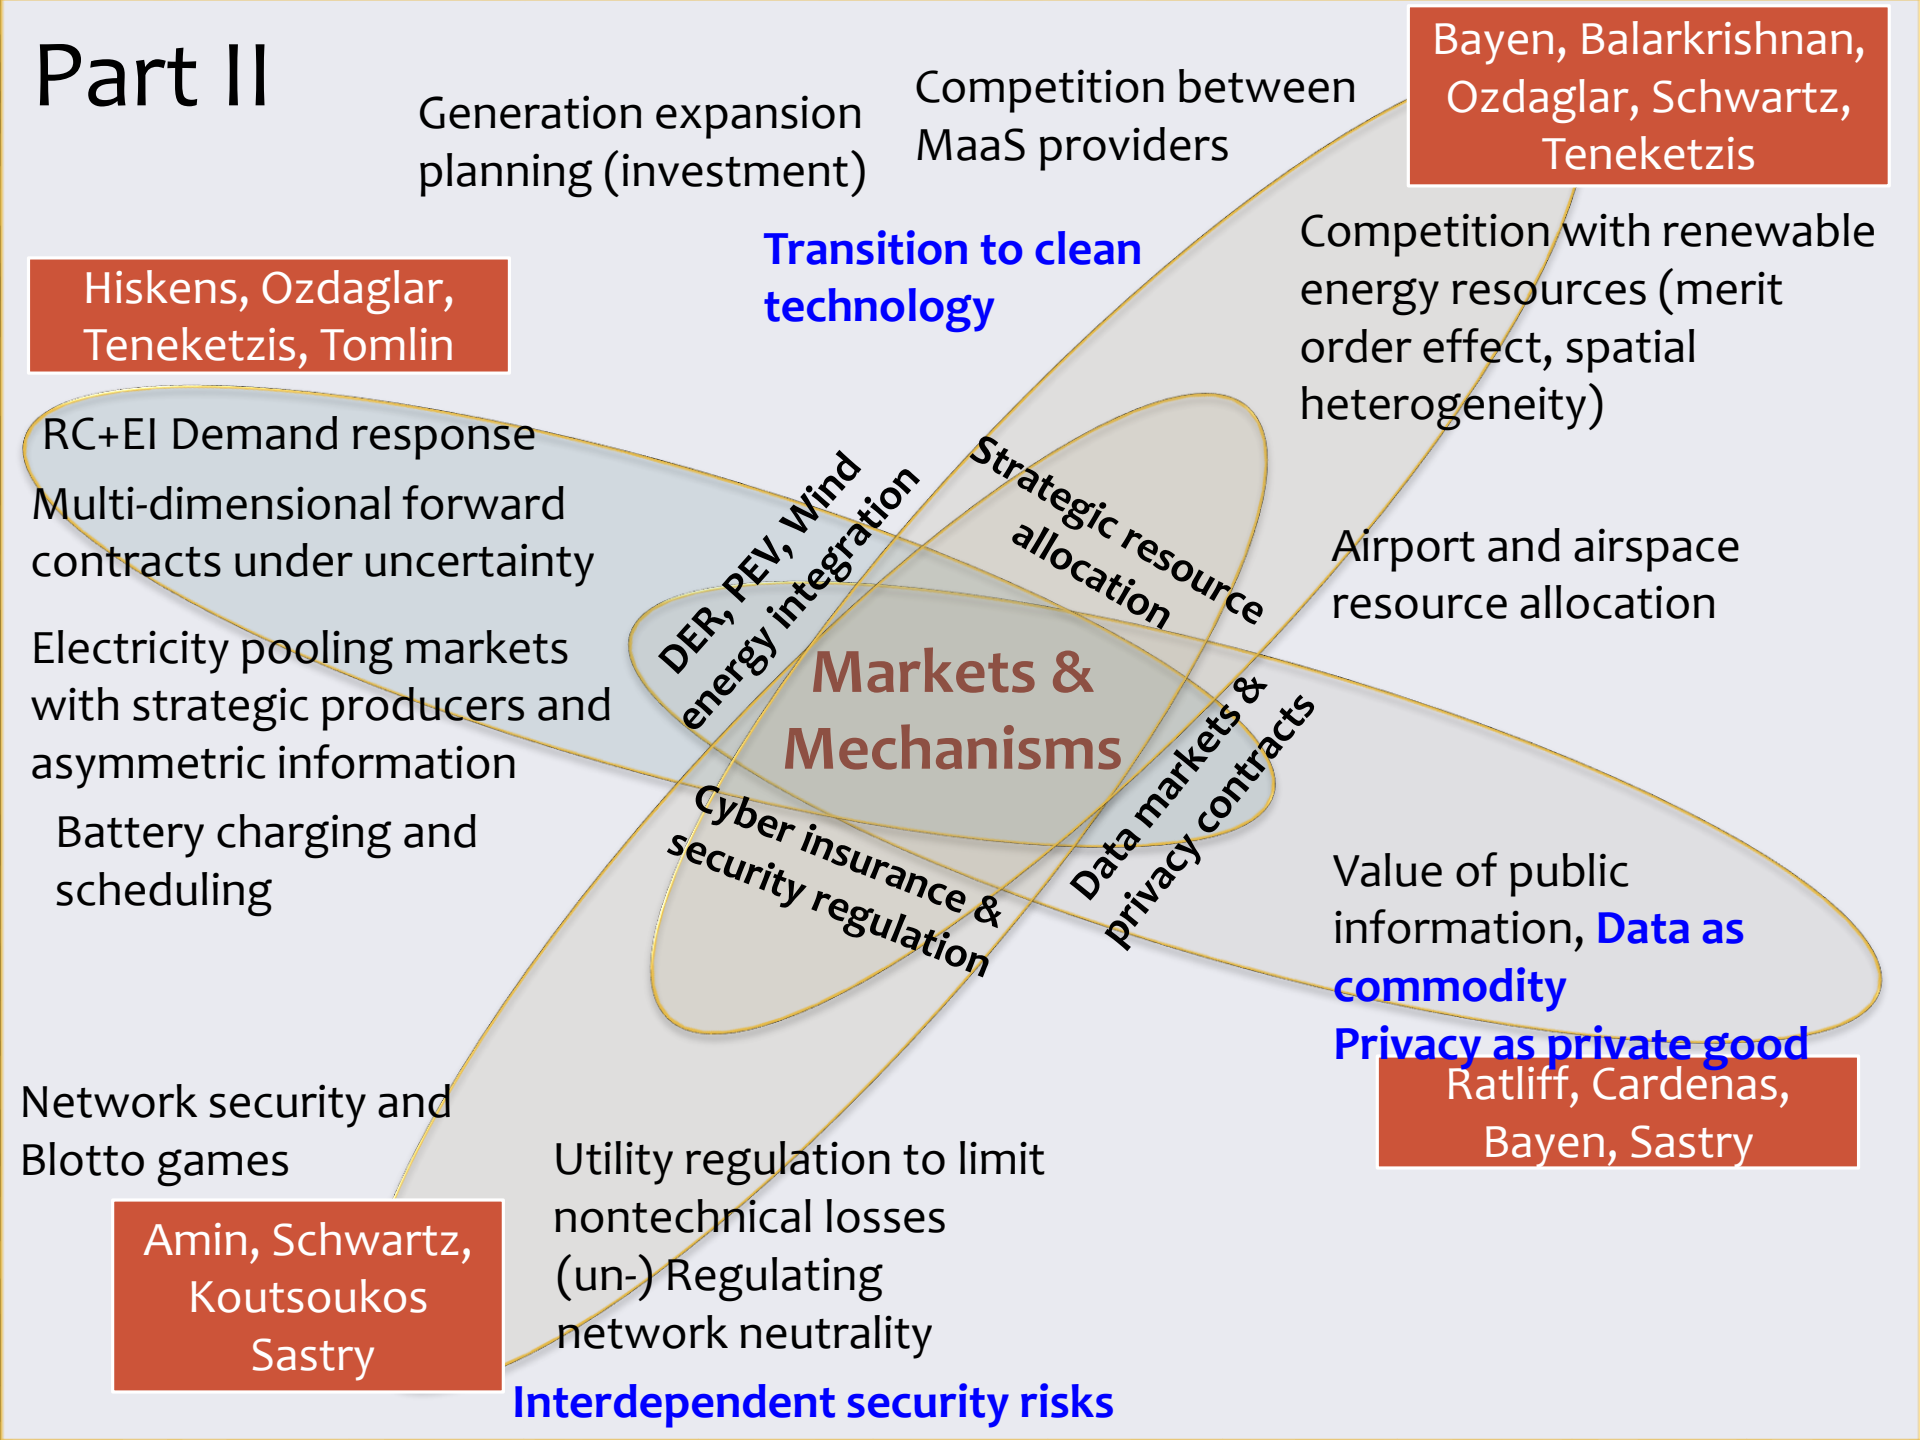
Generation expansion planning (investment)
Competition between MaaS providers
Bayen, Balarkrishnan, Ozdaglar, Schwartz, Teneketzis
Transition to clean technology
Competition with renewable energy resources (merit order effect, spatial heterogeneity)
Hiskens, Ozdaglar, Teneketzis, Tomlin
RC+EI Demand response
Multi-dimensional forward contracts under uncertainty
Electricity pooling markets with strategic producers and asymmetric information
Battery charging and scheduling
DER, PEV, Wind energy integration
Strategic resource allocation
Markets & Mechanisms
Cyber insurance & security regulation
Data markets & privacy contracts
Airport and airspace resource allocation
Value of public information, Data as commodity
Privacy as private good
Ratliff, Cardenas, Bayen, Sastry
Network security and Blotto games
Amin, Schwartz, Koutsoukos Sastry
Utility regulation to limit nontechnical losses
(un-) Regulating network neutrality
Interdependent security risks

Economic Incentives (EI) for promoting transition to clean tech & renewables

Four major topics of enquiry by CPS community:

1. Network control (Hiskens, Amin, Koutsoukos)
2. Demand shaping or load shifting (Tomlin, Yang, Hiskens, Sastry, Schwartz, Amin)
3. Energy storage (Hiskens, Yang, Tomlin)
4. **Dynamic** mechanisms for integrating renewables (Teneketzis, Ozdaglar)

… but we need an economic/EI framework to design optimal policy(-ies) to incentivize transition to clean tech that accounts for

* Dynamic and endogenous evolution of clean tech, including engineering and physical constraints of different energy technologies
* Evolution of energy production via different means (hence, carbon level)

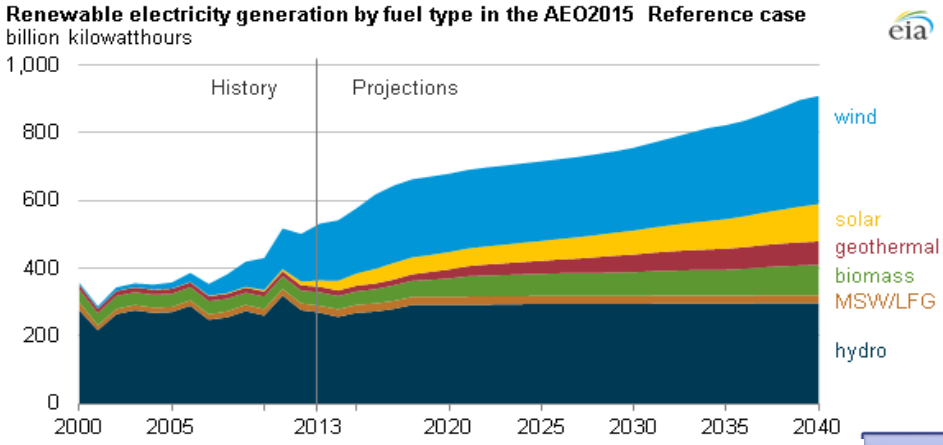

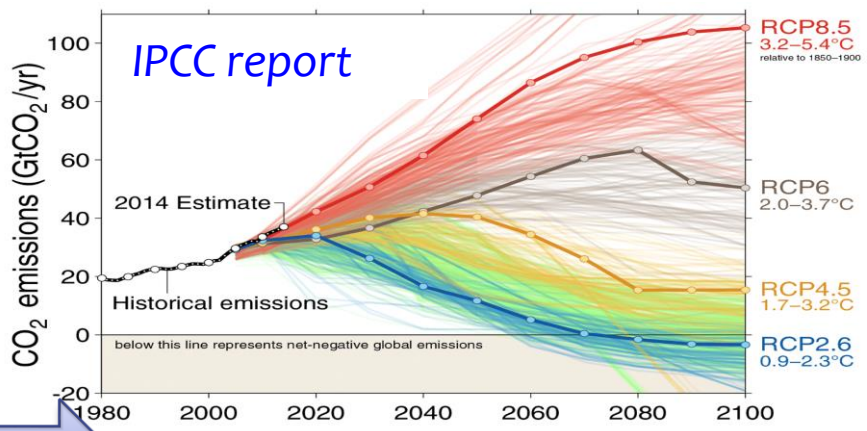

Projected technological change ? *Projected climate change*

EI framework for clean tech transition: an industry equilibrium approach

- Two-layer (sequential game) model: ***(Acemoglu, Ozdaglar, Yang)***
  - Lower layer (LL): Energy supply evolves over time according to an industry equilibrium arising from interaction between consumers & producers
  - Upper layer (UL): Social planner chooses optimal dynamic policy (*tax sequence*), accounting for LL industry equilibrium as a constraint
- Main features of the model:
  - Considers heterogeneous firms: conventional, clean, renewable energy
  - Firms have entry, exit, & upgrade options; face growth rate of renewables
  - Equilibrium production and carbon levels are endogenous
- Results: Industry eq. exists and characterized via dynamic programming
  - ➔ Firms stage payoffs are monotone in productivity terms
  - ➔ Entry /exit decisions based on "threshold rules"
  - ➔ Endogenously determines productivity distribution, supply, & carbon levels
  - ➔ LL solution (*effective prices*) constrain the UL optimal tax sequence.
- Outcome: We can characterize and compute optimal tax sequence
- These prescriptions are different from the ones obtained from models that ignore endogenous change and industrial equilibrium aspects

h-CPS data: disaggregation, privacy, value

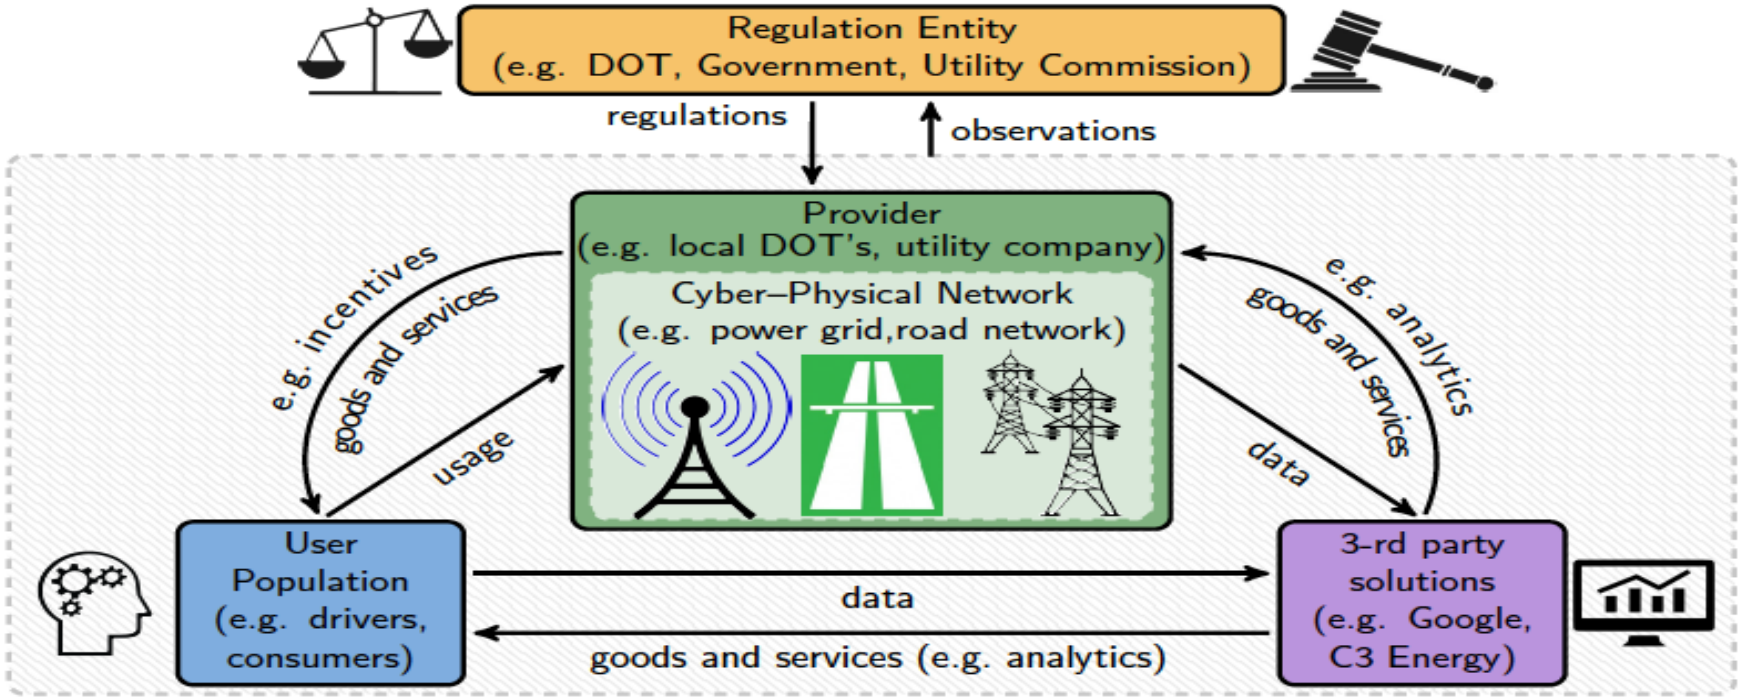

1. Data disaggregation

- Demand models and load forecasting
- Efficiency versus privacy
  - New notion of **inferential privacy**
  - Economic instruments for consumers with low vs. high privacy

(Ratliff, Dong, Mazumdar, Sastry)

2. Learning and games

- Urban parking (Seattle DOT)
- Estimation of queuing-based model to account for parking choices
- *Information structure* to maximize social welfare

(Calderone, Mazumdar, Ratliff)

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Coupled sequential decision problem

Waze | Google | Apple | INRIX | Bing (Microsoft)

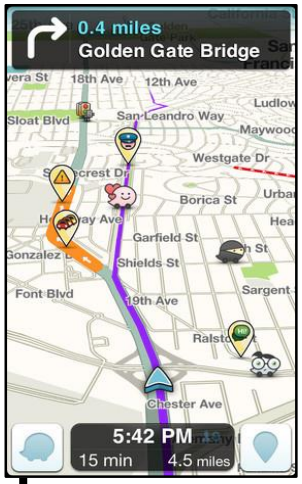

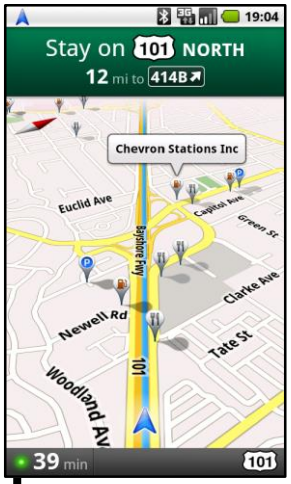

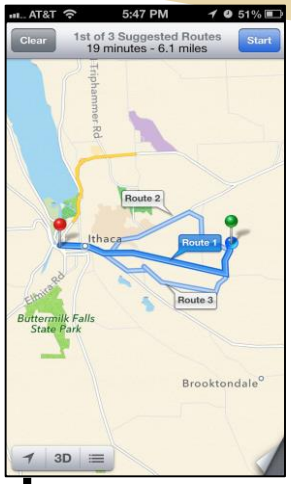

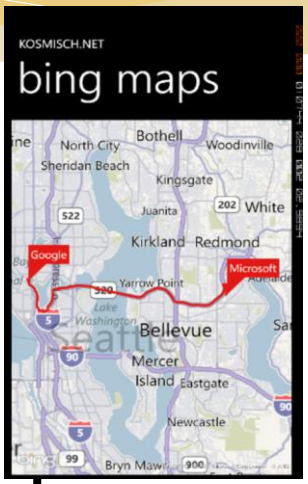

$p_{\text{Waze}} \sim x^{(t)}_{\text{Waze}}$ $p_{\text{Google}} \sim x^{(t)}_{\text{Google}}$ $p_{\text{Apple}} \sim x^{(t)}_{\text{Apple}}$ $p_{\text{INRIX}} \sim x^{(t)}_{\text{INRIX}}$ $p_{\text{Bing}} \sim x^{(t)}_{\text{Bing}}$

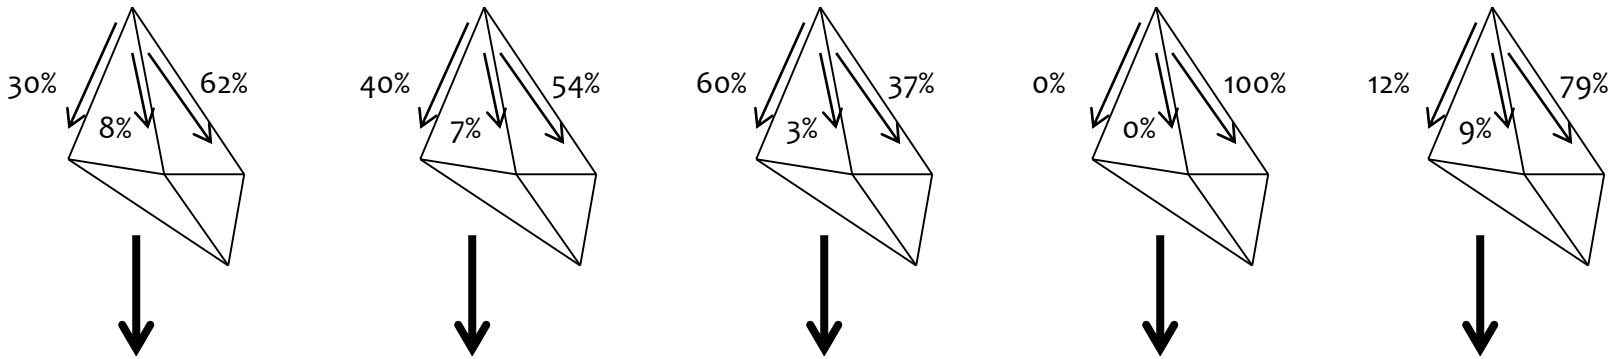

All users of each company "equal" by standards of the company i.e. same (shortest) travel time according to the company, "essentially" Nash.

(Bayen)

Routing games: distributed learning dynamics

- Routing games with multiple information providers and multiple populations (OD pairs)
- Make choice ---> Drive ---> Evaluate outcome ---> Learn
- Main questions:
  - Convergence of distributed learning dynamics; rates of convergence
  - Robustness to perturbations
  - Heterogeneous learning
- Convergence using regret analysis, stochastic approx.; convex optimization
- A new approach to design optimization algorithms using analysis of a class of continuous-time dynamical systems

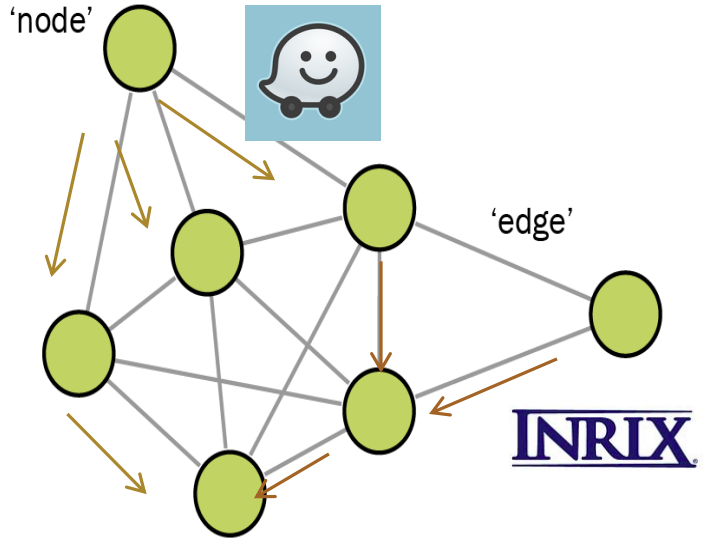

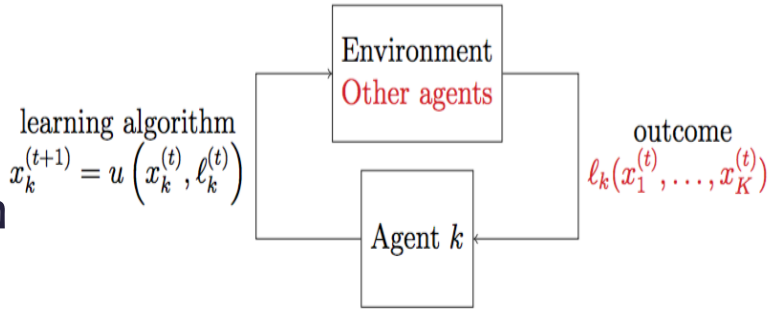

(Krichene and Bayen)

Effects of Information on Traffic Congestion

- Traffic information services (TIS) are changing how people make routing decisions
  - Inherent heterogeneities in TIS adoption and accuracy
- Previous work and our contribution:
  - [Arnott, De Palma, Lindsey]: effect of information using Vickrey's bottleneck model, but only for boundary cases (single informed player versus entire informed population)
  - [ben-Akiva, de Palma, Kaysi], [Mahmassani, Jaykrishnan]: identification of potential effects of information using traffic simulations
  - [Krichene and Bayen]: distributed learning and convergence dynamics
  - [Acemoglu, Makhdoumi, Malekian, and Ozdaglar]: "Informational Braess Paradox" and the effect of asymmetric info about available routes
  - [Liu, Amin, Schwartz]: Bayesian congestion games and effect of asymmetric info about network incidents
- How does heterogeneous information about traffic state (routes and incidents) affect the commuters' equilibrium route choices and costs

Effects of Heterogeneous Information on Traffic Congestion

- We introduce a Bayesian congestion game, in which players have private information about incidents, and each player chooses her route on a network of parallel links
- We characterize the Bayesian Wardrop Equilibrium of the game, and study how the cost to individual players and the social cost as a function of the fraction of highly-informed players.

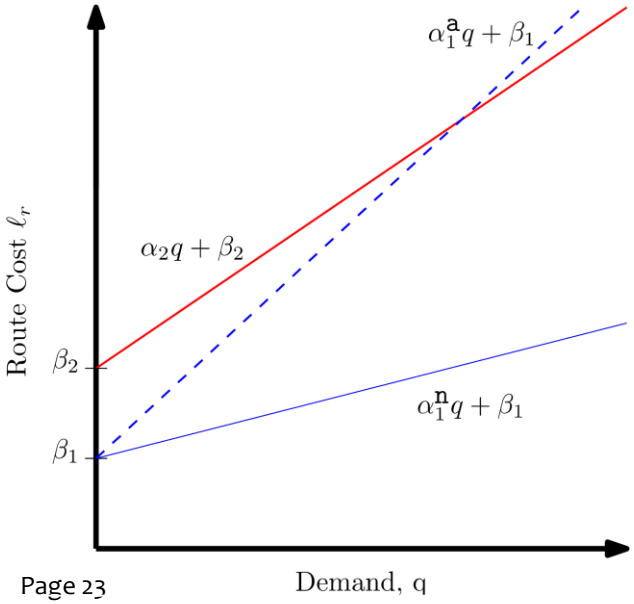

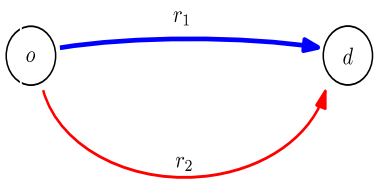

State-dependent route costs
- r_1: normal/accident states
- Drawn by Nature w/ fixed probability

Two commuter populations
- "H": receives signal
- "L": no signal

(Liu, Amin, Schwartz)

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Populations

- Total inelastic demand D
- Fraction λ^{H} in population H (High-information)
- $1-\lambda^{\mathrm{H}}$ in population L (Less-information)

Signals and Types

- Population H receives private signal that matches the state $y^{\mathrm{H}}=\theta$ with probability η^{H}
- Population L receives private signal $y^{\mathrm{L}}=\theta$ with probability $\eta^{\mathrm{L}}<\eta^{\mathrm{H}}$
- Signal determines populations' "types"

Model outline

Bayesian Congestion Game

- Models heterogeneous information about incidents and beliefs about other players
- Populations of players that receive different private information from different TIS (H and L)
- Key features:
  - Probability of incident p
  - Fraction of players with information λ^{H}
  - Accuracy of information η^{H}
- Recovers classical imperfect information and perfect information games at the boundaries

Bayesian Wardrop Equilibrium

For each player type, the expected costs of all utilized routes are equal and less than the expected cost of each unutilized route.

Beliefs

Each population has a belief $\mu^i(\theta, t^{-i}|t^i)$

- Probability distribution over the nature states θ and other population's type t^{-i}
- Conditioned on each population's type t^i
- Incorporates public information (common prior and common knowledge) and private information (TIS signal)
- Each population calculates expectation over its own belief

Bayesian congestion game

$$\Gamma_p = (\mathcal{I}, \mathcal{Q}, \Omega, \mathrm{T}, C, \mu),$$

- $\mathcal{I} = \{\mathrm{H}, \mathrm{L}\}$ Commuter populations
- $\mathcal{Q} = (\mathcal{Q}^i)_{i \in \mathcal{I}}$ Load vectors
- $\Omega = \{\mathbf{n}, \mathbf{a}\}$ Game states w/ element θ
- $\mathrm{T} = (\mathrm{T}_i)_{i \in \mathcal{I}}$ Type space for each population
- $C = (\ell_{\mathrm{r}}^{\theta})_{\mathrm{r} \in \mathcal{R}}$ Route cost functions
- $\mu = (\mu^i)_{i \in \mathcal{I}}$ Beliefs about states θ, and other populations' types t^{-i}

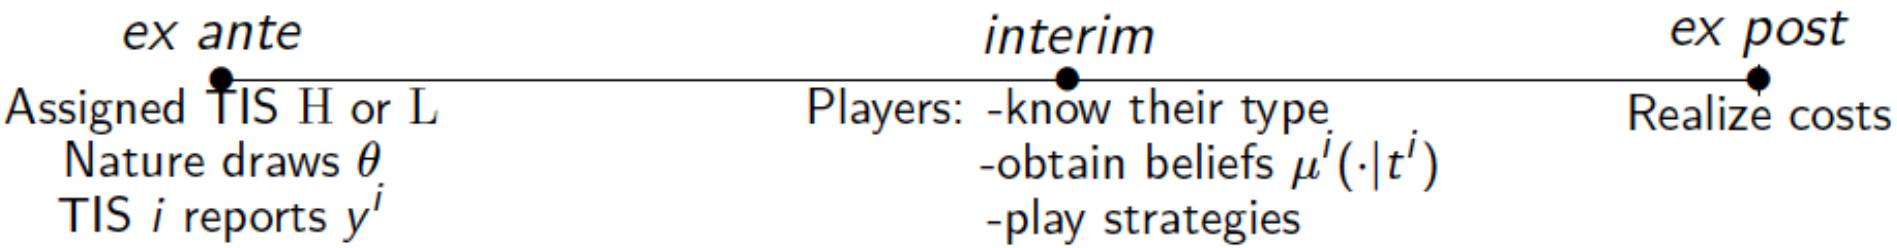

Figure : Timing of the game

Equilibrium Characterization

- Four qualitatively different equilibrium regimes
- Recover classical equilibria:
  - Complete info game when everyone is in population H
  - Imperfect info game when everyone is in population L

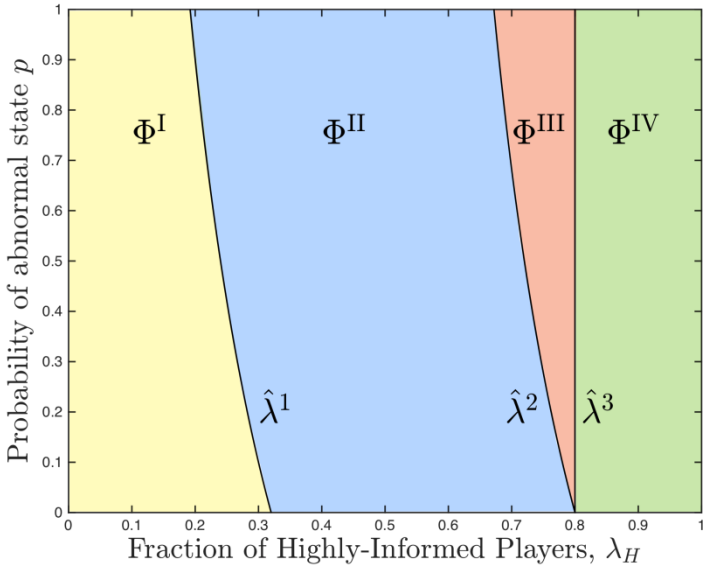

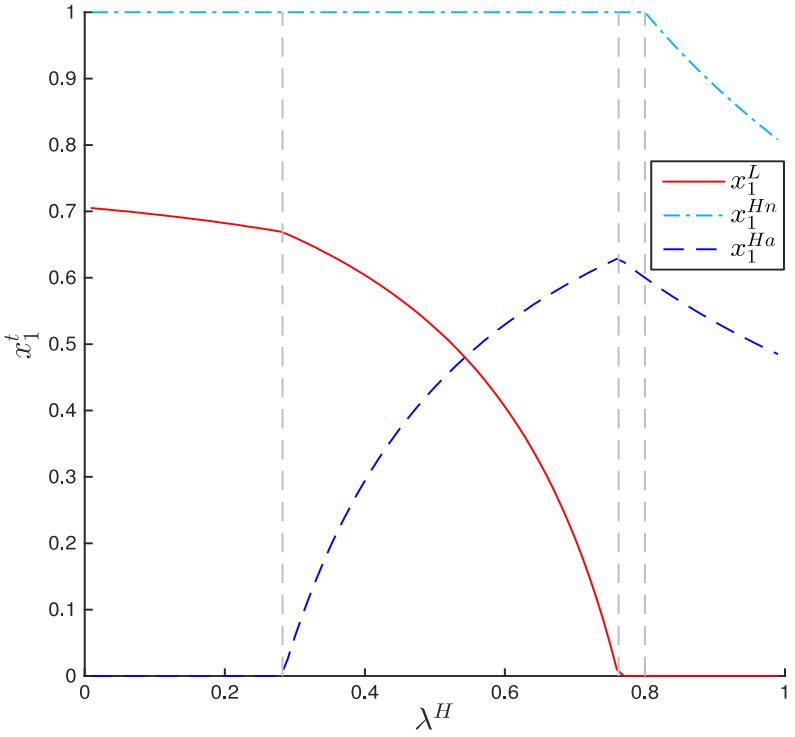

Individual Value of Information

- Value for Population H: as more players gain access to information, the value of information for population H players decreases
- Value for Population L: Benefits from other players having information *even though they don't receive information*
- Relative value of information: Positive up to a threshold, zero above, i.e. *there is no benefit of information if many others have it*

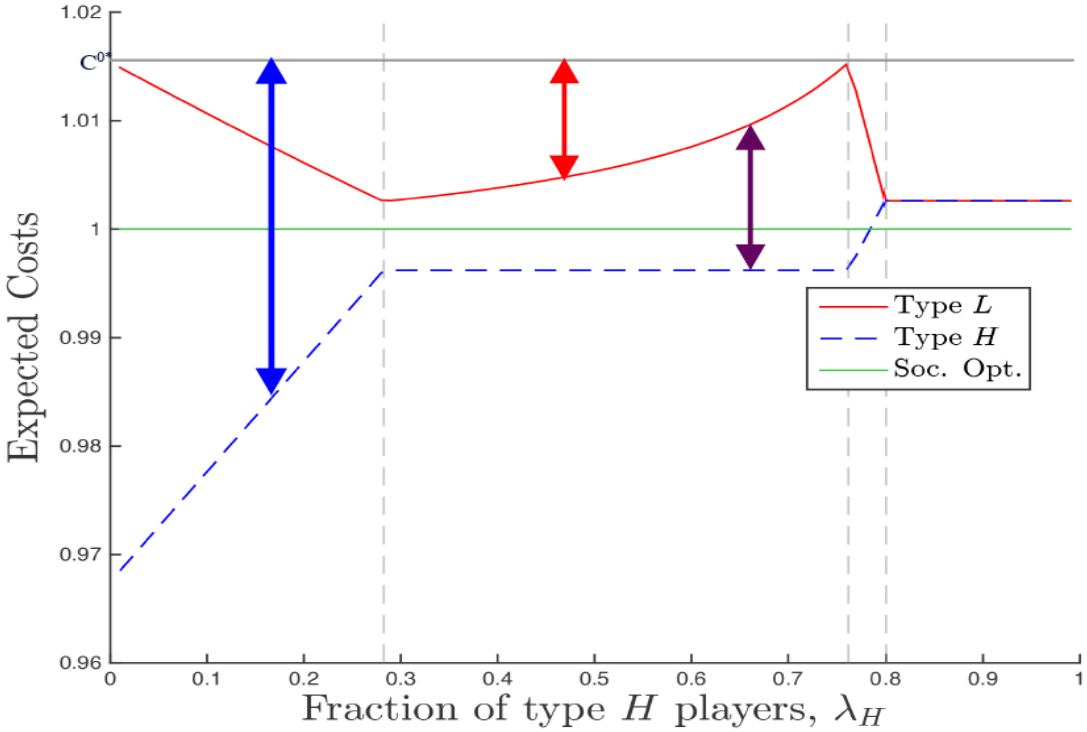

Social Value of Information

- There exists an "optimal" fraction of players with information
- This threshold is lower than the threshold where relative individual value goes to zero
- There exists a range of λ_H where it is individually advantageous for population L players to gain access to information, but harmful to society for them to do so

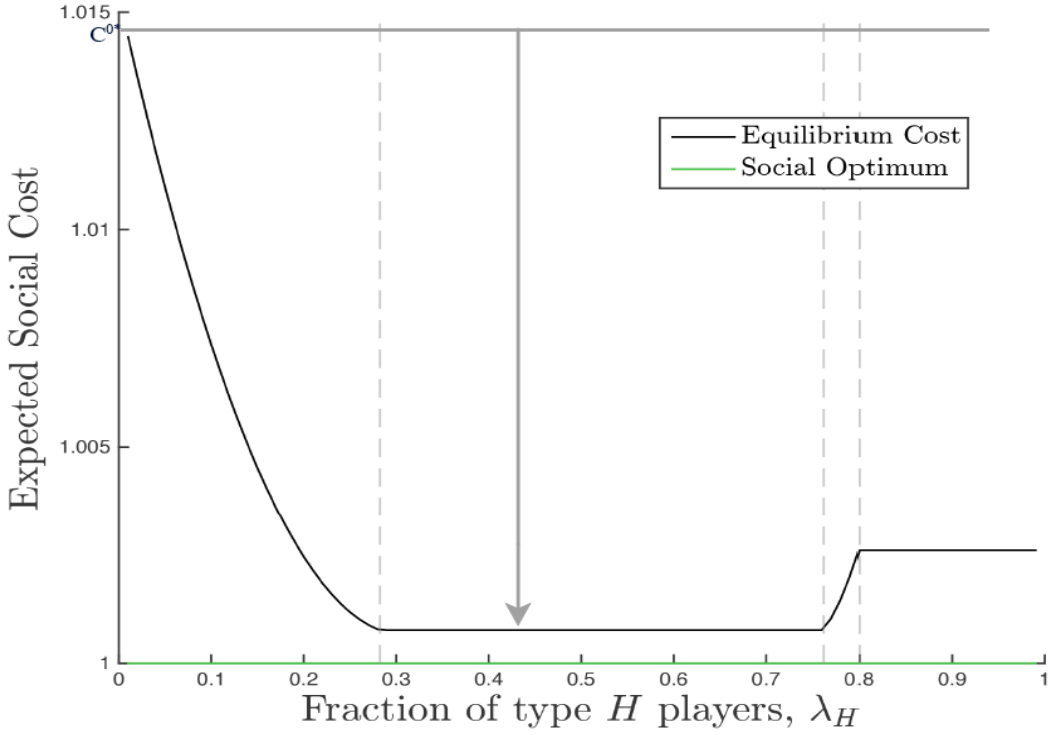

Part III: Resilient network control

Stochastic Hybrid Systems (SHS): state estimation, and control

* Random incidents, i.e., state dependent transitions and capacity fluctuations in freeway networks (PDMP): **Jin and Amin**
* Delay propagation in air-traffic networks (MJLS): **Gopalakrishnan and Balakrishnan**
* Non-intrusive load monitoring and utility learning (HMM and variants): **Ratliff, Dong, Sastry**
* Modeling of aircraft engine performance (Bayesian multiple linear regression, Gaussian processes): **Chati, and Balakrishnan**
* Secure state estimation under adversarial attacks (Kalman filters and switching variants): **Chang, Hu, and Tomlin**
* Quantifying user engagement in DR programs (nonparametric regression): **Balandat, Zhou, and Tomlin**
* Ensemble control of hysteretic loads (nonlinear hybrid systems): **Hiskens**

Air traffic dynamic with switching modes

Markov jump linear systems (MJLS)

- *Modes*: set of characteristic topologies
- *Continuous dynamics*: linear models estimated from real world data

Contribution: Stability analysis of MJLS with

- Periodic time-varying mode transition matrices
- Effect of temporal variations and continuous state resets

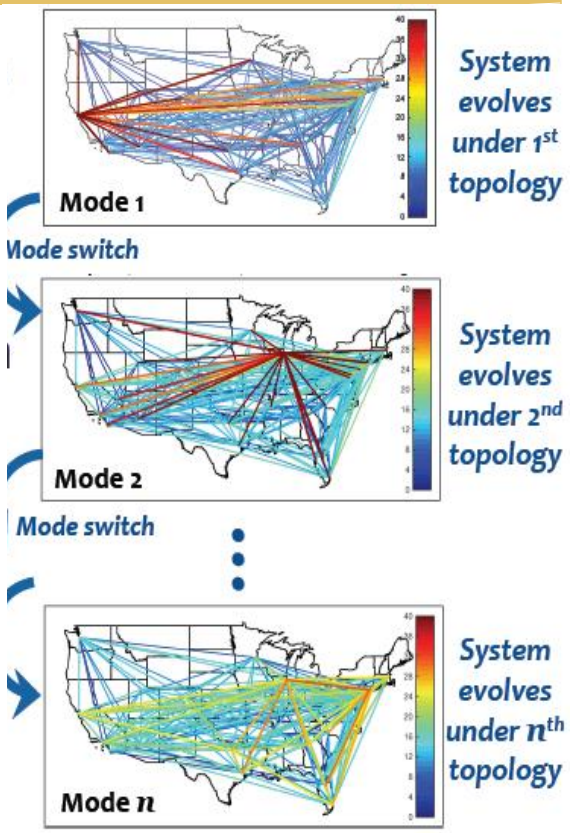

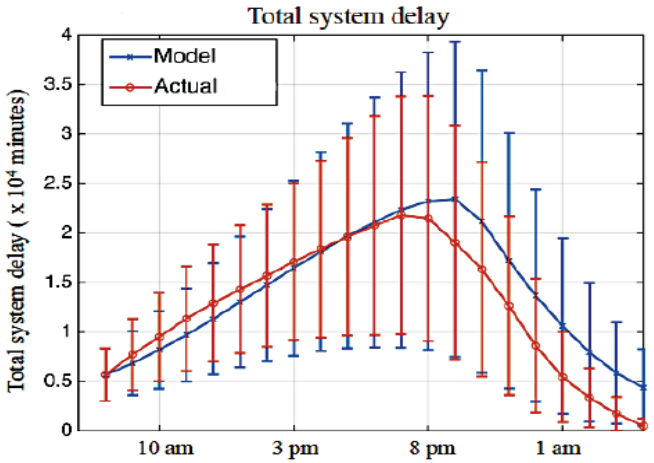

(Gopalakrishnan and Balakrishnan)

Traffic control under stochastic incidents

* **Incidents**: reduced capacity leads to perturbed dynamics (congestion)
* **Network operator**: Incident-aware dynamic network control strategies (both routing and ramp metering)

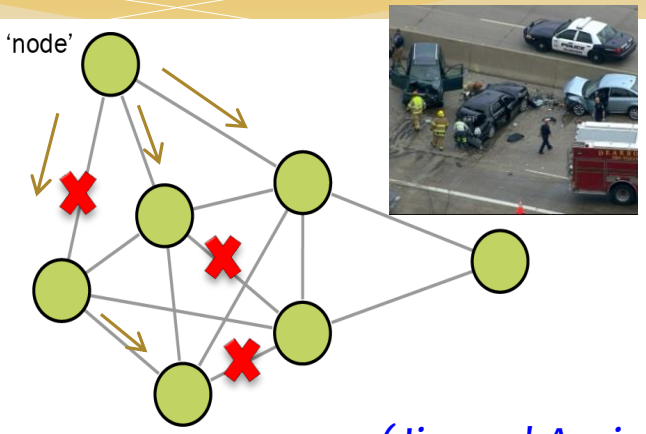

(Jin and Amin)

Outcomes:

- More intense but less frequent incidents worse than less intense but more frequent ones
- Better estimation of **effective capacity** can help improve operations:
  - Incorporate incident precursors in the control loop
  - Restrict access to incident locations
  - Avoid near capacity operations

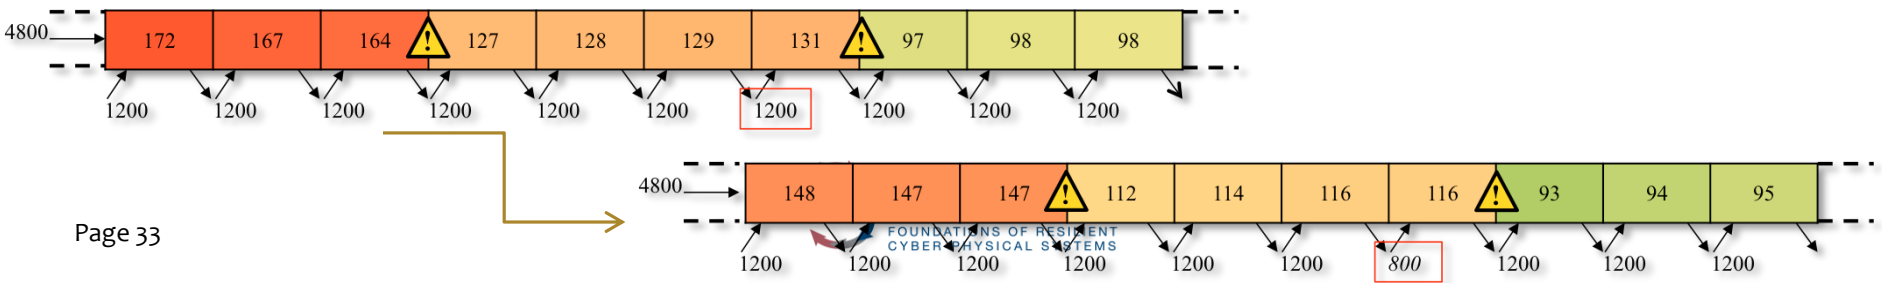

Traffic control under stochastic incidents

* Stochastic capacity model
  - Capacity switches between a set of values (e.g. between nominal & residual values).
  - Switches happen according to a Markov chain.
* Piecewise-deterministic queues (PDQs)
  - Stochastic switching saturation rates (modes)
  - Deterministic evolution between intermodal switches.
* PDQ model provides a way to quantitatively estimate effective capacity and delay due to stochastic incidents, and aid in design of incident-aware routing control strategies.

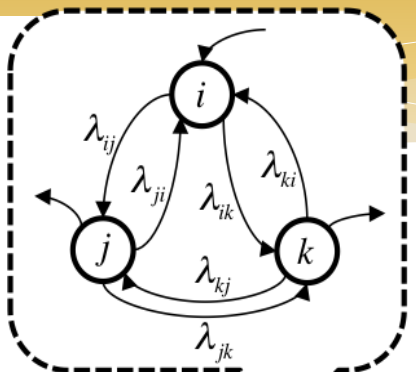

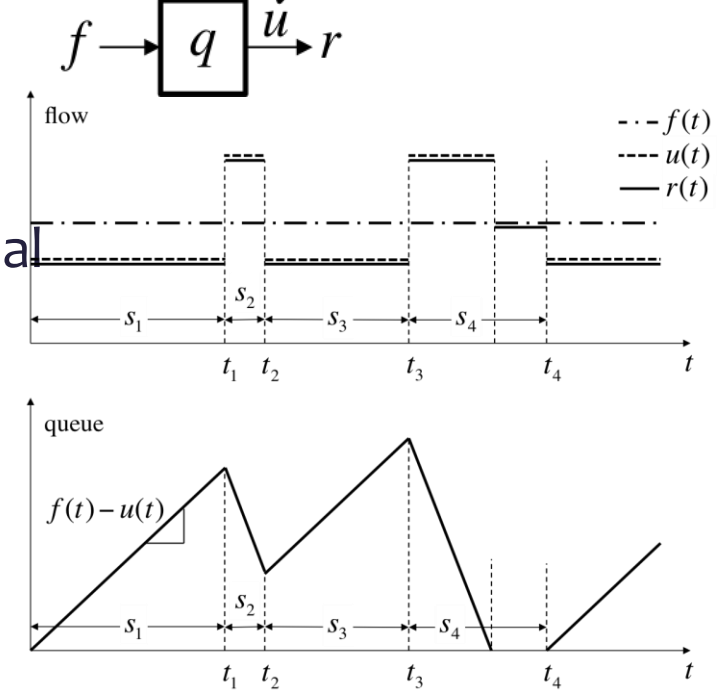

Traffic control under stochastic incidents

Interacting links:

- Links 1 & 2 have unreliable capacities: $\mathbf{u}_1$ in {1.6,0.8}, $\mathbf{u}_2$ in {1,0.5}.
- Link 3 has large capacity
- Outflow of link 1 = inflow of links 2 and 3

Main question: How to determine the splitting factor to minimize total travel cost (link travel cost + queuing cost)?

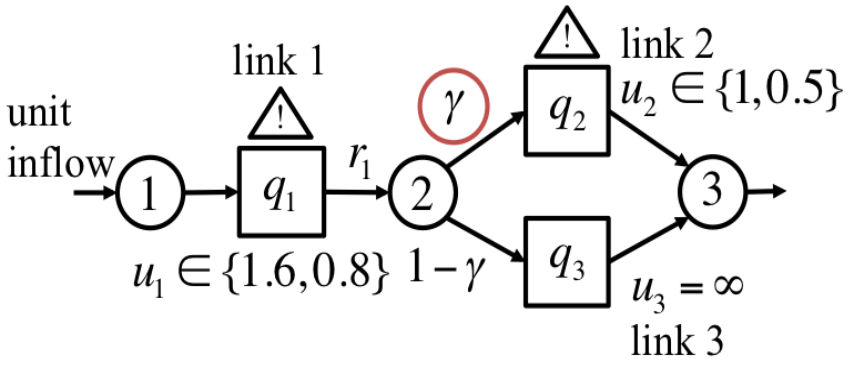

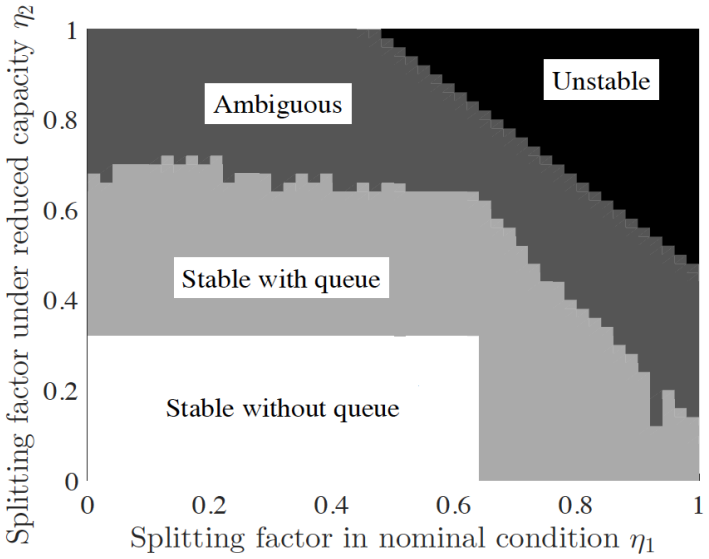

FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Conclusion

* We identified a unique set of domain-level resilience questions and approached them using different but complementary models and algorithms
* We derived new structural insights and design guidelines to improve resilience in infrastructures by applying RC+EI theory to multiple domains
* Future work (3 specific projects):
  * Network security: Information systems, incentives, and insurance
  * Economic foundations of new markets for energy and data: integration of dynamic mechanisms and control theoretic ideas
  * RCPS testbed: mapping resilient monitoring/diagnostics and control algorithms for water distribution, transportation, and power T&D