A Hierarchical Approach to CPS Resilience

based on Game Theory, Stochastic Control, and Theory of Incentives

Demosthenis Teneketzis[1]

(joint with Saurabh Amin[2] and Galina A. Schwartz[3])

[1]University of Michigan, Ann Arbor
[2]Massachusetts Institute of Technology
[3]University of California, Berkeley

FORCES Kickoff Meeting
Washington, D.C., April 12th, 2013

Outline

- Motivation: CPS resilience, security
- Research plan: Three-layer hierarchical approach
  - Upper layer: Game theory
  - Middle layer: Stochastic control & Theory of incentives
  - Lower layer: Control Theory
- Will concentrate on the upper and middle layers

Failures in CPS

- Simultaneous attacks [security failures]
  - Targeted cyber-attacks
  - Non-targeted cyber-attacks
  - Coordinated physical attacks
- Simultaneous faults [reliability failures]
  - Common-mode failures
  - Random failures due to nature
  - Operator errors
- Cascading failures
  - Failure of nodes in one subnet $\Rightarrow$ progressive failures in other subnets

Observation

Due to cyber-physical interactions, it is extremely difficult to distinguish reliability & security failures using *imperfect* diagnostic information.

Salient features of CPSs

CPSs are multi-agent systems, where

- Agents (players) are strategic, utility-maximizing entities
- Incomplete and also asymmetric (private) information is present
- CPSs are subject to security failures and reliability failures
- Defense strategies include both control and IT security tools
- Players face regulatory impositions for ensuring efficiency & safety

A hierarchical approach

The above features, along with the social objectives of resilient CPS operation, motivate a hierarchical approach.

Research plan: Three-layer hierarchical approach

Upper layer

- How the collection of CPS's agents deal with external strategic adversary(-ies)
- Network games that model both security failures and reliability failures

Middle layer

- How strategic agents contribute to CPS efficiency and safety, while protecting their conflicting individual objectives
- Joint stochastic control and incentive-theoretic design, coupled with the outcome of the upper layer game

Lower layer

- Control at each individual agent's site.

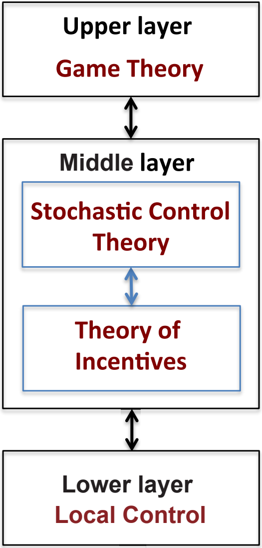

Upper hierarchical layer

Game with security-reliability failures

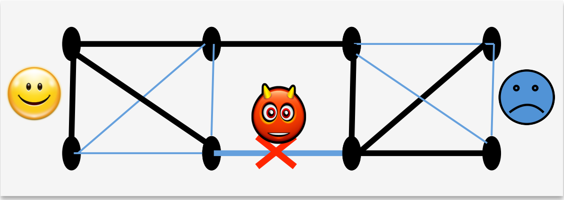

Game played on a graph $\mathscr{G} = (\mathscr{V}, \mathscr{E}, w)$ representing the topological structure of CPS

- Attacker(s)
  - Strategic adversary
  - Nature
- Defender: CPS network designer

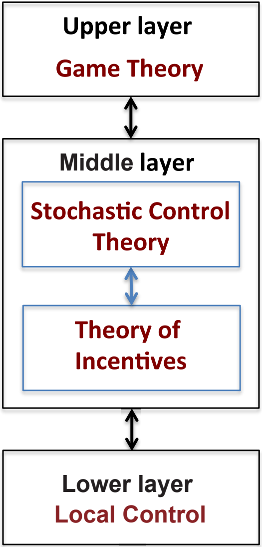

Game with security-reliability failures

Graph $\mathscr{G}$ representing CPS topology

- $\mathscr{V}$: Set of nodes
- $\mathscr{E}$: Set of edges
- w: Set of weights on edges

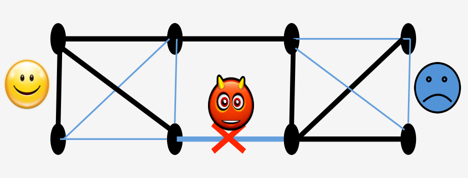

Attacker's strategy space

- $\mathscr{E}$: set of graph's edges
- Attacker chooses an edge $e \in \mathscr{E}$
  - Failure comes from nature with probability π
  - Failure comes from a strategic adversary with probability $(1 - \pi)$

Defender's strategy space

- $\mathscr{T}$: Set of graph's spanning trees
- Defender chooses $\tau \in \mathscr{T}$

Game with security-reliability failures

Payoffs for a choice of $\tau \in \mathscr{T}$ and $e \in \mathscr{E}$

$$\begin{aligned}\Pi_D(\tau, e) =& v(\tau) - (1-\pi)\left[w(e)\mathbf{1}_{\{e \in \tau\}}\right] \\ & - \pi\left[\sum_{e' \in \mathscr{E}} \gamma_{e'} w(e')\mathbf{1}_{\{e' \in \tau\}}\right] \\ \Pi_A(\tau, e) =& w(e)\mathbf{1}_{\{e \in \tau\}}\end{aligned}$$

- $v(\tau)$: value of an operational spanning tree $\tau \in \mathscr{T}$
- $w(e)$: Weight/importance of edge $e \in \mathscr{E}$
- $\mathbf{1}_{\{e' \in \tau\}}$: Indicator function of the even $\{e \in \tau\}$
- $\gamma_{e'}$: Probability of reliability failure of $e' \in \mathscr{E}$

Upper hierarchical layer - Game Theory

Assumptions

- Imperfect information: defender faces aggregate failure probabilities:

$$P(f_e) = \underbrace{\pi\gamma_e}_{\text{reliability}} + \underbrace{(1-\pi)\beta_e}_{\text{security}}, \quad \forall e \in \mathscr{E},$$

  - Given failure probabilities due to nature: $\gamma = (\gamma_{e_1}, \ldots, \gamma_{e_m})$
  - Equilibrium failure probabilities due to attacker: $\beta = (\beta_{e_1}, \ldots, \beta_{e_m})$
- Common knowledge: Payoff functions Π_A and Π_D

Objectives

- Determine Nash equilibria (NE) of the one-stage game within the class of mixed strategies
- Determine equilibria for the finitely or infinitely repeated game

References

A. Washburn, K. Wood (1995)
Two-Person Zero-Sum Games for Network Interdiction.
Operations Research, 43(2), 243–251.

G. A. Schwartz, S. Amin, A. Gueye, J. Walrand (2011)
Network design game with both reliability and security failures.
49th Annual Allerton Conference, 675 – 681.

J. Salmeron, K. Wood, R. Baldick (2004)
Analysis of electric grid security under terrorist threat.
IEEE Transactions on Power Systems, 19, 905–912.

A. Schrijver (2003)
Combinatorial optimization: polyhedra and efficiency.
Springer-Verlag.

Upper layer → Middle layer

How to embed the outcomes of upper layer into the middle layer failure models for the design of resilient CPS strategies using stochastic control and incentive-theoretic formulations?

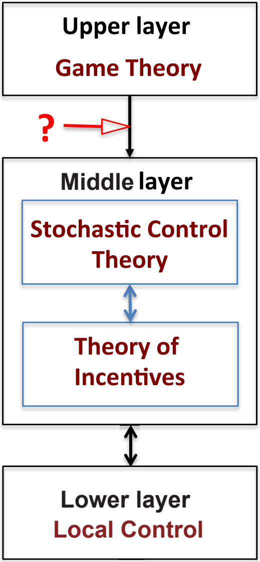

Upper layer → Middle layer

Outcome of upper layer game

- Equilibrium strategies for attacker and defender (α, β)
- Edge failure probabilities: $P(f_e) = \pi\gamma_e + (1-\pi)\beta_e, \quad \forall e \in \mathscr{E}$

Embedding $P(f_e)$ into middle layer model

- Physical: structural failures
- Cyber: sensor-actuator failures

Middle hierarchical layer

Resulting failure models are used to design of resilient strategies.

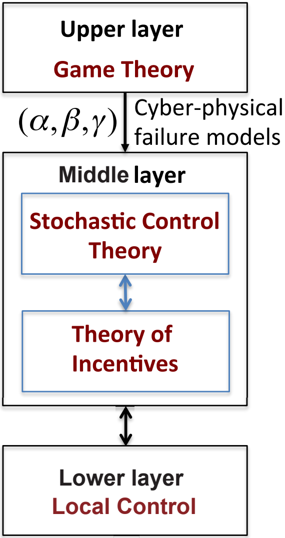

Middle hierarchical layer

Stochastic control and incentives

- Stochastic control: Performance benchmark against CPS failures
- Theory of Incentives: implement in appropriate equilibria the optimal control strategies of the stochastic control problem

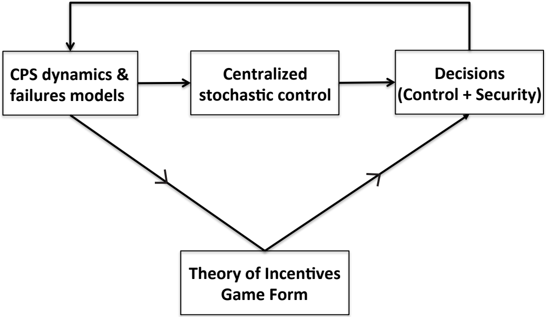

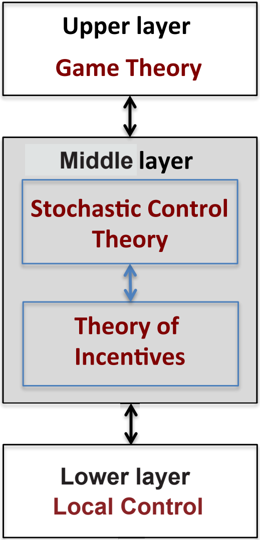

Middle hierarchical layer - CPS model

Agent i's dynamics modeled by

- A controlled stochastic vector difference equation
- A controlled multi-dimensional Markov chain

$$X_{t+1}^{i} = f_t^i \left(X_t^i, U_t^1, \ldots, U_t^N, W_t^i, P_{s,t}^i, P_{u,t}^i \right)$$

- N: $\#$ of agents in CPS
- $\mathscr{N} = \{1, \ldots, N\}$: set of agents
- $X_i^t \in \mathscr{X}_i$: state of agent i at time t and $\mathscr{X}_i$ finite
- U_t^i: control action of agent i
- W_t^i: noise in component i at t
- $P_{s,t}^i$ and $P_{u,t}^i$: probabilities of structural failure & actuation failure at t

Middle hierarchical layer - CPS model

State of CPS with N agents at time t

$$\underline{\mathrm{X}}_t = \left(X_t^1, X_t^2, \ldots, X_t^N\right)$$

Sensing model

$$Y_t^i = h_t^i\left(X_t^i, W_t^{o,i}, P_{i,t}^o\right), i \in \mathcal{N}$$

- Y_t^i: observation of agent i at t
- $W_t^{o,i}$: observation noise of i at t
- $P_{i,t}^o$: probability of sensing failure at t

Middle hierarchical layer - CPS model

Decision strategies

$$U_t^i = g_t^i \left(Y_{1:t}^1, Y_{1:t}^2, \ldots, Y_{1:t}^N, U_{1:t-1}^1, U_{1:t-1}^2, \ldots, U_{1:t-1}^N \right), i \in \mathcal{N}, t = 1, \ldots, T$$

- T: time horizon (finite or infinite)
- $Y_{1:t}^i = \left(Y_1^i, Y_2^i, \ldots, Y_t^i\right)$
- $U_{1:t-1}^i = \left(U_1^i, U_2^i, \ldots, U_{t-1}^i\right)$
- $g^i = \left(g_1^i, g_2^i, \ldots, g_T^i\right)$: control/decision strategy of agent i
- $g = \left(g^1, g^2, \ldots, g^N\right)$: control strategy for the CPS

Middle hierarchical layer - CPS model

Reward Functions

- Reward function for agent i

$$R^i = \sum_{t=1}^{T} R_t^i\left(X_t^i, U_t^1, U_t^2, \ldots, U_t^N\right)$$

- Total reward

$$R = \sum_{i=1}^{N}\sum_{t=1}^{T} R_t^i\left(X_t^i, U_t^1, U_t^2, \ldots, U_t^N\right)$$

CPS model: An example

- CPS system - system consisting of N energy suppliers
  - Each supplier is strategic (selfish, self-utility optimizer)
  - Each supplier has private information (e.g. production technology)
  - Efficient operation so as to achieve a *social objective*
- X_t^i: energy producing capability of power supplier i at t
- U_t^i: energy produced by power suppliers i at t
- $X_{t+1}^i = f_t^i\left(X_t^i, U_t^i, W_t^i, P_{s,t}^i, P_{u,t}^i\right)$, i.e., X_{t+1}^i depends on X_t^i, U_t^i, failures due to nature, failures due to strategic adversary, repairs.
- Profit of power supplier i at time t

$$R_t^i(X_t^i, U_t^1, U_t^2, \ldots, U_t^N) = \lambda_t\left(U_t^1, U_t^2, \ldots, U_t^N\right) \cdot U_t^i - \hat{c}_t^i\left(X_t^i\right) \cdot U_t^i$$

  - $\lambda_t\left(U_t^1, U_t^2, \ldots, U_t^N\right)$: price charged per unit of produced energy
  - $\hat{c}_t^i\left(X_t^i\right)$: cost per unit of energy produced when state is X_t^i.

Middle hierarchical layer - Objectives

Determine $g = (g^1, g^2, \ldots, g^N)$ to maximize $E^g[R]$, subject to

- Informational constraints (agent i's information at t is $(Y^i_{1:t}, U^i_{1:t})$)
- Taking strategy behavior into account

To achieve the objective

- Derive performance benchmark using stochastic control
- Achieve performance benchmark by a mechanism/game form which satisfies the problem's constraints using the theory of incentives

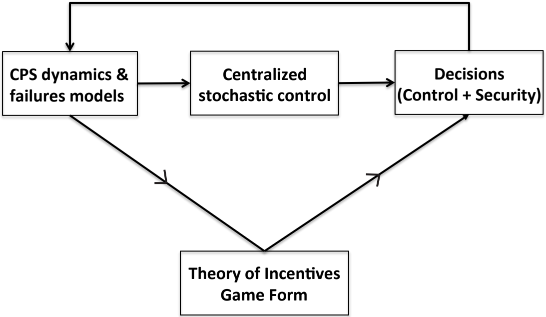

Middle hierarchical Layer - Stochastic control

Consider a central authority that has all the information, including

- Agents' utilities/reward functions
- Observations & control actions, i.e. $\mathscr{I}_t = \left(Y^1_{1:t}, \ldots, Y^N_{1:t}, U^1_{1:t}, \ldots, U^N_{1:t}\right)$
- CPS dynamics

Stochastic control problem

- Central authority chooses $g = \left(g^1, g^2, \ldots, g^N\right)$ to maximize $E^g[R]$ subject to
  - Sensor-actuator failures
  - Structural failures
- Solution provides a performance benchmark
- Achievable if all agents were willing to cooperate & share information
- However, CPS agents are strategic, selfish!

References

L. Schenato, M. Franceschetti, K. Poolla, S. Sastry (2007)
Foundations of Control and Estimation over Lossy Networks.
Proceedings of the IEEE, 95(1), 163 - 187.

S. Amin, G.A,. Schwartz, S. Sastry (2007)
Security of interdependent and identical networked control systems.
Automatica, 49(1), 186-192.

P. R. Kumar, P. Varaiya (1986)
Stochastic systems: estimation, identification and adaptive control.
Prentice-Hall.

Middle hierarchical layer - Achieving the benchmark

Theory of incentives / Mechanism design

- $\mathbf{E}$ environment space (space of agents' utilities, network topologies)
- $\mathbf{U}$ Action / alloc. / control space
- $(\mathbf{M}, h)$: game form/mechanism
  - $\mathbf{M}$: message / strategy space
  - h: outcome function
- μ: message correspondence
- $\forall \mathbf{e} \in \mathbf{E}$, $(\mathbf{M}, h, \mathbf{e})$ is the game induced by $(\mathbf{M}, h)$

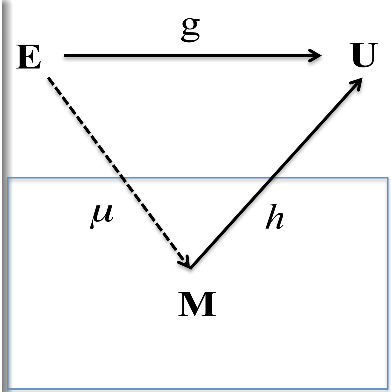

Middle hierarchical layer - Incentives

Let $\mathbf{M}^*(\mathbf{e}) = \{\mathbf{m}^* \in \mathbf{M} : \mathbf{m}^* \text{ is an equilibrium message/strategy of } (\mathbf{M}, h, \mathbf{e})\}$

Objective: Design (M, h) so that

$$\forall e \in E, \forall m^* \in M^*(e), h(m^*)(e) = g(e)$$

That is, design a game form/mechanism (if exists) that accounts for the

- Information structure of the CPS,
- Agents' strategic behavior
- Achieves the same performance as the performance benchmark (i.e., the solution of the stochastic control problem)

Incentives: Achieving the upper bound

Approach

- Restrict attention to direct revelation mechanisms invoking the revelation principle.
- Revelation principle: If a game form $(\mathbf{M}, h)$ implements $g : \mathbf{E} \to \mathbf{U}$ in a certain equilibrium concept $\hat{\Lambda}$ (e.g. BNE), then there is a direct revelation mechanism $(\mathbf{E}, h^*)$ which has the following property:

  Reporting one's true environment $\mathbf{e}$ is an equilibrium message/strategy of $(\mathbf{E}, h^*, \mathbf{e})$ in the same equilibrium concept $\hat{\Lambda}$, and $h^*(\mathbf{e}) \in g(\mathbf{e})$ for all $\mathbf{e} \in \mathbf{E}$.
- We are looking for truthful implementation of g (optimal control strategy for the stochastic control problem).
- We consider agents $i \in \mathcal{N} = \{1, 2, \ldots, N\}$ with quasi-linear utilities

$$\mathbf{V}_t^i(X_t^i, U_t^1, \ldots, U_t^N, (tx)_t^i) = R_t^i(X_t^i, U_t^1, \ldots, U_t^N) - (tx)_t^i$$

Dynamic Incentives: Achieving the upper bound

Determine a dynamic direct revelation mechanism $(\mathbf{E}, h_1, h_2, \ldots, h_T)$ [if it exists] that has the following properties:

(i) It is incentive compatible (i.e., truth telling is a BNE of the game induced by the mechanism)

(ii) It is budget-balanced

$$\sum_{i=1}^{N} (tx)_t^i = 0 \quad \forall t \text{ OR } \sum_{t=1}^{T} \sum_{i=1}^{N} (tx)_t^i \quad = 0 \quad \text{at truthful equilibrium}$$

(iii) Decisions/control actions at truthful equilibrium are the same as the decisions made by g (the optimal control law).

References

S. Athey and I. Segal (2012)
Optimal Collusion with Private Information.
RAND Journal of Economics.

D. Bergemann and J. Valimaki (2010)
The Dynamic Pivot Mechanism.
Econometrica.

J. Escobar and J. Toikka (2012)
Efficiency in Games with Markovian Private Information.
working paper, MIT.

A. Pavan , I. Segal, and J. Toikka (2012)
Dynamic Mechanism Design.
working paper, Stanford University.

P. Courty and H. Li (2000)
Sequential Screening.
Review of Economic Studies, 67(4), 697 - 717.

M. Battaglini (2003)
Optimality and Renegotiation in Dynamic Contracting.
Mimeo, Princeton University.

M. Battaglini (2003)
Long-Term Contracting with Markovian Consumers.
American Economic Review, 95(3), 637 - 658.

A. Atkenson and R. Lucas (1993)
On Efficient Distribution with Private Information.
Review of Economic Studies, 59, 427 - 453.

D. Fudenberg, D. Levine, and E. Maskin (1994)
The Folk Theorem with Imperfect Public Information.
Econometrica, 62, 997 - 1039.

S. Athey and K. Bagwell (2008)
Collusion with Persistent Cost Shocks.
Econometrica, 76(3), 493 - 540.

S. Athey and D. Miller (2007)
Efficiency in Repeated Trade with Hidden Valuations.
Theoretical Economics.

D. Miller (2004)
The Dynamic Cost of Ex-Post Incentive Compatibility in Repeated Games of Private Information.
Mimeo, UCSD.