Game Theoretic Analysis of Security and Sustainability

Bo An

boan@ntu.edu.sg

School of Computer Science and Engineering

Nanyang Technological University

August 22, 2017@IJCAI'17, Early Career Spotlight

NOTE: 1) Some slides might be messy as I used animation in the ppt slides.

2) You can read the accompanying paper for more details: http://www.ntu.edu.sg/home/boan/papers/IJCAI17_Spotlight.pdf

AMI *Agent Mediated Intelligence* Research Group

Game Theory for Security in the Field

Security Games

- Security allocation: (i) Target weights; (ii) Opponent reaction
- Stackelberg game: Security forces commit first
- Strong Stackelberg Equilibrium
- Our contributions:
  - *Algorithms for solving large scale games*
  - *Learning adversary behavior*
  - *Applications in the real world*

Attacker

Defender

	Attacker Action #1	Attacker Action #2	……	……
Defender Action #1	7, -3	-6, 2	……	……
Defender Action #2	-3, 4	4, -6	……	……
……	……	……	……	……
……	……	……	……	……

Analyzing Complex Security Games

Dynamic Payoff

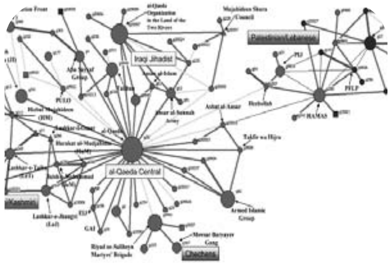

Network Games

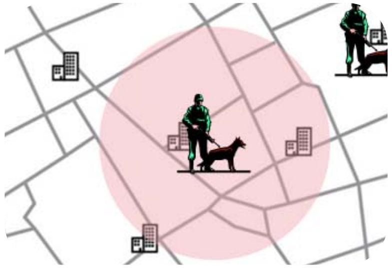

Protection Externality

Uncertainty

Strategic Deception

Cyber Security

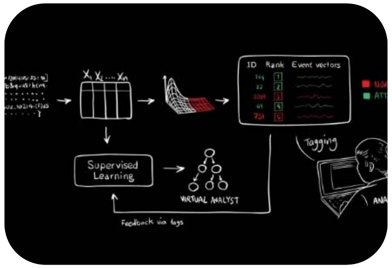

Adversarial Machine Learning

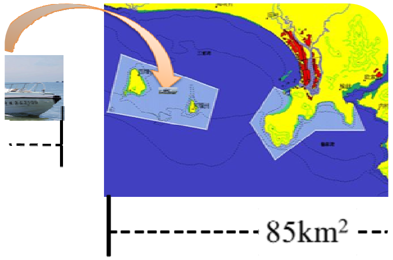

Coral Reef

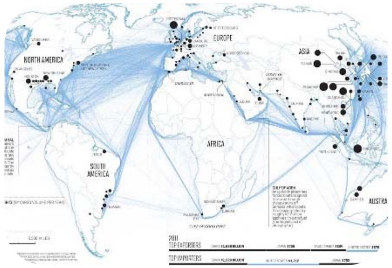

Nuclear Smuggling

Elections

- Combining techniques from AI, Game Theory, Operations Research …
- Marry theory with practice
- Approaches can be applied to other domains
  - *With proper tuning & extension*

SGs with Dynamic Payoffs [AAAI'14]

- Target value changes over time
  - *Utility of attacking a target decreases with # of protecting resources*
- Idea: Dynam
  - *Transfer reso*
  - *A resource in*

Boston Marathon Bombings

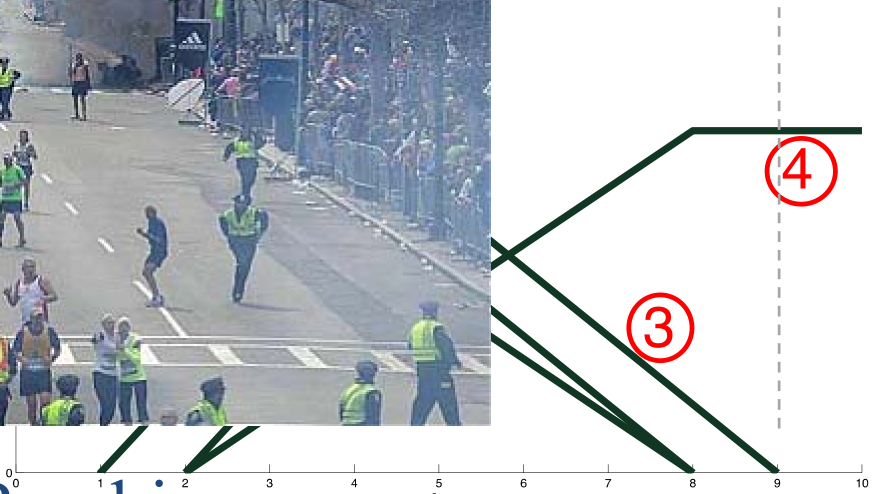

Varying target value $v_i(t)$

SCOUT-A: Negligible Transfer Time

- Context: Resources can be transferred quickly
- Find the minimax assignment of resources at each time point
  - *Infeasible to find the minimax assignment at each time point since time is continuous*
- 'Minimax assignment' does not change continuously

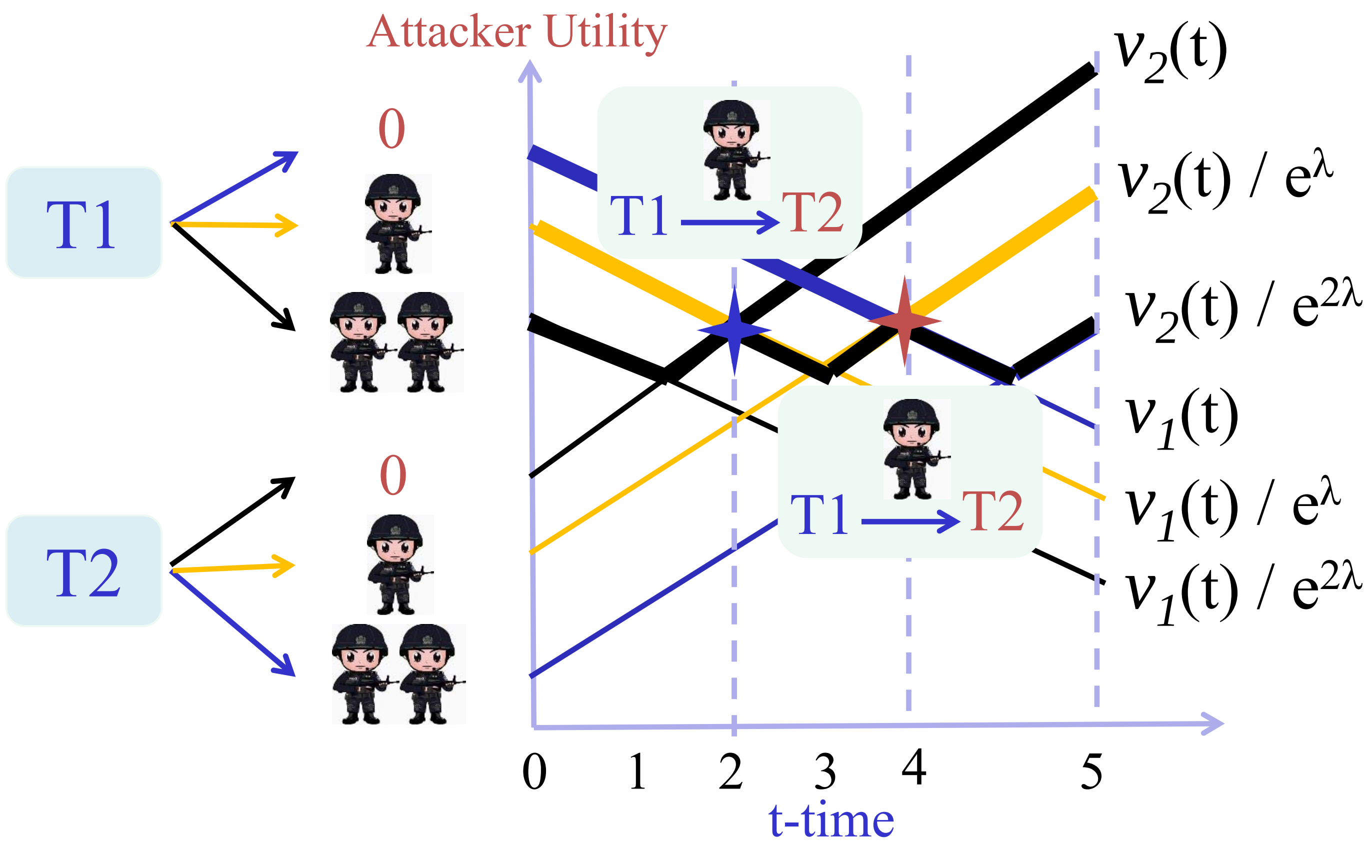

SCOUT-A computes the time point at which a minimax assignment 'expires', then finds the 'next' minimax assignment.

SCOUT-C: Nonzero Transfer Time

Theorem:
For any game with *continuous* defender strategy space, we can construct an equivalent game with *discrete* defender strategy space.

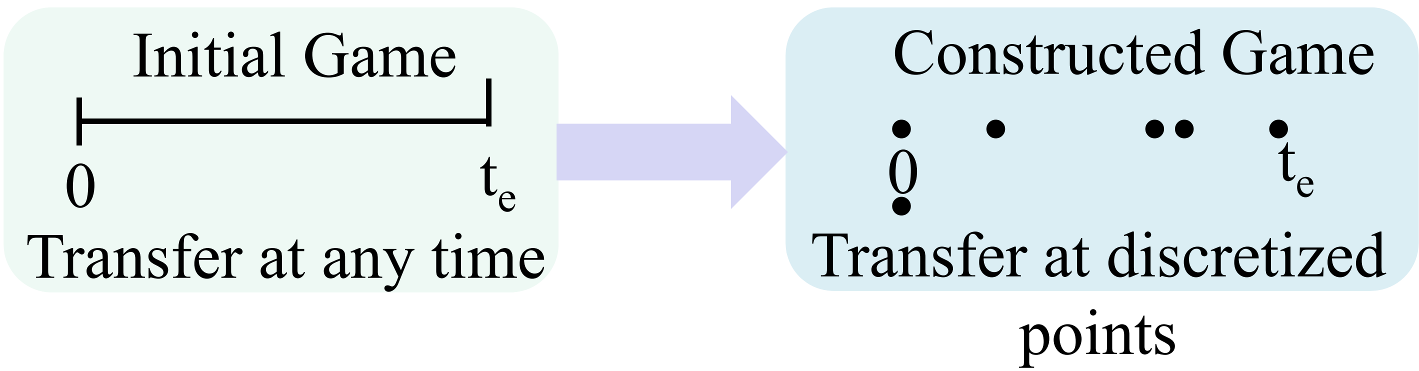

The equilibrium of the constructed game is also an equilibrium of the initial game.

Mixed Strategy EQ for Dynamic Payoff Games [IJCAI'15]

- Dynamic Payoffs: Targets' importance changes

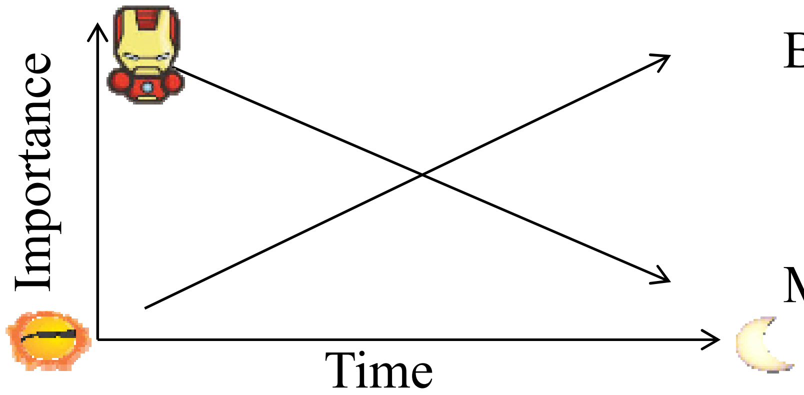

- Mixed strategy
  - *Distribution over pure strategies (Infinite & Continuous support set)*
- Mixed strategy ➡ Compact representation
  - *Coverage functions* $c_i(t)$*: Probability of protecting target i at time t*

Pure strategy 1 (0.5)	Target 1: 6:00 – 12:00 Target 2: 12:00 – 18:00
Pure strategy 2 (0.5)	Target 1: 6:00 – 10:00 Target 2: 10:00 – 18:00

$$\Rightarrow c_1(t) = \begin{cases} 1, t \in [6:00, 10:00) \\ 0.5, t \in [10:00, 12:00) \\ 0, t \in [12:00, 18:00] \end{cases}$$

Theorem:

Coverage functions can be implemented by sampling pure strategies with finite transfers

ADEN: Nonzero Transfer Time

- A resource may be in transfer at time t
- Approach: Discretize game by an $(\varepsilon,\delta)-Mesh$

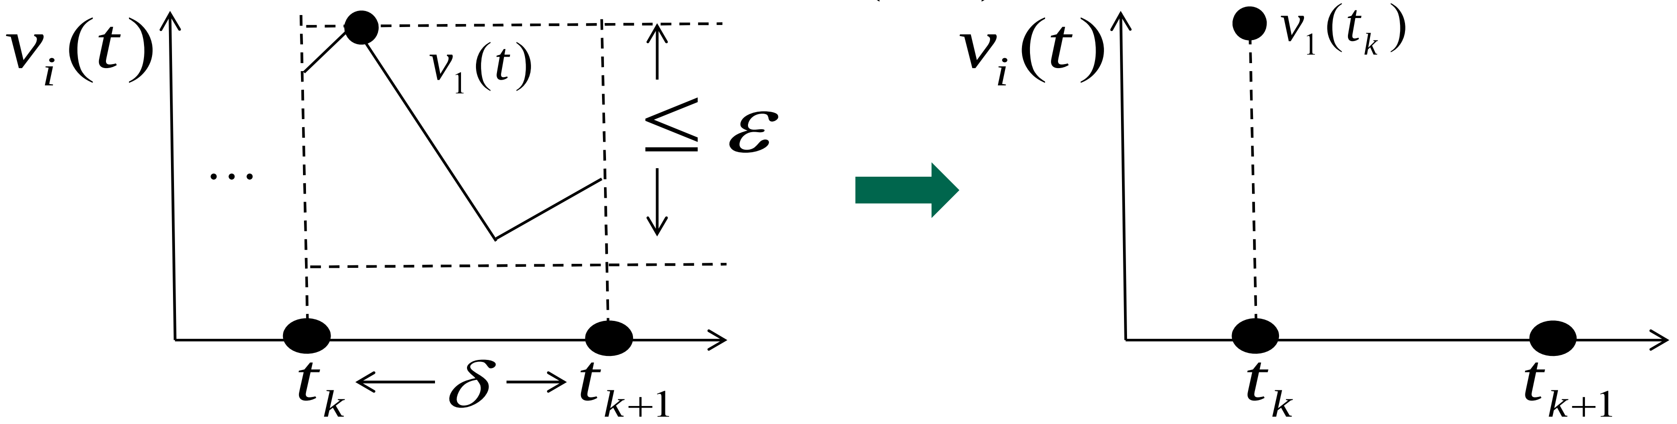

Theorem: Discretizing game by $(\varepsilon,\delta)-Mesh$ leads to a loss of at most ε

Weight of point (k,i)

Value of target i at time t_k in the bridge game $v'_i(t_k)$

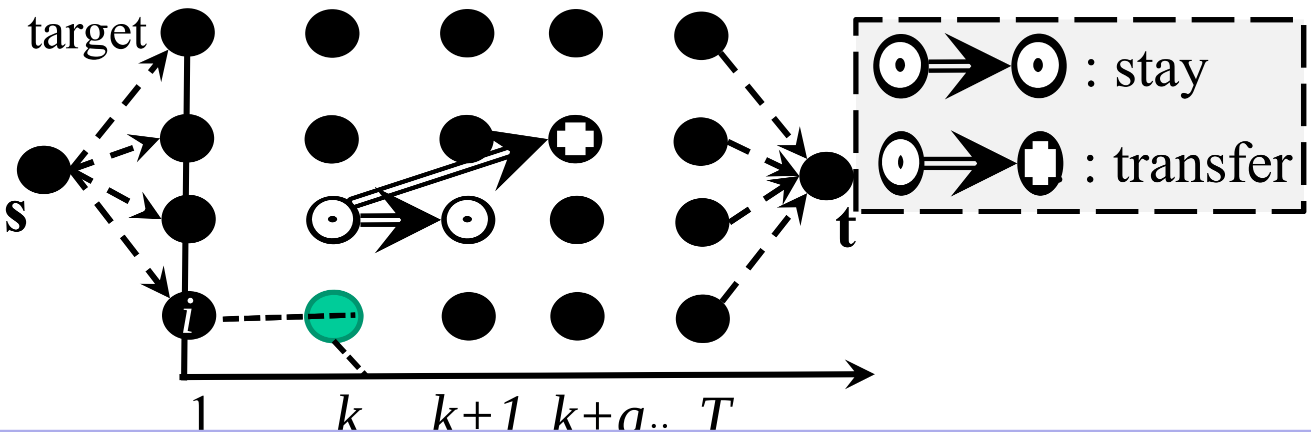

Theorem: A mixed strategy ⟷ An m-unit fractional flow

Complexity: $O(\frac{t_e}{\delta}n^2)$ (FPTAS for Lipschitz continuous value functions)

Network Games: Detecting Terrorist Plots [AAAI'16a]

- **Paris Shootings on January 7, 2015**
  - *Two Kouachi brothers stormed into the Paris office of Charlie Hebdo and gunned down 12 people*
  - *A few hours later, Amedy Coulibaly killed a policewoman in Montrouge and four hostages at a kosher supermarket in east Paris*
  - *A coordinated plot planned by al-Qaeda in the Arabian Peninsula (AQAP)*
- **Monitoring potential terrorists!**
- **Terrorist planners (e.g., ISIS):**
  - *Arouse a connected subgroup terrorist network*
  - *Conduct surveillance*
- **Security agencies (e.g., DGSE)**
  - *Decide how to allocate limited security resources to monitor the terrorists*

Chérif Kouachi

Saïd Kouachi

Charlie Hebdo

Amedy Coulibaly

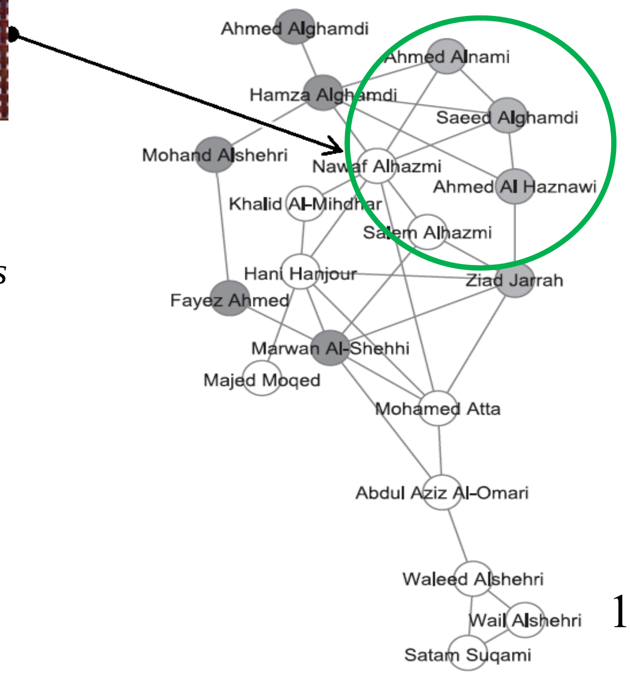

Computing Optimal Monitoring Strategy

- Utility function:
  - *If overlap, defender wins:* $U_a = U_d = 0$
  - *If not, attacker succeeds:* $U_a = P(A_{42}), U_d = -P(A_{42})$
  - *Attacker utility of choosing subgraph A:*

$$P(A) = \sum_{v\in A}(\tau_v + \delta \sum_{u \in N_A v} \tau_u)$$

neighborhood of v in subgraph A

the extent of positive network externality

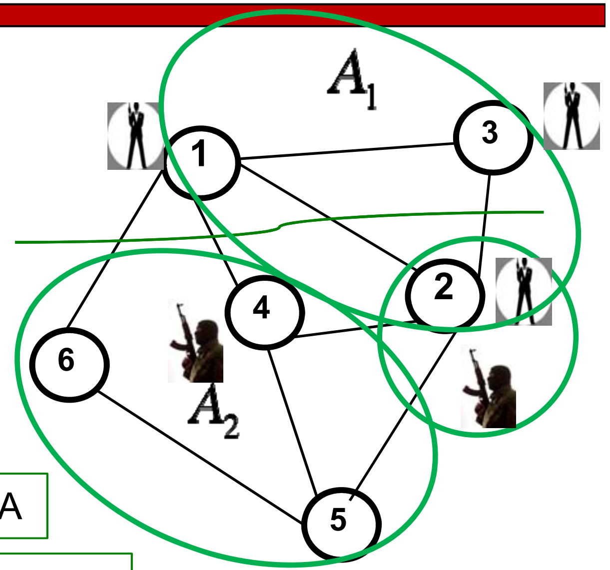

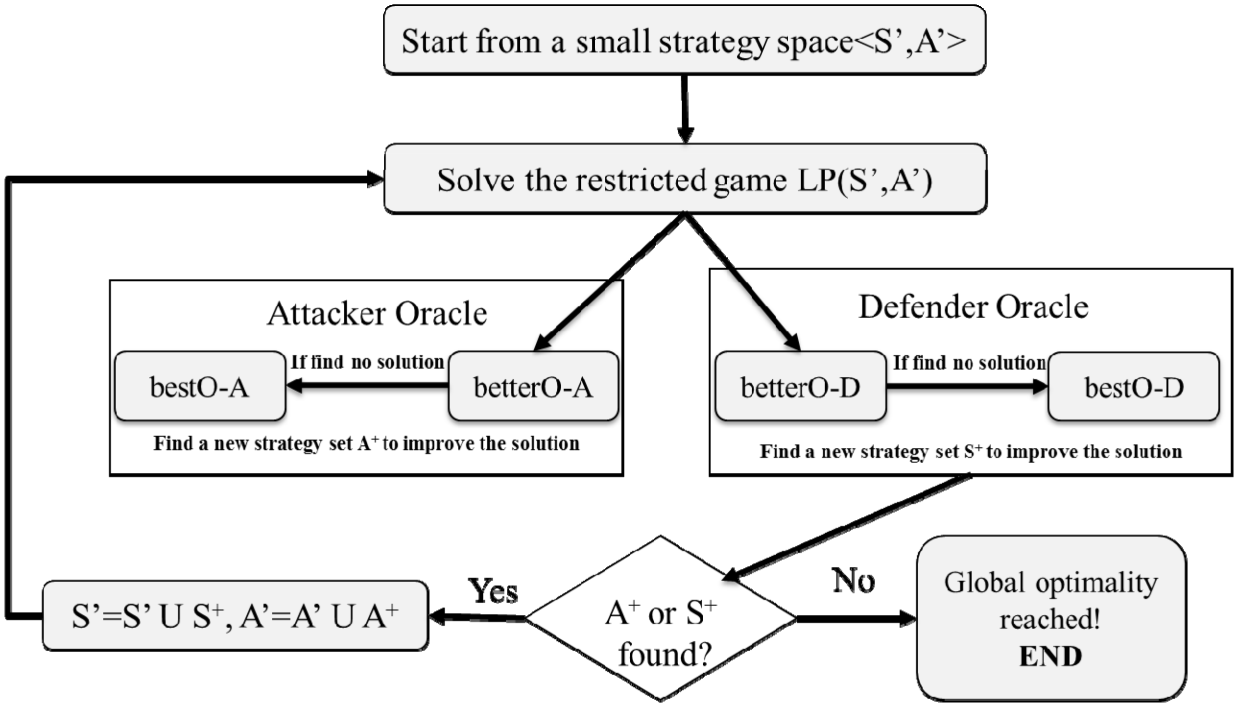

Theorems:

- BestO-D/bestO-A is NP-hard;
- BetterO-A/betterO-D guarantees a (1-1/e)-approximation ratio

Network Games: Coalitional Security Games [AAMAS'16]

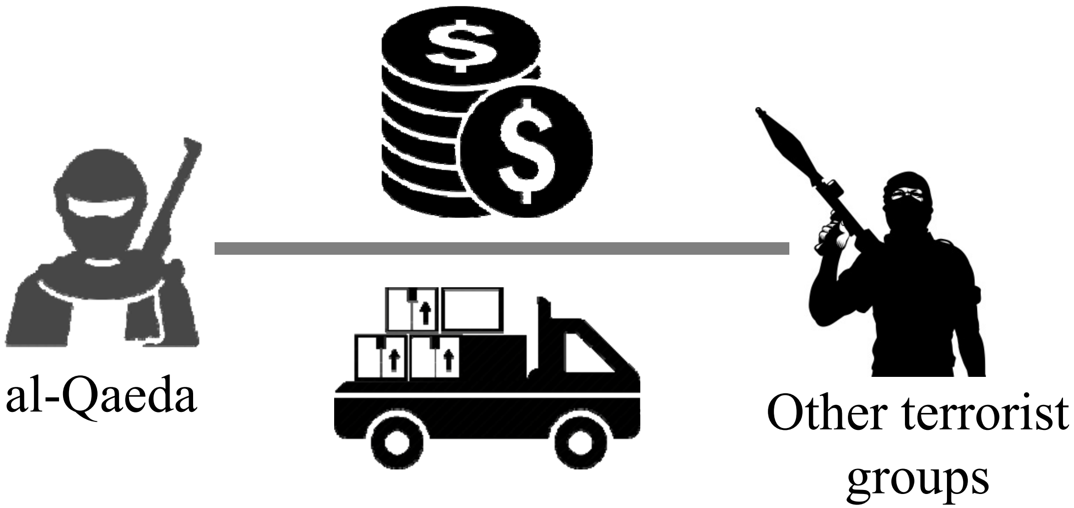

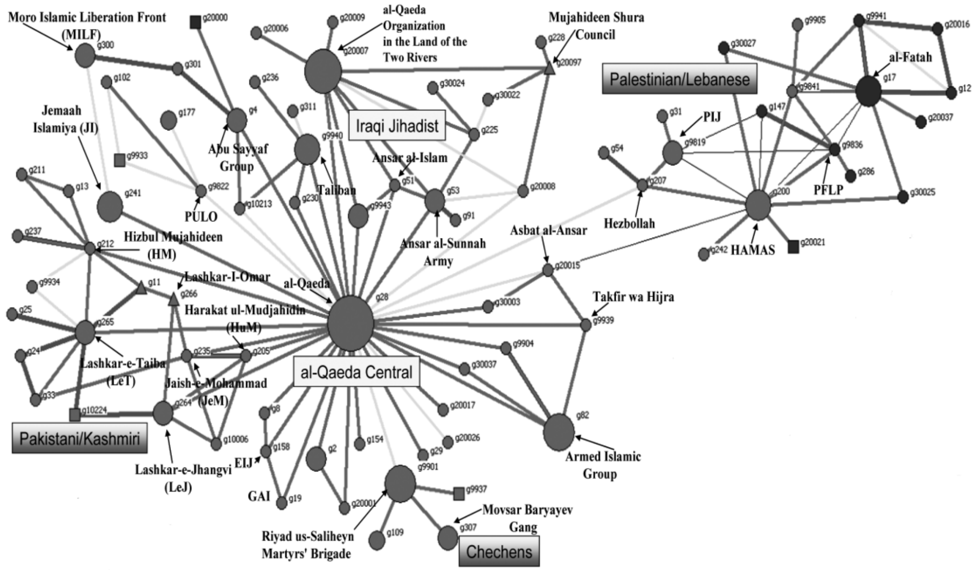

"The ***interconnected*** *nature of terrorist organizations necessitates that we pursue them across the geographic spectrum to ensure that all* ***linkages*** *between the strong and the weak organizations are* ***broken****, leaving each of them* ***isolated****,* ***exposed****, and* ***vulnerable*** *to defeat."*

-CIA

National Strategy for Combating Terrorism

CSG: Strategies, Utility and Equilibrium

- **Strategies**
  - *Defender: blocking a set of edges*
  - *Attacker: playing coalitional game and forming a coalitional structure*

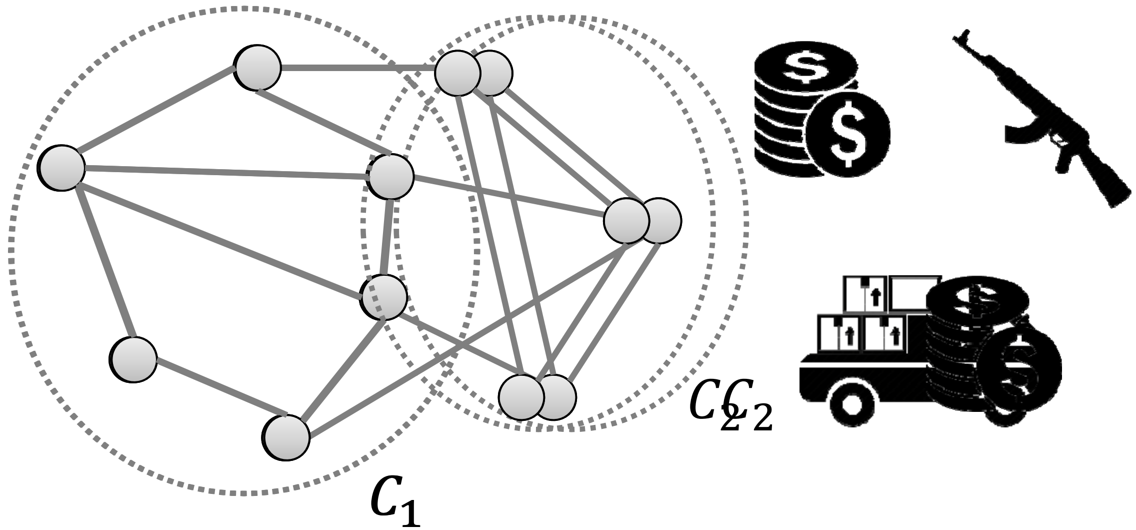

blocking costs: 10

- **Utility**
  - *Coalition value: capability of attacking targets (knapsack problem)*

Value: 10

$$v(C_1) = 20$$
$$v(C_2) = 10$$

Value: 20

CSG: Branch and Price Algorithm

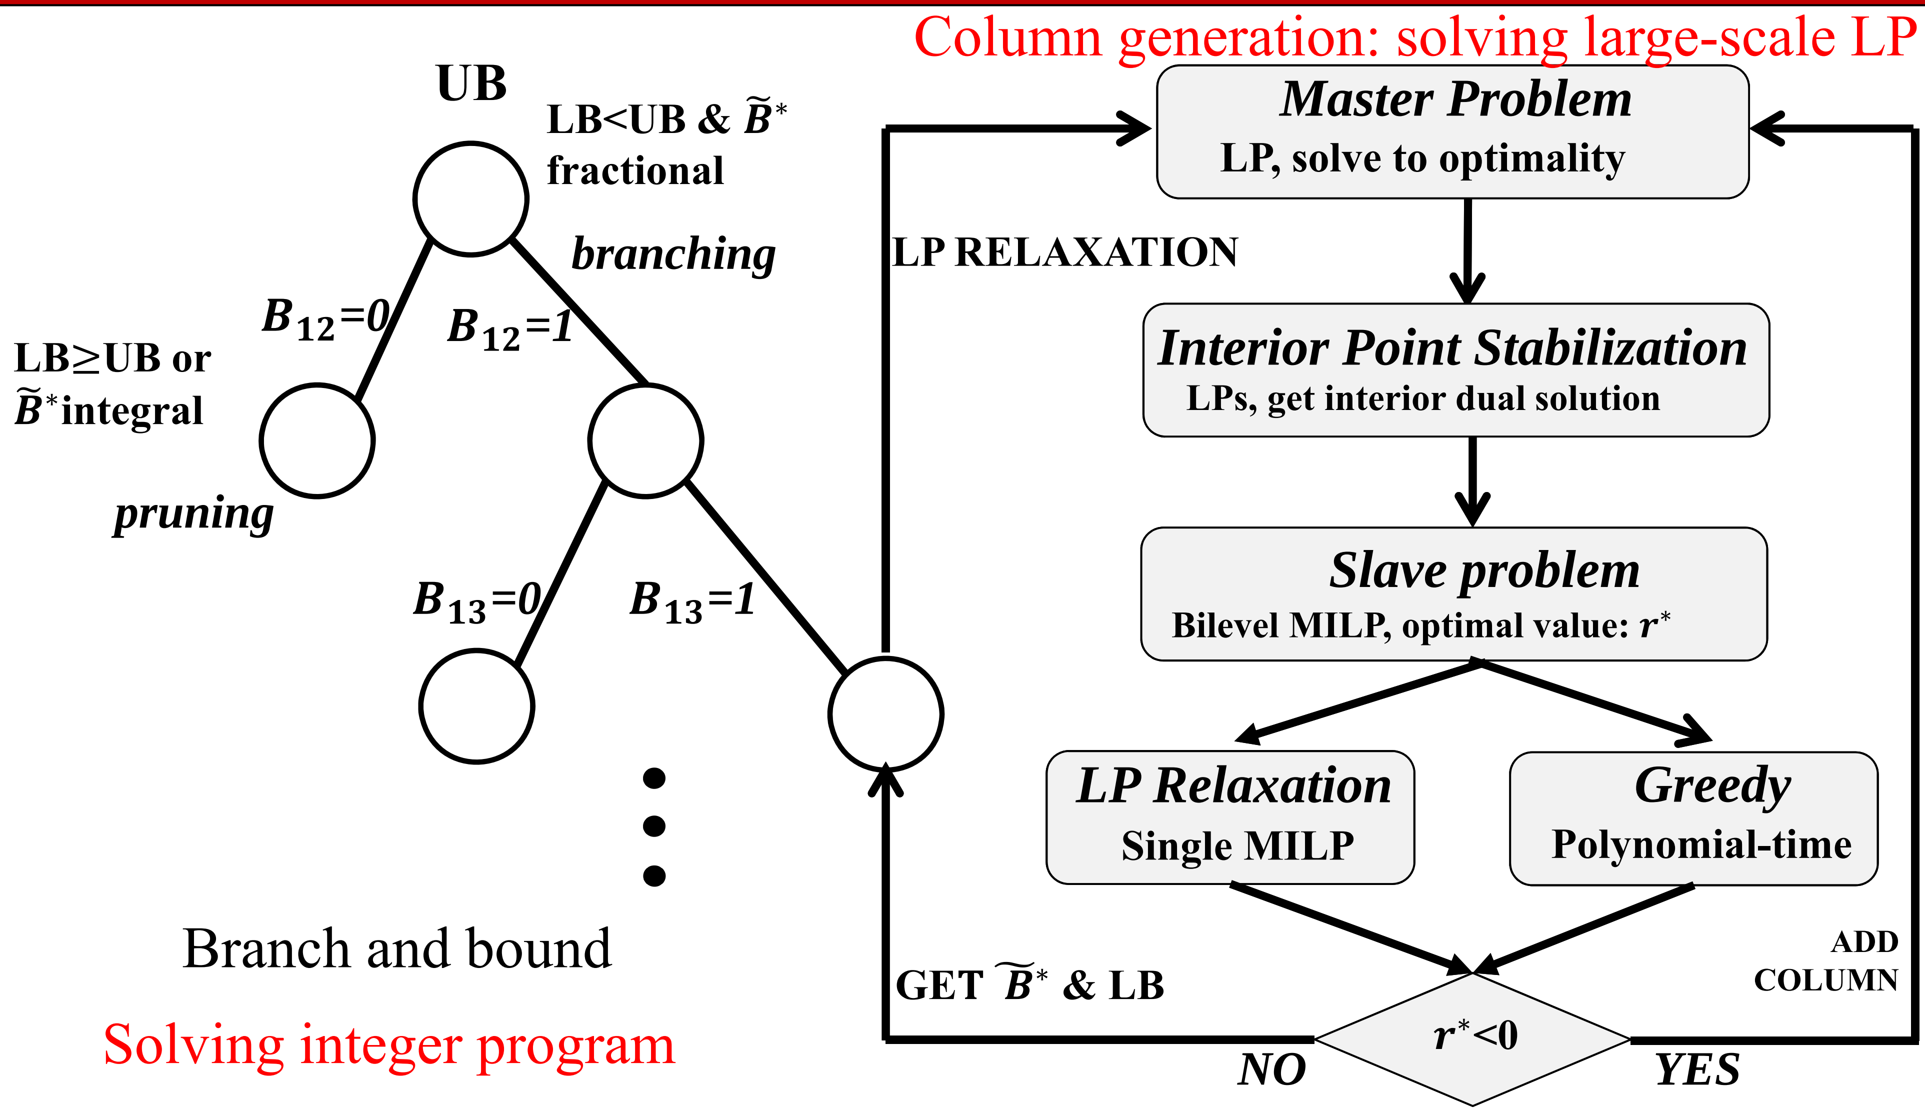

Theorems:

- CSG problem is MAX SNP-hard
- The algorithm achieves constant factor approximation

Security Games with Protection Externalities [AAAI'15]

- **Protection Externalities (PE)**
  - *One resource protects multiple targets simultaneously*
  - *Many real-world scenarios: ferry ship…*

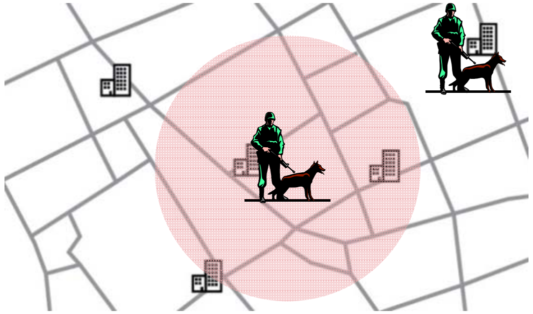

- **NP-hard** to compute the equilibrium
  - *Reduction from Set Cover*

- A **Column Generation** based solution algorithm:
  - ① *N target-defined LPs (t-LP)*
  - ② *Column Generation for t-LP*
    - An MILP formulation for the slave procedure
    - A constant-factor greedy approach to speed up
  - ③ *An upper bound LP for pruning (u-LP)*

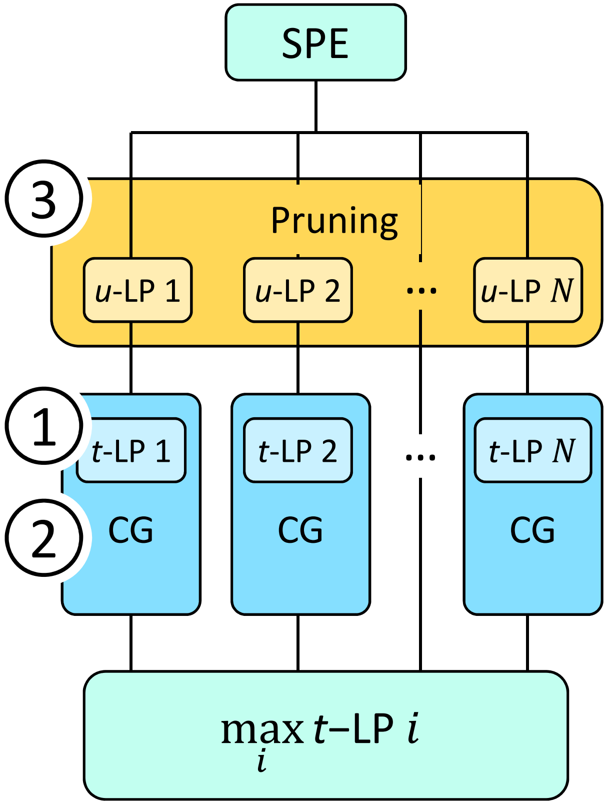

SPE on a Plane [AAAI'17a]

- **A more realistic setting**
  - A planar topology
  - Resource allocation in a continuous space (not restricted to targets)

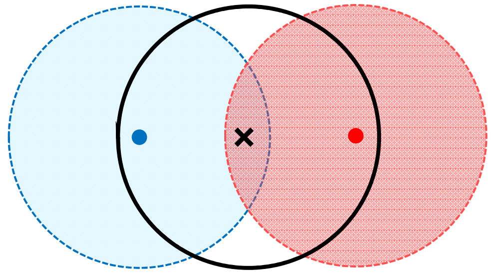

- **Easier? Harder?**
  - NP-hard ☹

    Reduction from *Euclidian Disc Cover*: if n given points on a plane can be covered by m identical disks.

  - Inapproximable: no PTAS unless P=NP ☹
  - Approximable for zero-sum games: a PTAS ☺

A PTAS based on Grid Shifting

- Divide the plane with a grid $\mathcal{G}_1$
- Shift $\mathcal{G}_1$ to obtain grids $\mathcal{G}_2, \dots, \mathcal{G}_l$
- A new pure strategy space $\tilde{S}$ in which each pure strategy fits into one grid (not overlapping grid lines)
- Compute $\tilde{x}^*$, the optimal defender strategy under $\tilde{S}$
- (Polynomial time computable: ellipsoid method + Dynamic programming for the separation oracle)
- $U(\tilde{x}^*)$ is a $\left(1 - \frac{2}{l}\right)$-approximation to the optimal solution under the original strategy space

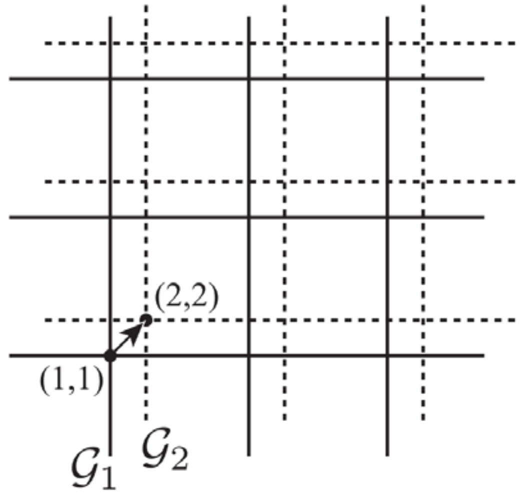

None zero-sum games

- Inapproximable in general but solutions with guaranteed quality can be obtained efficiently if:
  - The game is quasi-zero-sum (as is in most real scenarios)
  - Solutions are restricted to robust ones

Network Flow Interdiction Games [IJCAI'16]

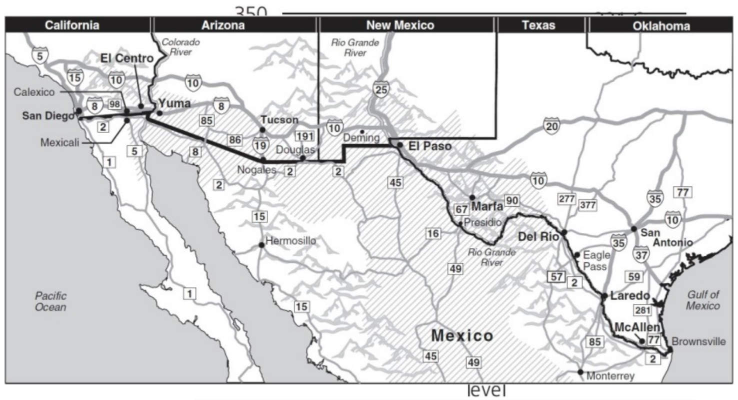

- **Drug Smuggling**
  - *Illegal drug trafficking is a world wide issue*
  - *Checkpoints are operated to prevent drug trafficking*

- **Challenges**
  - *Attacker strategy is a flow*
  - *No compact representation of attacker strate*
  - *Double oracle cannot be applied*

- **Solutions**
  - *Network flow interdiction game model*
  - *Column and constraint generation algorithm*

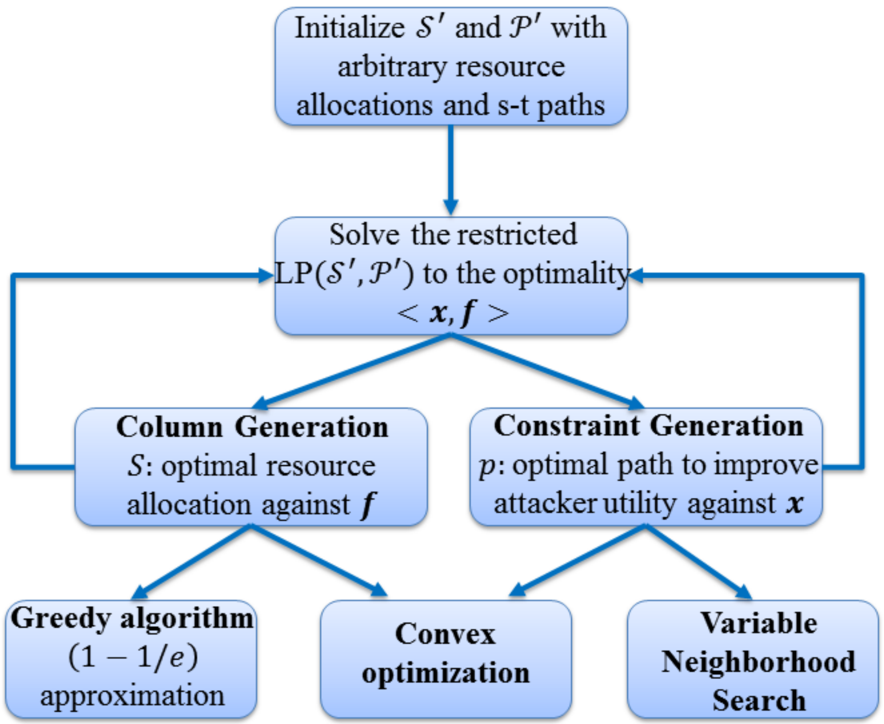

Repeated Network Interdiction Problems [IJCAI'17a]

- **Limitation of Existing Models**
  - *Human adversaries are not fully rational*
  - *Defender has few prior knowledge of adversary*
  - *Interactions between agents are frequent*
- **Repeated Network Interdiction Game**
- **Online Learning Framework**

Objective: designing an *online* policy to minimize the defender's *regret*:

$$R_T = \max_{S \in \mathcal{S}} \sum_{t=1}^{T} U_d(S, \boldsymbol{f}^t) - \mathbb{E}\Big[\sum_{t=1}^{T} U_d(S^t, \boldsymbol{f}^t)\Big]$$

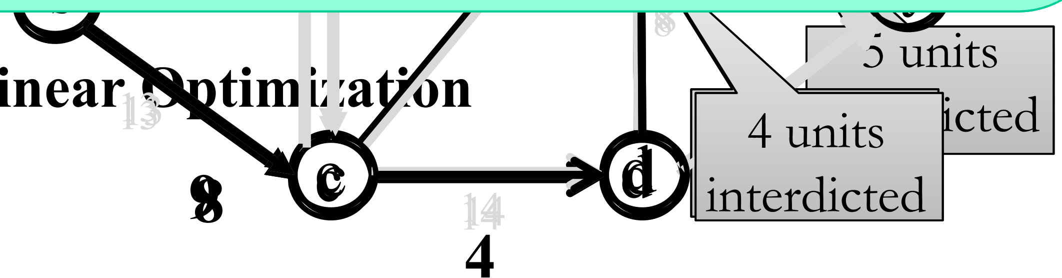

- **Transformation to Online Linear Optimization**

SBGA

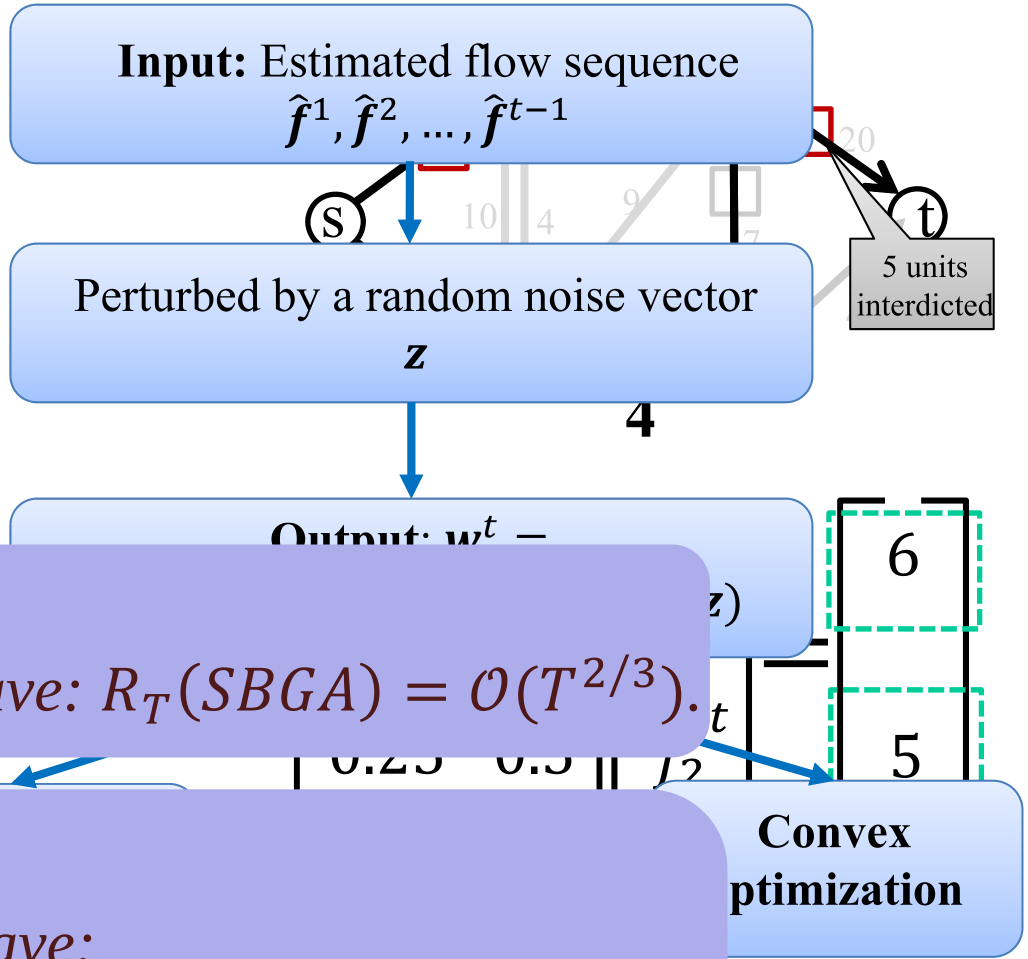

- **Exploration**
  - *Unbiased estimator*
- **Exploitation**
  - *Follow the perturbed leader (FPL)*
- **Theoretical Guarantees**

Theorem:
With proper learning parameters, we have: $R_T(SBGA) = \mathcal{O}(T^{2/3})$.

Theorem:
With proper learning parameters, we have:
$\delta \cdot OPT - Reward(SBGA) \leq \mathcal{O}\left(T^{\frac{2}{3}}\right)$, *where* OPT *is the utility of the optimal adaptive defender policy and* δ *is a constant depending on the interdiction probabilities.*

Strategic Secrecy in Security Games [IJCAI'17b]

- **Dilemma of Plainclothes**
  - *How to explain the frequency use of plainclothes in practice?*
- **Strategic Secrecy vs. Commitment**
- **Perfect Bayesian Equilibrium (PBE)**
- **Computing PBE**
- **PBE vs SSE**

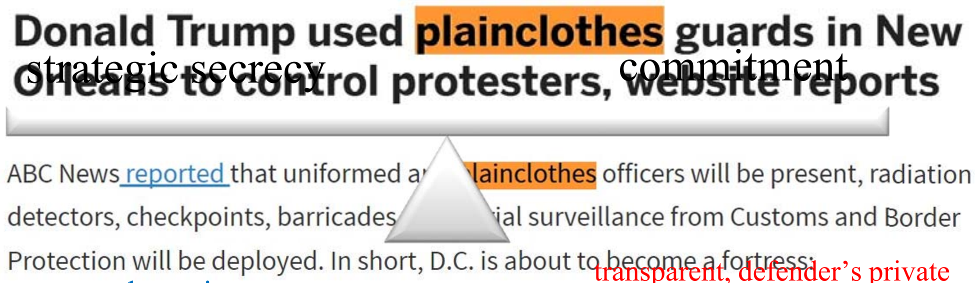
Donald Trump used plainclothes guards in New Orleans to control protesters, website reports

ABC News reported that uniformed and plainclothes officers will be present, radiation detectors, checkpoints, barricades ... ial surveillance from Customs and Border Protection will be deployed. In short, D.C. is about to become a fortress.

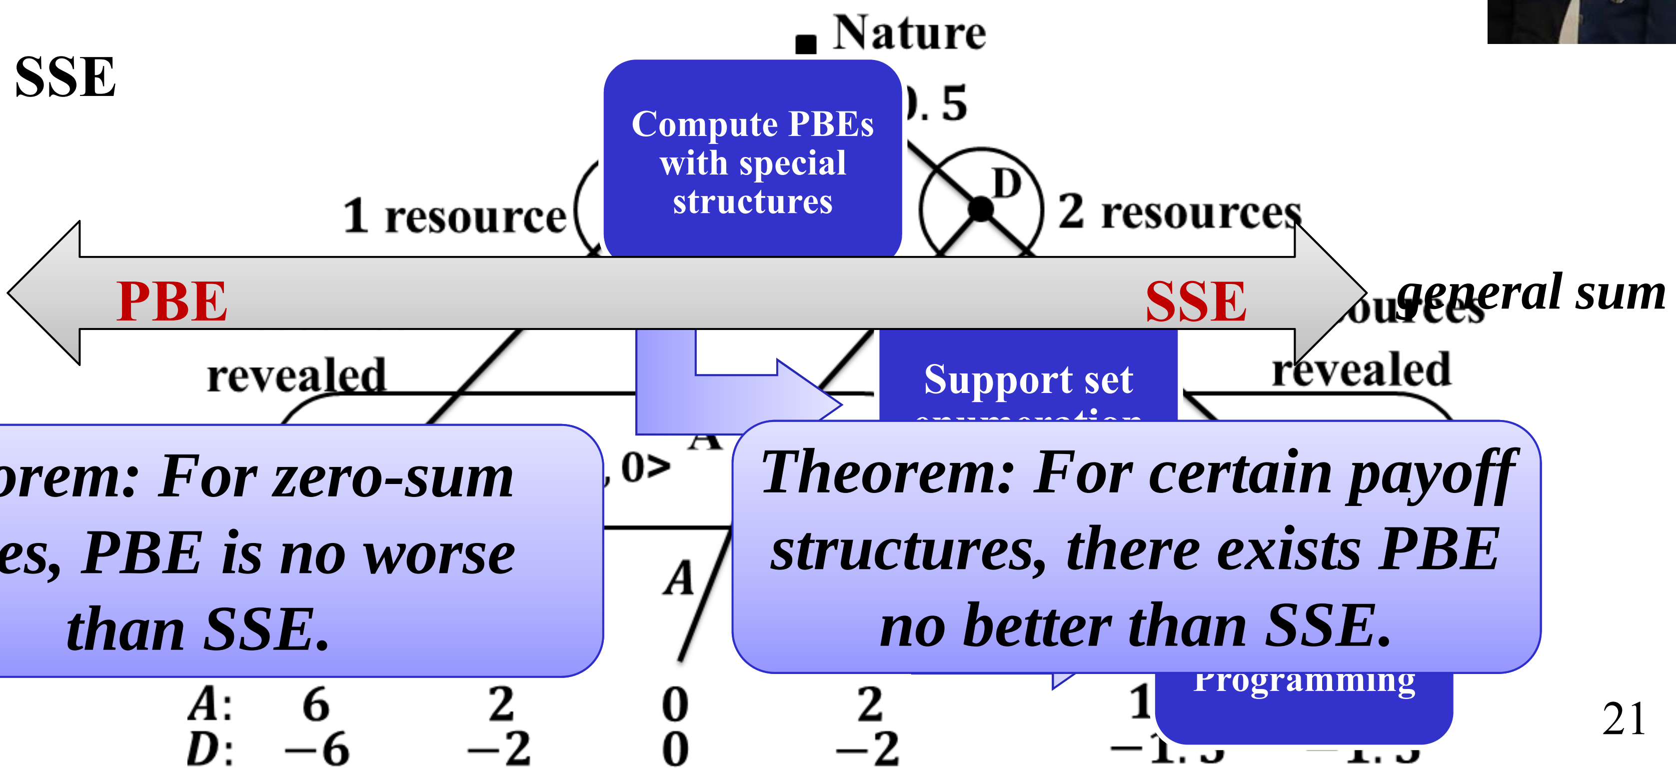

Theorem: For zero-sum games, PBE is no worse than SSE.

Theorem: For certain payoff structures, there exists PBE no better than SSE.

Cyber Security: Spear Phishing Attacks [AAAI'16b]

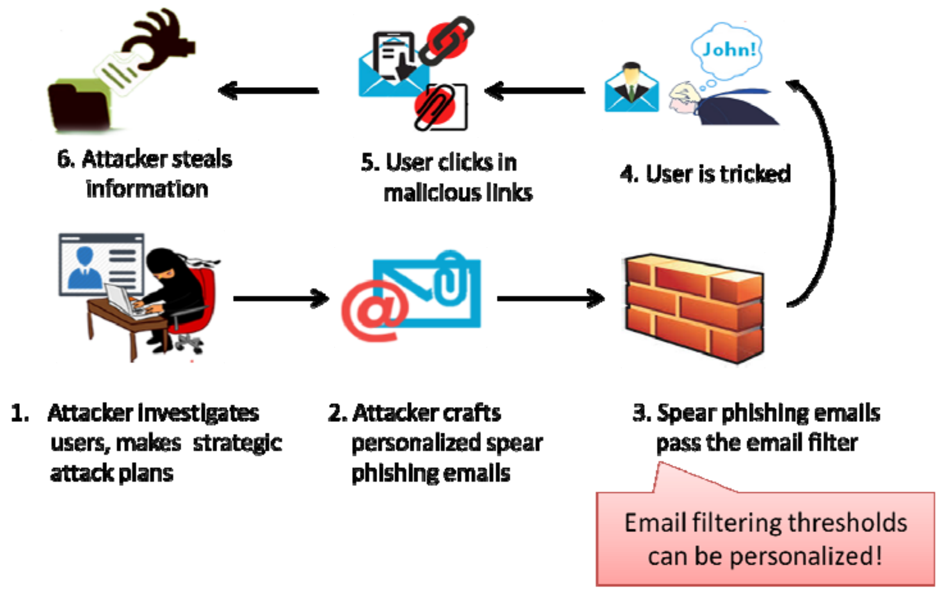

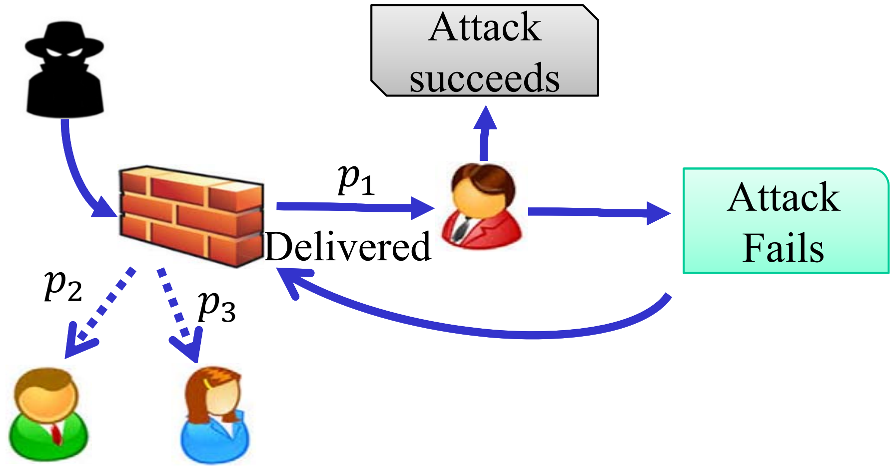

- Sequential Attack
  - *Attack targeted users sequentially in order to save attack cost*
  - *Stop attacking when the credential is accessed*
- Methods and Algorithms
  - *Attacker's problem: MDP*
  - *Defender's problem: Optimally control the transitions of the attacker's MDP*

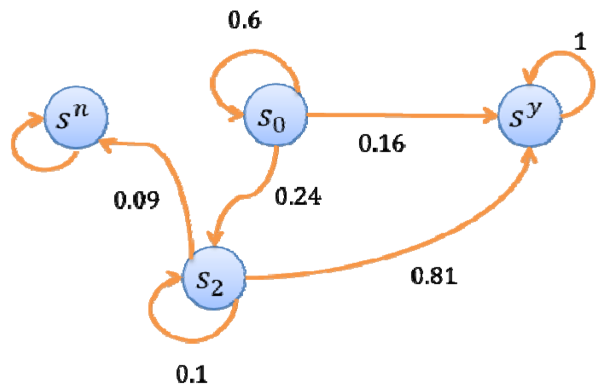

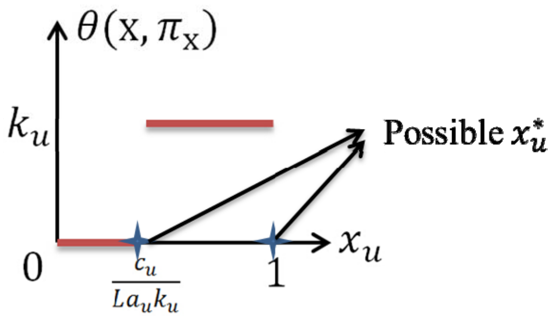

Label Contamination Attack [IJCAI'17c]

Motivating Domains:

amazon mechanical turk™
Artificial Artificial Intelligence

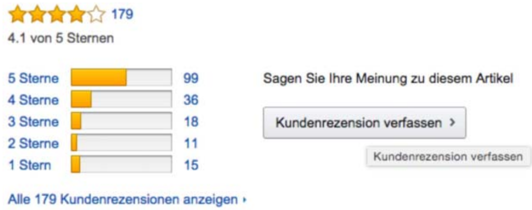

Label Flipping:

Decision Boundaries

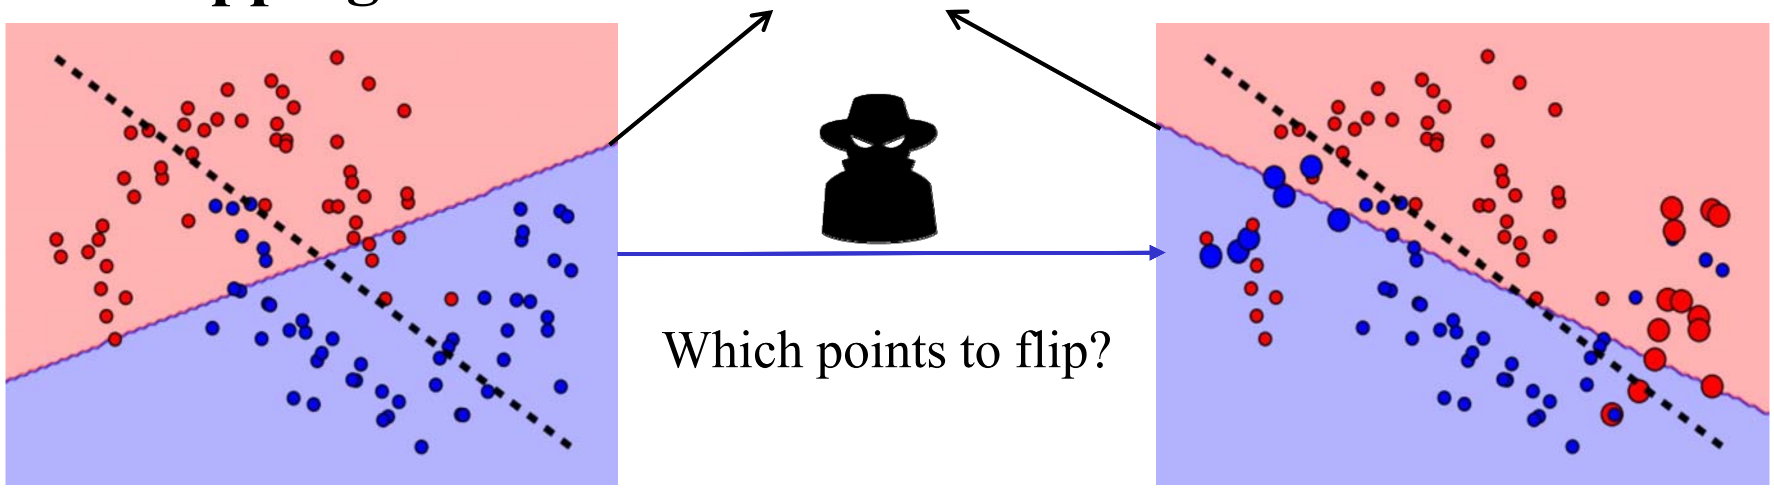

0 Flip

20% Flips

Black-Box Learning Model:

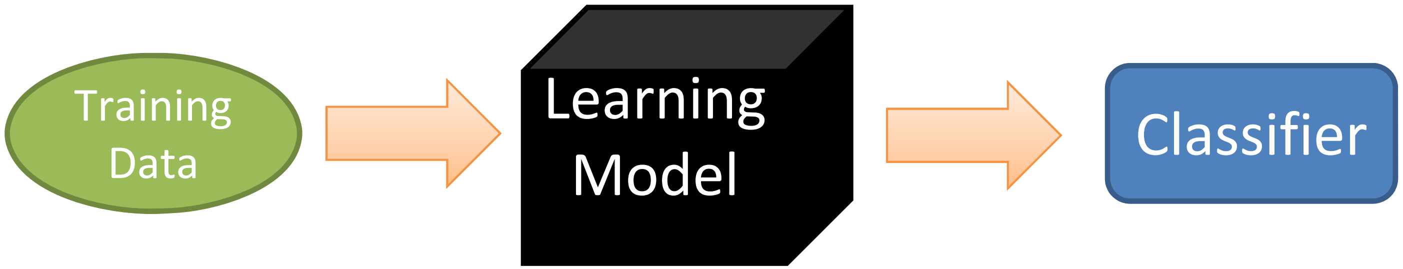

Optimal Attack and Potential Defense

Formulation: Bilevel integer program (NP-hard)

$$\max_{\mathbf{z}} \quad \frac{\mathbf{w}^\intercal \mathbf{w}^*}{||\mathbf{w}||||\mathbf{w}^*||}$$

Distance between learner and attacker's model

$$\text{s.t.} \quad f \in \arg\min_{g \in \mathcal{H}} \quad C \sum_{i=1}^{n} L(y_i', g(\mathbf{x}_i)) + \frac{1}{2}||g||^2$$

Learning: empirical risk minimization

$$\left.\begin{array}{l} \sum_{i=1}^{n} z_i \leq B \\ y_i' = y_i(1 - 2z_i), \forall i \in [n] \\ z_i \in \{0, 1\}, \forall i \in [n] \end{array}\right\}$$

Budget constraints

Approach: Projected gradient ascent

- Relax z_i to interval [0,1]
- Update: $\mathbf{z}^t \leftarrow \text{Proj}(\mathbf{z}^{t-1} + \eta \nabla_{\mathbf{z}^{t-1}} U)$
- Obtain near-optimal attacks

Use substitute models to compute candidate attacks

Cyber Security for Smart Traffic Control

➢ Cyber Attacks

- ❑ E.g., hack the sensors and send fake data to the controller, remotely hack into the controller and take control.
- ❑ Example from Hollywood movies (e.g., The Italian Job)
- ❑ *Real-world example*
  - ❖ Israeli students attacked Waze APP with fake traffic jam in 2014

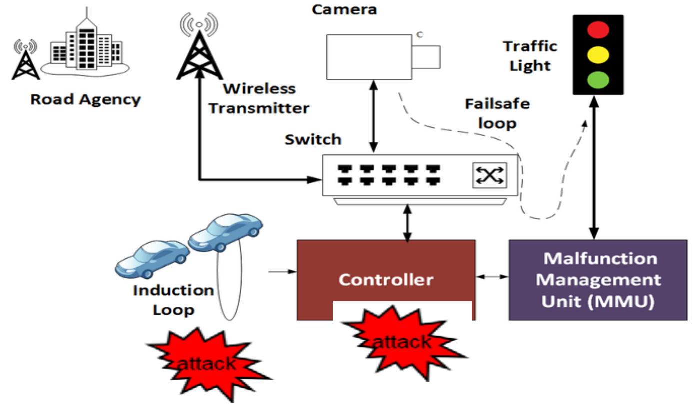

➢ STRICT: Secure TRaffIc ConTrol

- ❑ Identification, defence
- ❑ Complex system, heterogeneous interaction, dynamic

<u>A game theoretic defense mechanism against data poisoning attacks</u>

- ➢ Attacker can poison sensor data
- ➢ A verification based online defense mechanism

<u>Optimal escape interdiction on transportation networks</u> [IJCAI'17d]

- ➢ A defender-attacker security game model
- ➢ Dynamically relocate security resources (e.g., police cars) to interdict attacker

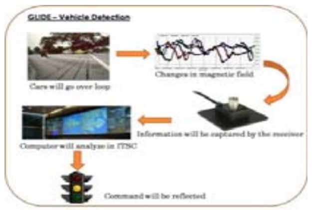

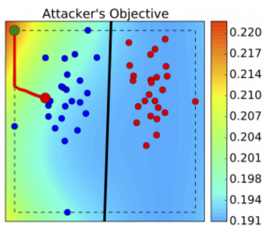

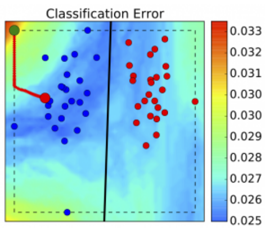

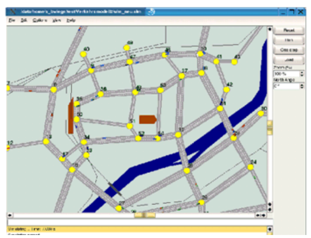

Urban Intelligence

- Rapid growth and increasing issues
  - *> 5 billion urban population by 2060*
  - *Unprecedented burden to transportation, energy, wastes, security…*
- Conflicts:
  - *Individual needs vs. social welfare*
  - *Flexibility? Fairness? Efficiency?*
- Optimal policy making is crucial
  - *Optimal pricing for taxi systems*
  - *Planning and pricing for electric vehicle (EV) charging networks*

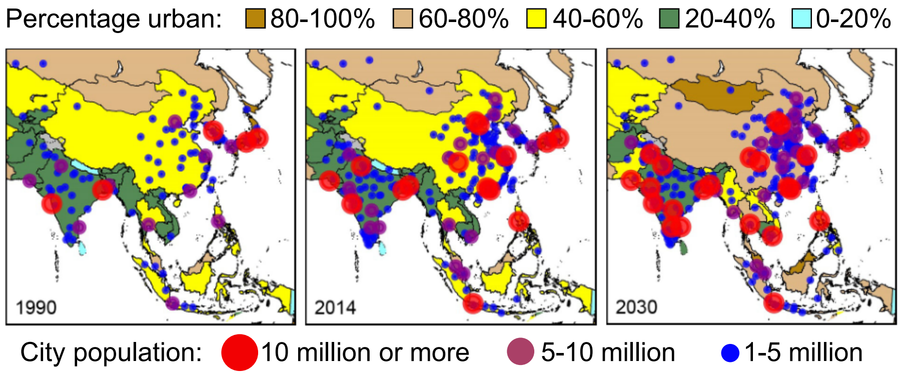

Taxi System Efficiency Optimization [IJCAI'13, AAMAS'15]

- Peak time dilemma in Beijing
  - ***Low speed*** *in peak time &* ***distance-based pricing scheme***

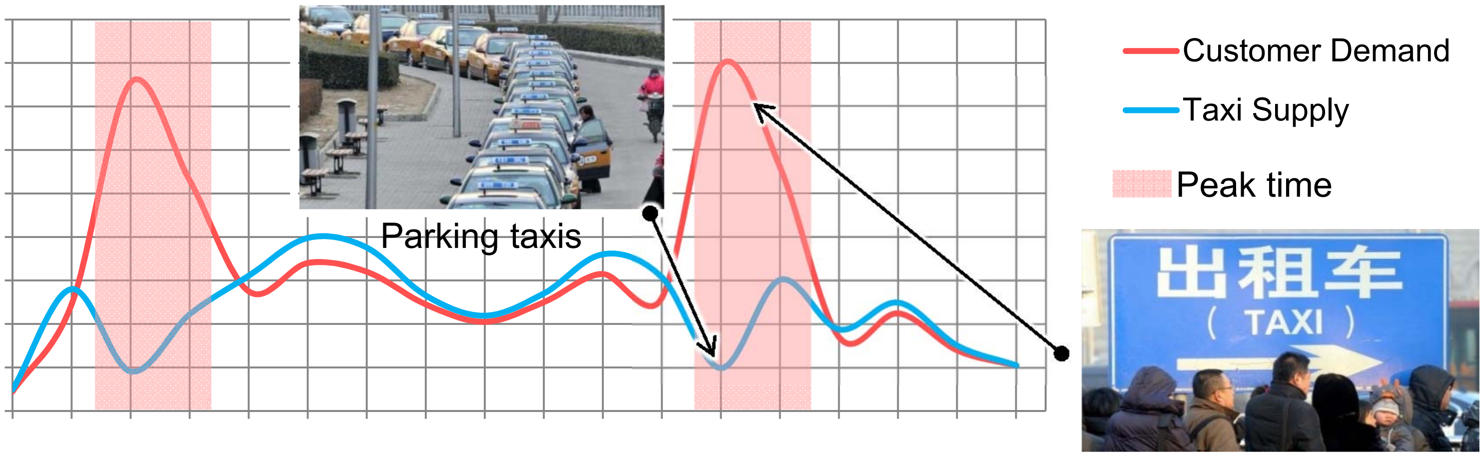

- Solution: Increase **fare price** in peak time

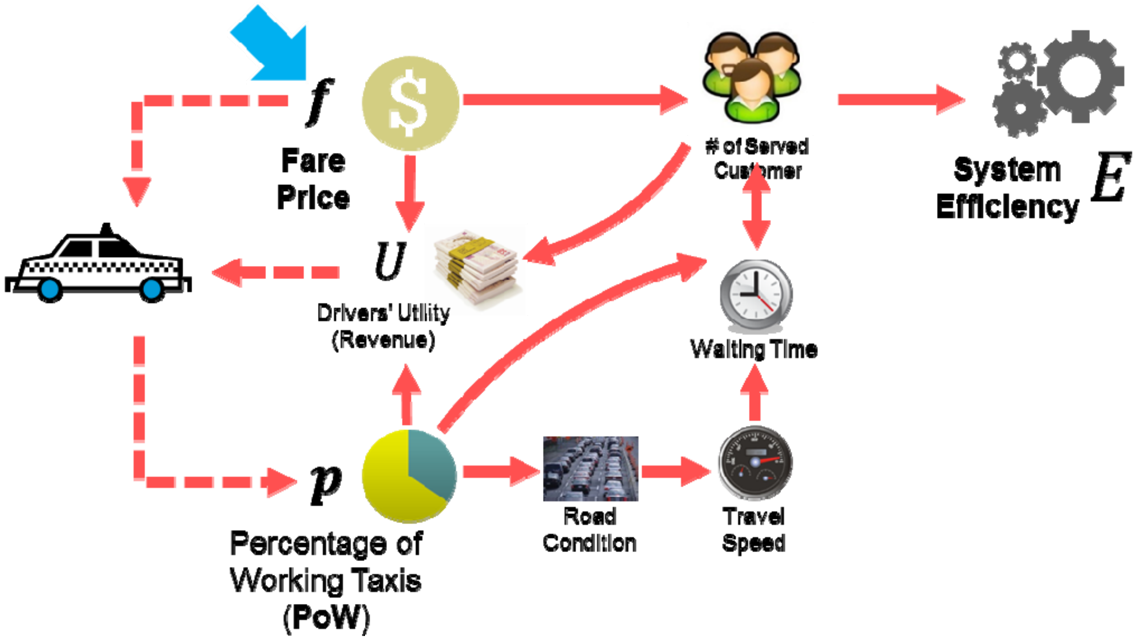

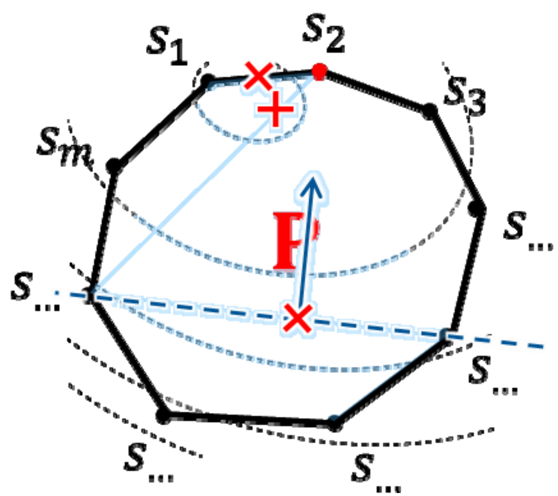

EV Charging Station Placement [IJCAI'15, AAMAS'16]

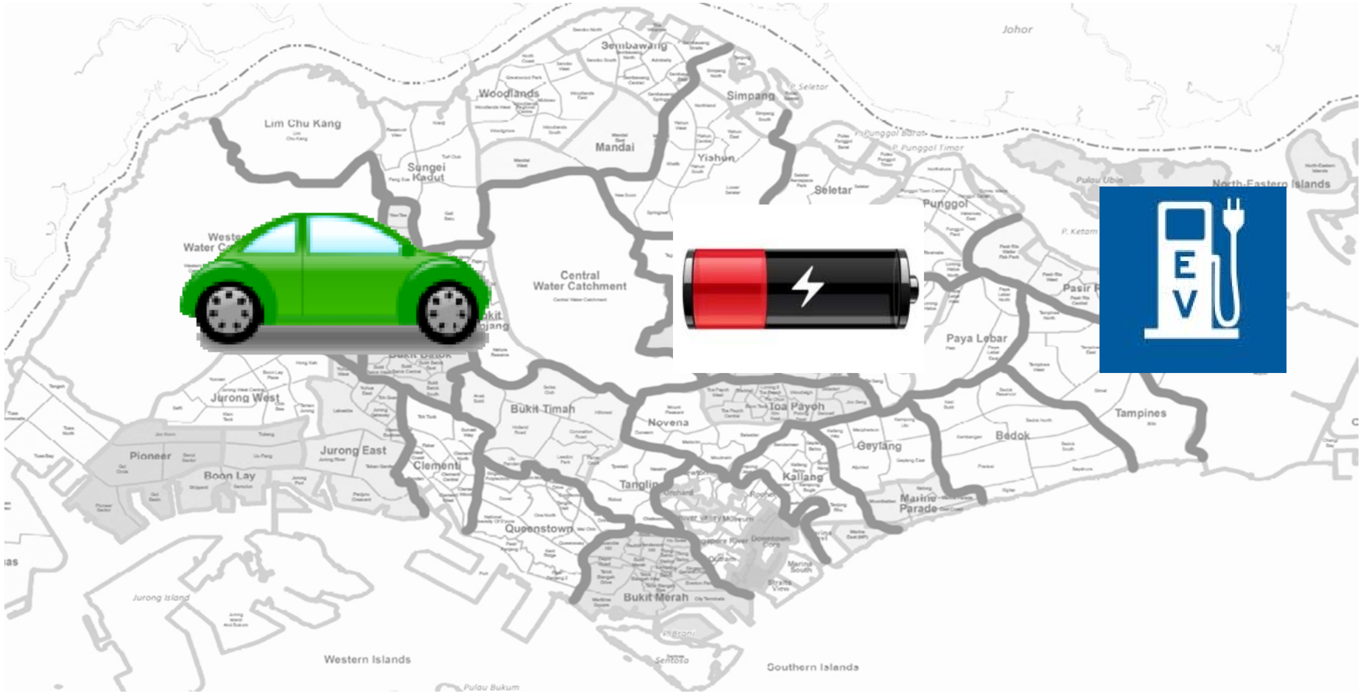

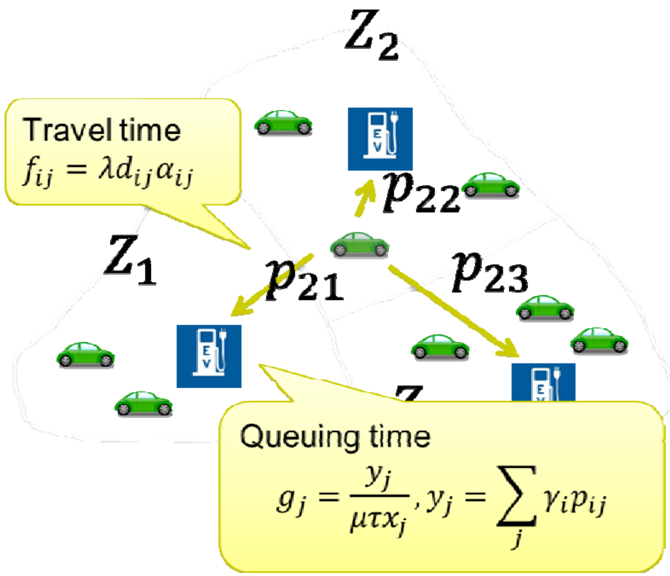

- How to assign ***m*** chargers to ***n*** zones to minimize social cost ***SC*** of all the EVs
- Bilevel optimization with Infinite number of constraints
  - *Nash Equilibrium in non-atomic congestion game*
- Optimal pricing – nonlinear optimization

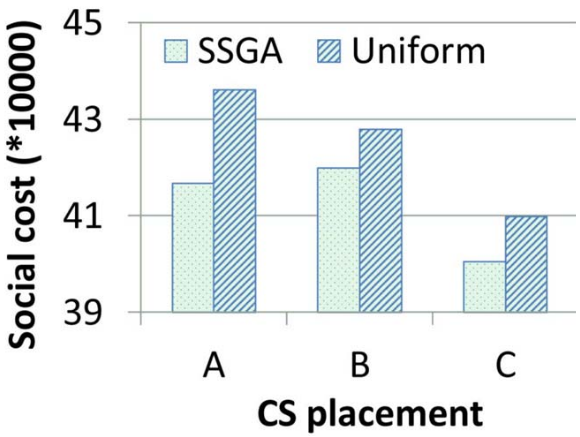

Summary

- Security Games
  - *High impact research and deployed applications*
  - *Lot of follow on research & applications: data-rich domains, cyber*
  - *Future challenges: more realistic modes, scalability, robustness, optimality*
- Computational Sustainability
  - *Applications involving human beings*
  - *Interaction with machine learning*
  - *Future challenges: behaviour modelling, scalability, robustness*
- Acknowledgement
  - *Mentor:*

Milind Tambe

  - *Collaborators: Yevgeniy Vorobeychik, Christopher Kiekintveld, Long Tran-Than, Branislav Bosansky, Jianye Hao, Yair Zick, many others, and students*
  - *Funding agencies*