# Game Theory for Security:

*Key Algorithmic Principles, Deployed Systems, Research Challenges*


***Milind Tambe***

**Helen N. and Emmett H. Jones Professor in Engineering**

***University of Southern California***


with:

***Current/former PhD students/postdocs:***

Bo An, Matthew Brown, Francesco Delle Fave, Fei Fang, Benjamin Ford, William Haskell, Manish Jain, Albert Jiang, Debarun Kar, Chris Kiekintveld, Rajiv Maheswaran, Janusz Marecki, Praveen Paruchuri, Jonathan Pearce,James Pita, Thanh Nguyen, Yundi Qian, Eric Shieh, Jason Tsai, Pradeep Varakantham, Haifeng Xu, Amulya Yadav, Rong Yang, Zhengyu Yin, Chao Zhang

TEAMCORE
USC

***Other collaborators:***

Fernando Ordonez (USC & U Chile), Richard John (USC), David Kempe (USC), Shaddin Dughmi (USC) &

Craig Boutilier (Toronto), Jeff Brantingahm (UCLA), Vince Conitzer (Duke), Sarit Kraus (BIU, Israel), Andrew Lemieux (NCSR),Kevin Leyton-Brown (UBC), M. Pechoucek (CTU, Czech R), Ariel Procaccia (CMU), Tuomas Sandholm (CMU), Martin Short (GATech), Y. Vorobeychik (Vanderbilt), ….

# Global Challenges for Security:
# Game Theory for Security Resource Optimization

# Example Model: Stackelberg Security Games

Security allocation:

- Target weights
- Opponent reaction

**Adversary**

**Defender**

| | Target #1 | Target #2 |
|---|---|---|
| Target #1 | 4, -3 | -1, 1 |
| Target #2 | -5, 5 | 2, -1 |

# Example Model: Stackelberg Security Games

Security allocation:

- Target weights
- Opponent reaction

**Adversary**

**Defender**

| | Target #1 | Target #2 |
|---|---|---|
| Target #1 | 4, -3 | -1, 1 |
| Target #2 | -5, 5 | 2, -1 |

# Example Model: Stackelberg Security Games

Security allocation:

- Target weights
- Opponent reaction

**Adversary**

**Defender**

| | Target #1 | Target #2 |
|---|---|---|
| Target #1 | 4, -3 | -1, 1 |
| Target #2 | -5, 5 | 2, -1 |

# Stackelberg Security Games

## Security Resource Optimization: *Not 100% Security*

- Random/mixed strategy:
  - *Increase cost/uncertainty to attackers*
- Stackelberg game:
  - *Defender commits to mixed strategy*
- Stackelberg Equilibrium: Optimal random?

**Adversary**

**Defender**

| | Target #1 | Target #2 |
|---|---|---|
| Target #1 | 4, -3 | -1, 1 |
| Target #2 | -5, 5 | 2, -1 |

# Research Contributions: Game Theory for Security

**Computational game theory:**

- *Massive games*

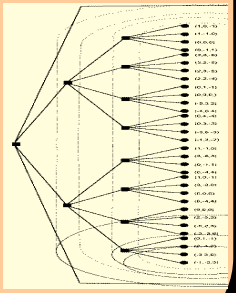

**Behavioral game theory:**

- *Exploit human behavior models*

**+ Planning under uncertainty, learning…**

## *Computational Game Theory in the Field*

# Applications: Deployed Security Assistants

**Ports & port traffic**
*US Coast Guard*

**Airports, access roads & flights**
*TSA,*
*Airport Police*

**Urban transport**
*LA Sheriff's/TSA*
*Singapore Police*

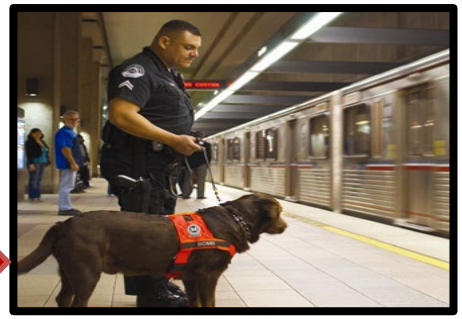

**Environment**
*US Coast Guard*
*& others*

# Key Lessons: Security Games

Decision aids based on computational game theory in daily use

- Optimize limited security resources against adversaries

Applications yield research challenges: Science of security games

- Scale-up: Incremental strategy generation & Marginals
- Uncertainty: Integrate MDPs, Robustness, Quantal response

Current applications (wildlife security): AI challenge problems

- Global challenges: Merge planning/learning & security games

# Outline: "Security Games" Research (2007-Now)

***Research Challenges:***
***Scale-up, Uncertainty***

**Airports**

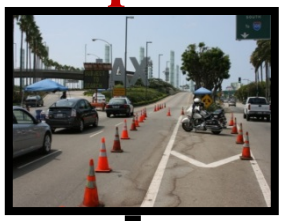

**Flights**

**Roads**

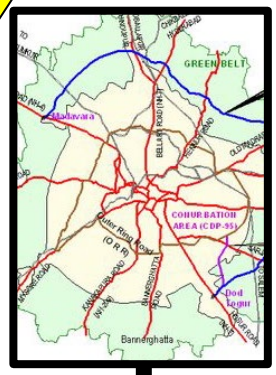


**Ports**

**Trains**

**Environment**

2007 — 2009 — 2011 — 2012 — 2013 — 2013-

***Evaluation I:*** *Scale up? Handle uncertainty?*

***Evaluation II:*** *Real-world deployments*

*(Patience)*

***Publications:***
*AAMAS, AAAI, IJCAI…*
*2007 onwards*

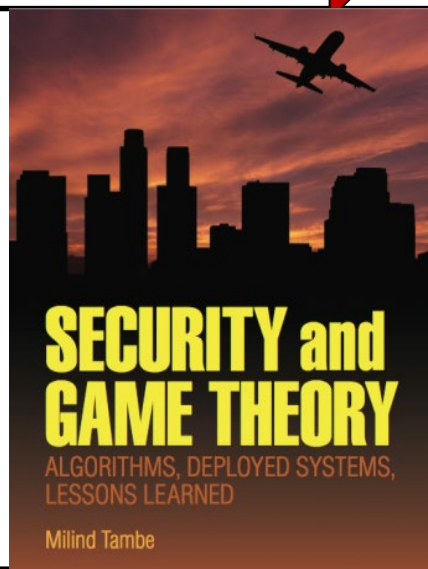

# Airport Security: Mapping to Stackelberg Games

**ARMOR: LAX (2007)**

**GUARDS: TSA (2011)**

- **6 plots against LAX**

GLASGOW 6/30/07

# ARMOR Operation [2007]

## Generate Detailed Defender Schedule

*Pita Paruchuri*

| | Target #1 | Target #2 |
|---|---|---|
| Defender #1 | 2, -1 | -3, 4 |
| Defender #2 | -3, 3 | 3, -2 |

***Mixed Integer Program***

***Pr(Canine patrol, 8 AM @Terminals 2,5,6) = 0.17***

***Pr(Canine patrol, 8 AM @ Terminals 3,5,7) = 0.33***

***Canine Team Schedule, July 28***

| | Term 1 | Term 2 | Term 3 | Term 4 | Term 5 | Term 6 | Term 7 | Term 8 |
|---|---|---|---|---|---|---|---|---|
| **8 AM** | | Team1 | | | Team3 | Team5 | | |
| **9 AM** | | | Team1 | Team2 | | | | Team4 |
| **10 AM** | | Team3 | | Team5 | | Team2 | | |

# ARMOR MIP [2007]

## Generate Mixed Strategy for Defender

*Pita* *Paruchuri*

| | Target #1 | Target #2 |
|---|---|---|
| Defender #1 | 2, -1 | -3, 4 |
| Defender #2 | -3, 3 | 3, -2 |

$$\max \sum_{i \in X} \sum_{j \in Q} R_{ij} \times x_i \times q_j$$

***Maximize defender expected utility***

$$s.t. \quad \sum_i x_i = 1$$

***Defender mixed strategy***

$$\sum_{j \in Q} q_j = 1$$

***Adversary response***

$$0 \le (a - \sum_{i \in X} C_{ij} x_i) \le (1 - q_j) M$$

***Adversary best response***

# ARMOR Payoffs [2007]

## Previous Research Provides Payoffs in Security Game Domains

| | Target #1 | Target #2 |
|---|---|---|
| Defender #1 | 2, -1 | -3, 4 |
| Defender #2 | -3, 3 | 3, -2 |

$$\max \sum_{i \in X} \sum_{j \in Q} R_{ij} \times x_i \times q_j$$

*Maximize defender expected utility*

***+ Handling Uncertainty***

# ARMOR MIP [2007]

## Solving for a Single Adversary Type

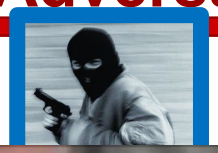

Newsweek

*ARMOR…throws a digital cloak of invisibility….*

*Handling Uncertain*

*nder*

*ected utility*

*fender strategy*

*versary strategy*

$$0 \le (a - \sum_{i \in X} C_{ij} x_i) \le (1 - q_j) M$$

*Adversary best response*

# IRIS: Federal Air Marshals Service [2009]
# Scale Up Number of Defender Strategies

- *1000 Flights, 20 air marshals: $10^{41}$ combinations*
  - *ARMOR out of memory*
- *Not enumerate all combinations:*
  - *Branch and price: Incremental strategy generation*

# IRIS: Scale Up Number of Defender Strategies [2009]
# Small Support Set for Mixed Strategies

**Small support set size:**

- Most xi variables zero

$$\max_{x,q} \sum_{i \in X} \sum_{j \in Q} R_{ij} x_i q_j$$

$$s.t. \sum_i x_i = 1, \sum_{j \in Q} q_j = 1$$

$$0 \le (a - \sum_{i \in X} C_{ij} x_i) \le (1 - q_j) M$$

$$x_i \in [0 \ldots 1], \; q_j \in \{0,1\}$$

*1000 flights, 20 air marshals:*

$10^{41}$ *combinations*

| | Attack 1 | Attack 2 | Attack … | Attack 1000 |
|---|---|---|---|---|
| 1,2,3 | 5, 10 | 4, 8 | … | 20,9 |
| 1,2,4.. | 5,-10 | 4,-8 | … | -20,9 |
| 1,3,5.. | 5,-10 | -9,5 | … | -20,9 |
| … | | | | |
| … | | | | |

x123=0.0

x124=0.239

x135=0.0

x378=0.123

***$10^{41}$ rows***

# IRIS: Incremental Strategy Generation
# Exploit Small Support

*Jain Kiekintveld*

***Master***

| | Attack 1 | Attack 2 | Attack… | Attack 6 |
|---|---|---|---|---|
| **1,2,4** | 5,-10 | 4,-8 | … | -20,9 |

| | Attack 1 | Attack 2 | Attack… | Attack 6 |
|---|---|---|---|---|
| **1,2,4** | 5,-10 | 4,-8 | … | -20,9 |
| **3,7,8** | -8, 10 | -8,10 | … | -8,10 |

***Slave (LP Duality Theory)***

*Best new pure strategy:*
*Minimum cost network flow*

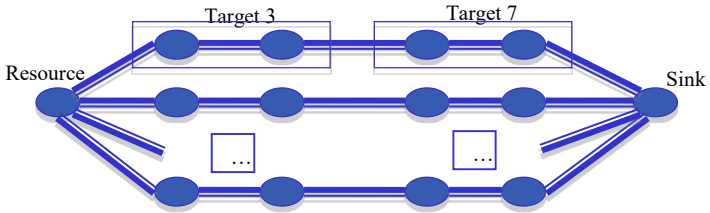


***Converge: GLOBAL OPTIMAL***

| | Attack 1 | Attack 2 | Attack | Attack 6 |
|---|---|---|---|---|
| **1,2,4** | 5,-10 | | | ,9 |
| **3,7,8** | -8, 10 | | | 0 |
| **…** | | | | |

***500 rows***
***NOT $10^{41}$***

# IRIS: Deployed FAMS (2009-)

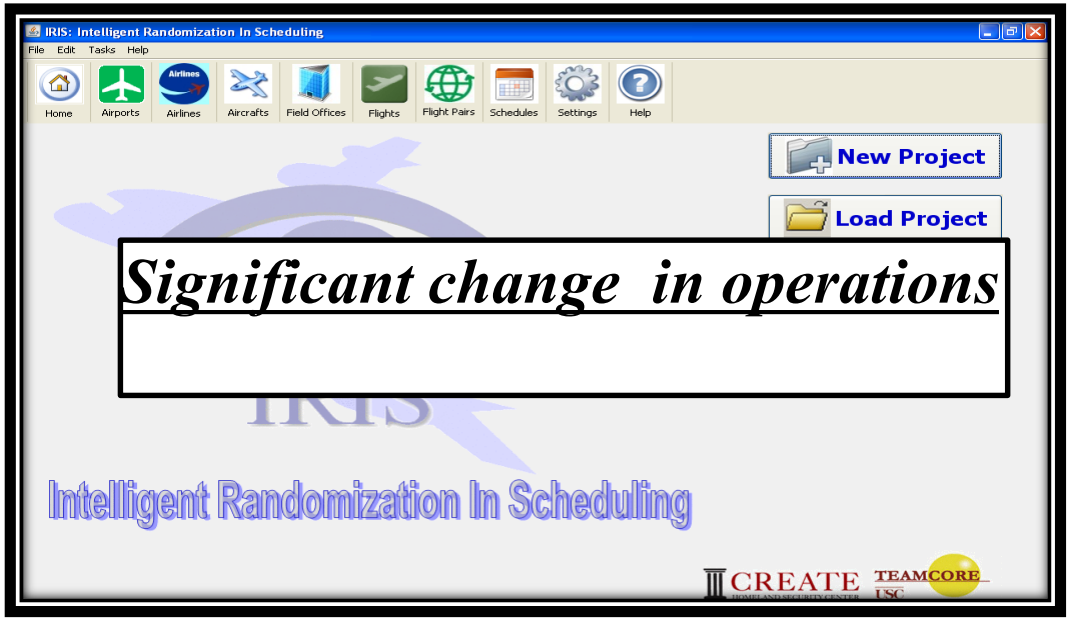

***"…in 2011, the Military Operations Research Society selected a University of Southern California project with* FAMS on randomizing flight schedules for the prestigious Rist Award*…"***

**-R. S. Bray (TSA)**

***Transportation Security Subcommittee***

***US House of Representatives 2012***

# Networks: Mumbai Police Checkpoints[2013]*

*Jain*

150 edges; 2 Checkpoints
150-choose-2 strategies

*With V Conitzer

# Networks: Mumbai Police Checkpoints[2013]*
# Incremental Strategy Generation

Jain

## *Double oracle: Converge to a global optimal*

| | Path #1 | Path #2 | Path #3 |
|---|---|---|---|
| Checkpoint strategy #1 | 5, -5 | -1, 1 | -2, 2 |
| Checkpoint strategy #2 | -5, 5 | 1, -1 | -2, 2 |

### *Defender oracle*

| | Path #1 | Path #2 |
|---|---|---|
| Checkpoint strategy #1 | 5, -5 | -1, 1 |
| Checkpoint strategy #2 | -5, 5 | 2, -1 |

### *Attacker oracle*

| | Path #1 | Path #2 | Path #3 |
|---|---|---|---|
| Checkpoint strategy #1 | 5, -5 | -1, 1 | -2, 2 |
| Checkpoint strategy #2 | -5, 5 | 1, -1 | -2, 2 |

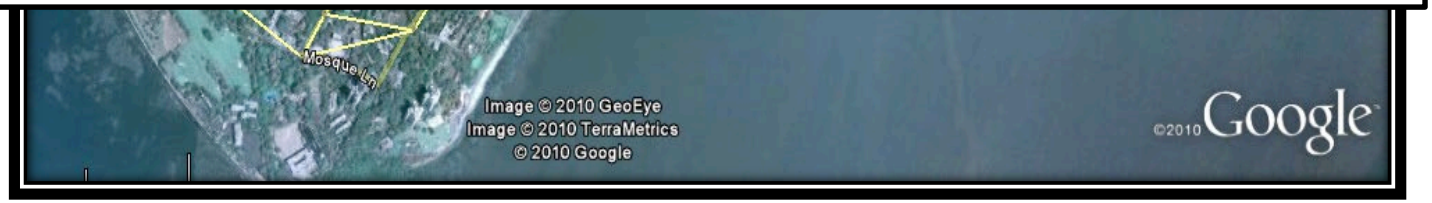


*With V Conitzer

# Double Oracle[2013]

## Incremental Strategy Generation: Exploit Small Support

*Jain*

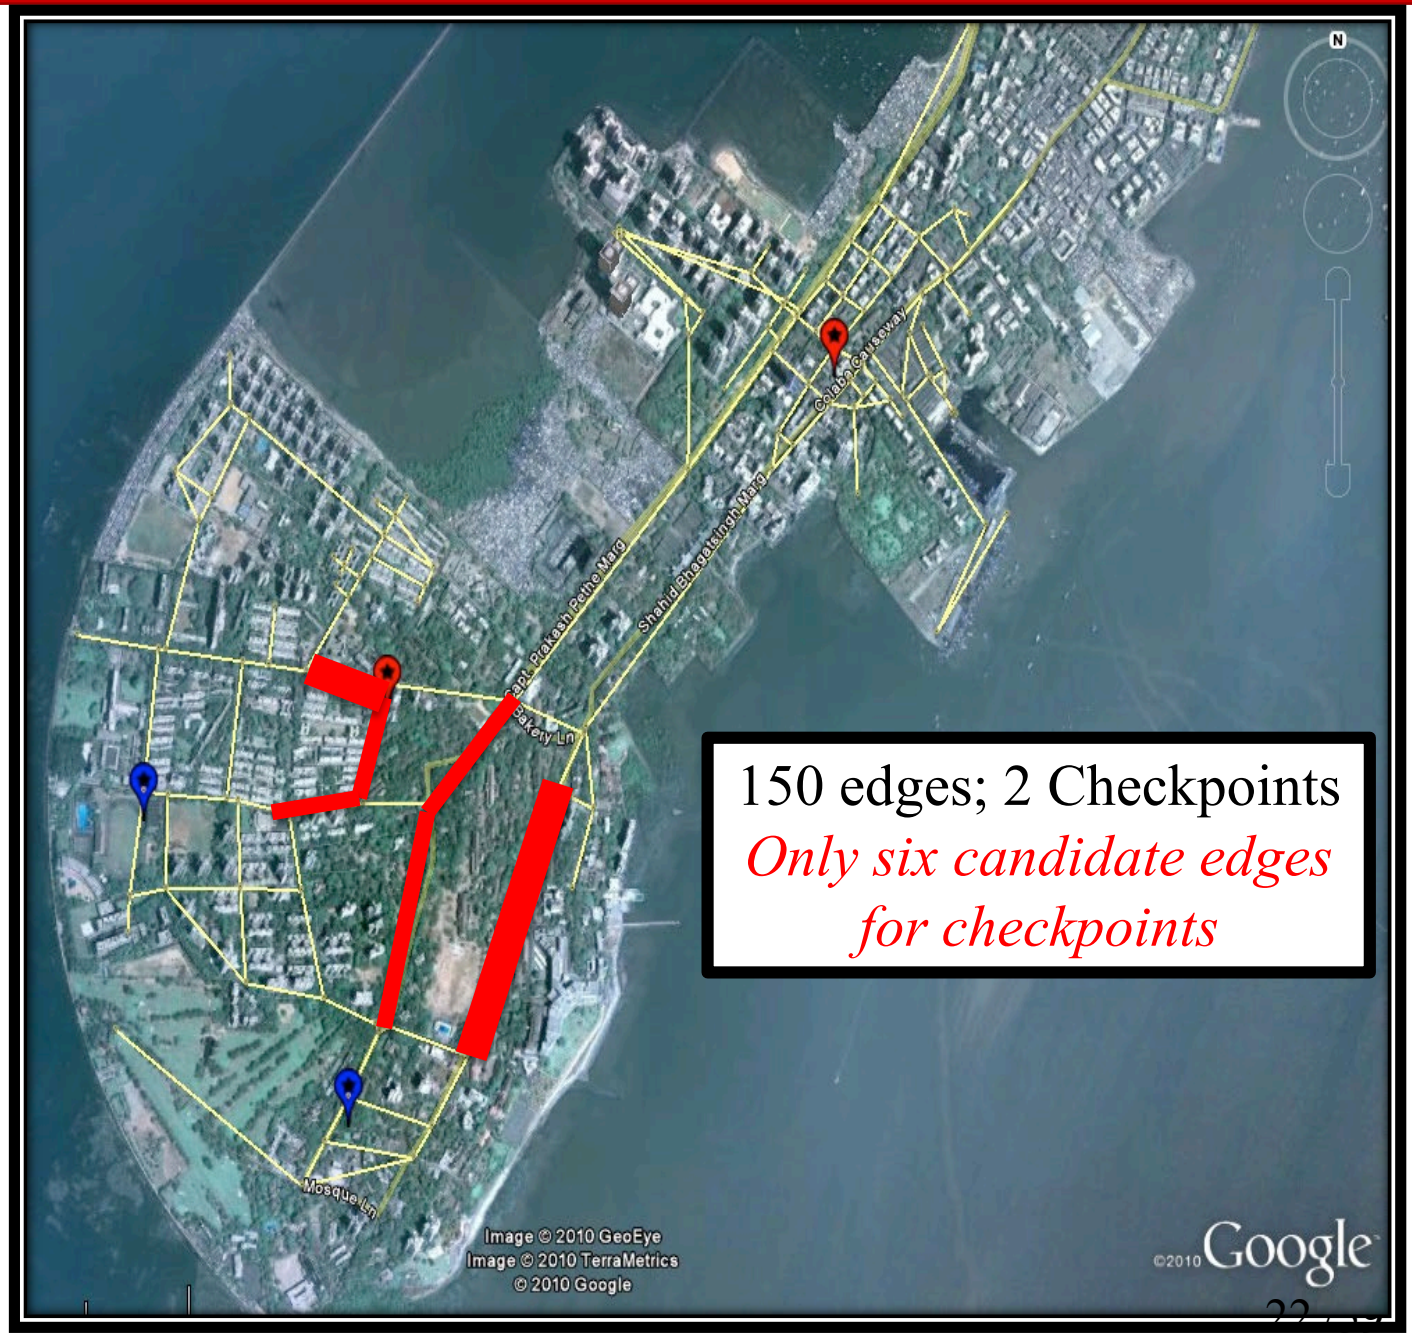


150 edges; 2 Checkpoints

*Only six candidate edges for checkpoints*

# Mumbai Police Checkpoints[2013]
# Results of Scale-up

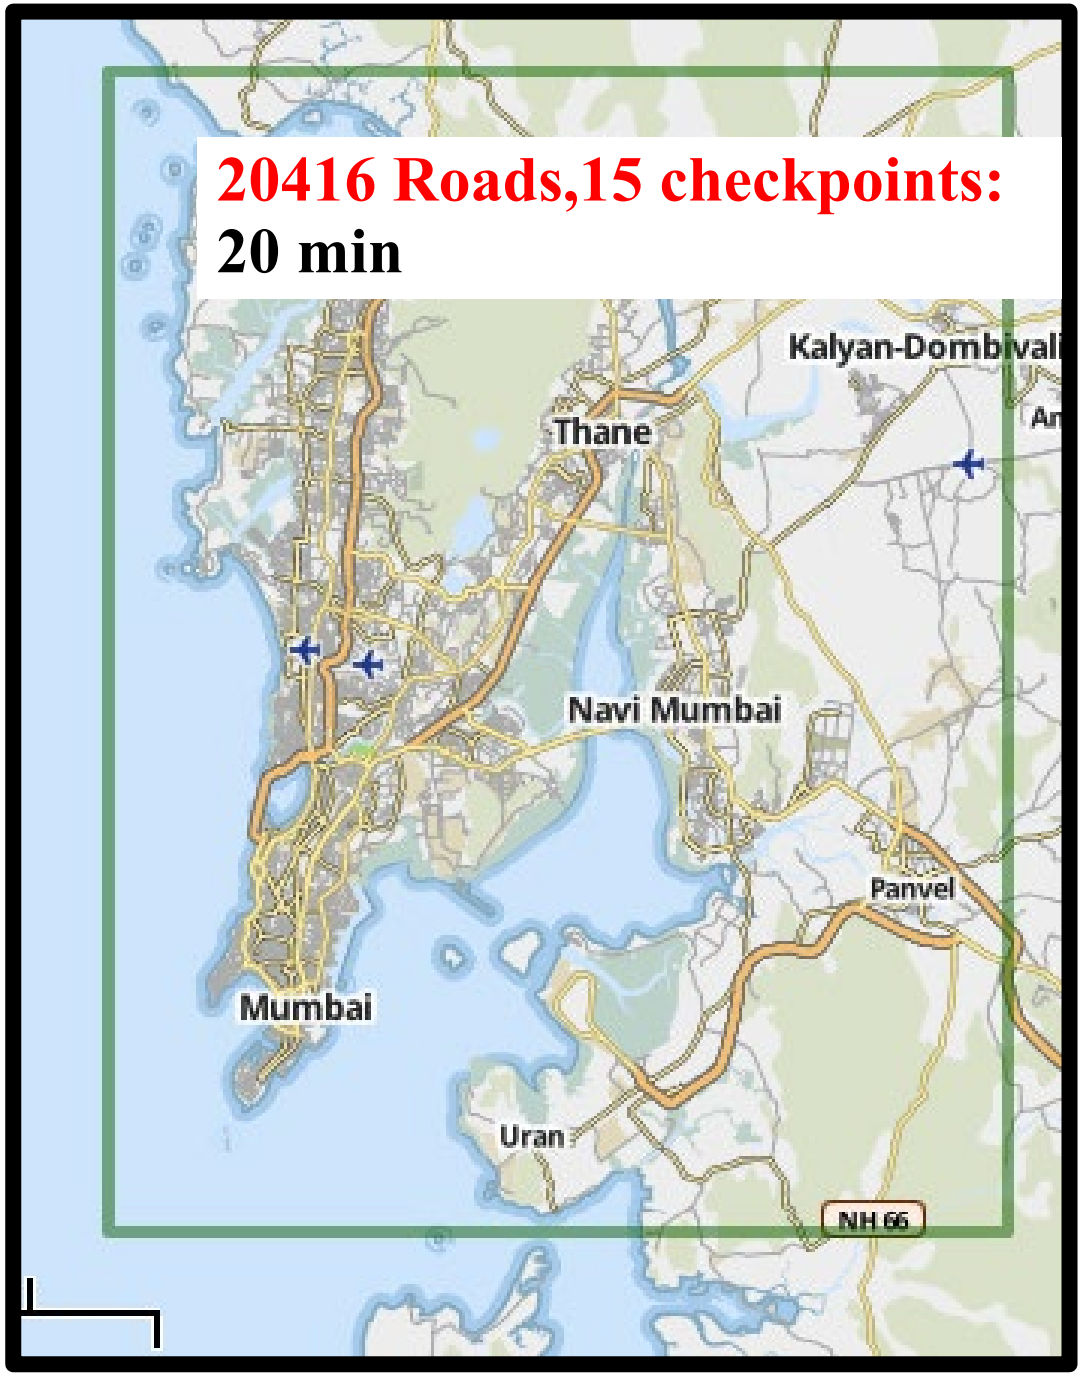

# Double Oracle: Social, Cyber Networks[2013]
# Incremental Strategy Generation

Tsai

## *Social networks: e.g., counter-insurgency*

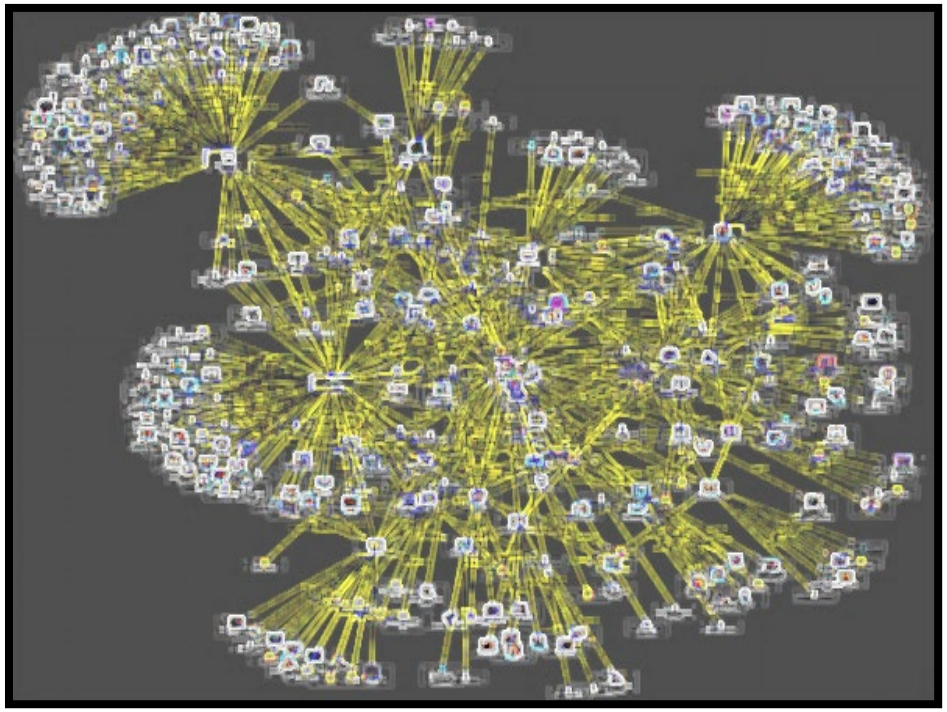

## *Cyber networks:*

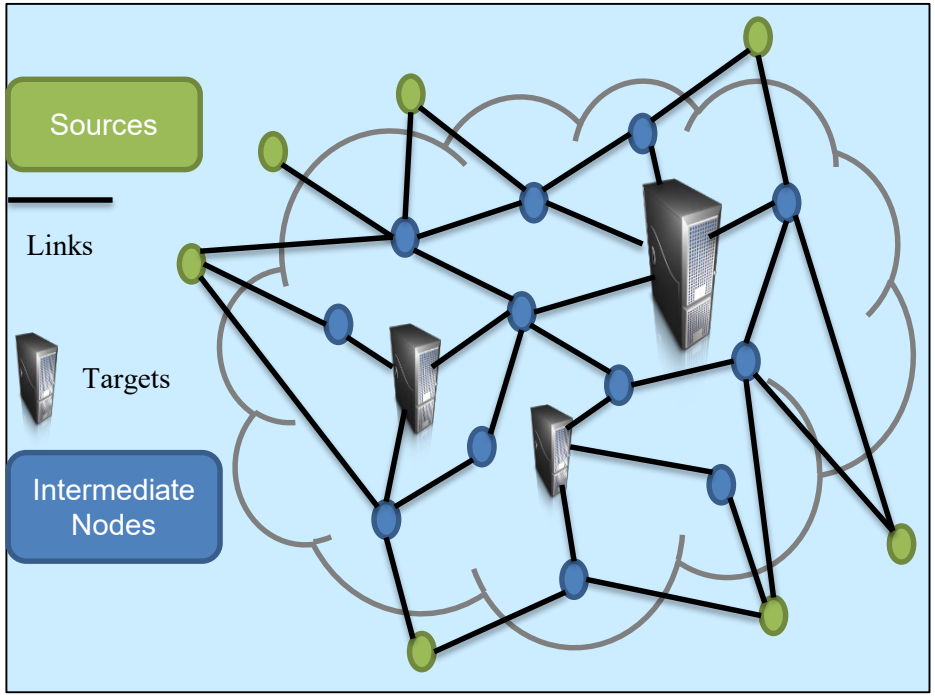

# Outline: "Security Games" Research

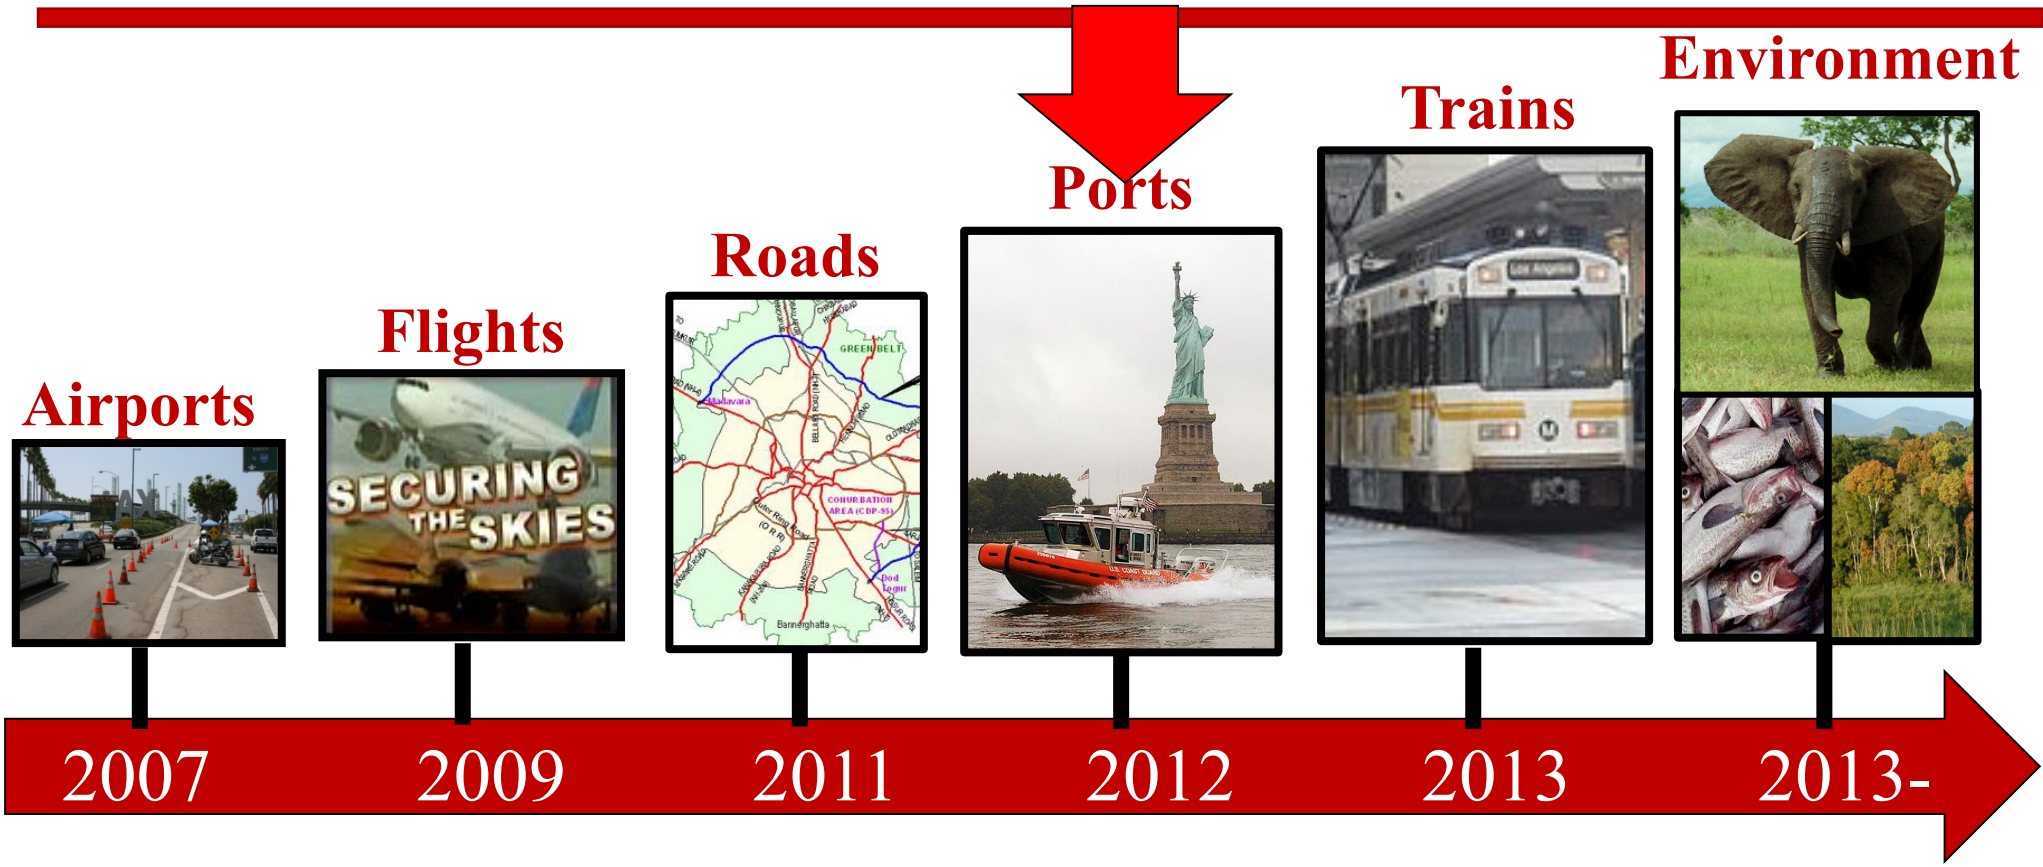

# Port Security Threat Scenarios

**US Ports: $3.15 trillion economy**

USS *Cole* after suicide attack

Attack on a ferry

French oil tanker hit by small boat

# PROTECT: Randomized Patrol Scheduling [2013]

## Port Protection (Scale-up) and Ferries (Continuous Space/time)

# PROTECT: Randomized Patrol Scheduling [2013]

## Port Protection (Scale-up) and Ferries (Continuous Space/time)

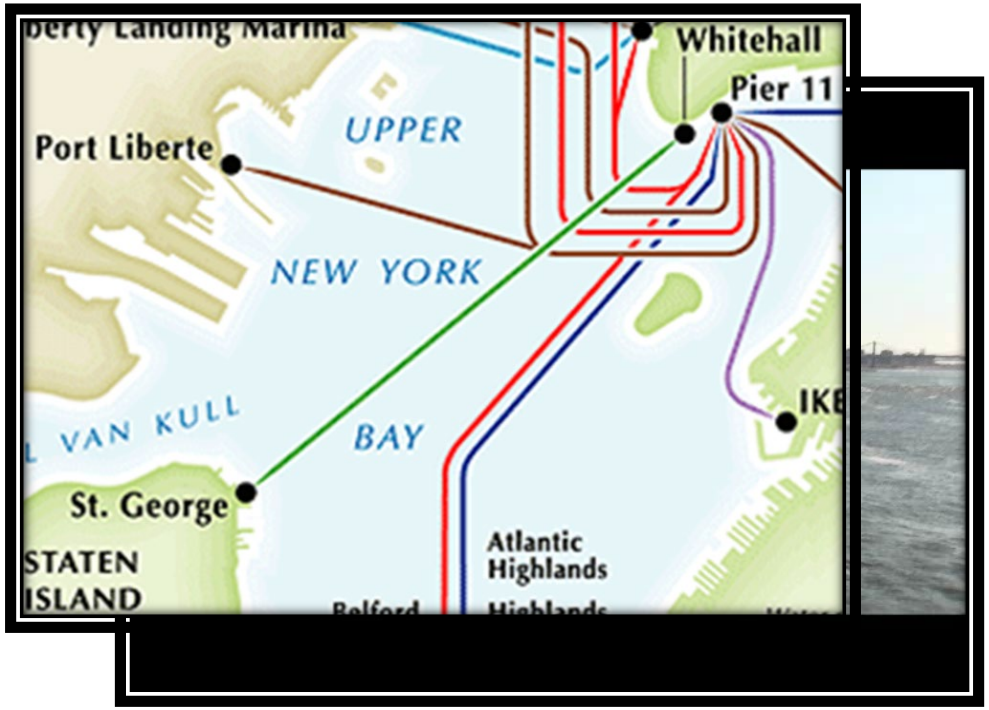

# Ferries: Scale-up with Mobile Resources & Moving Targets
# Transition Graph Representation

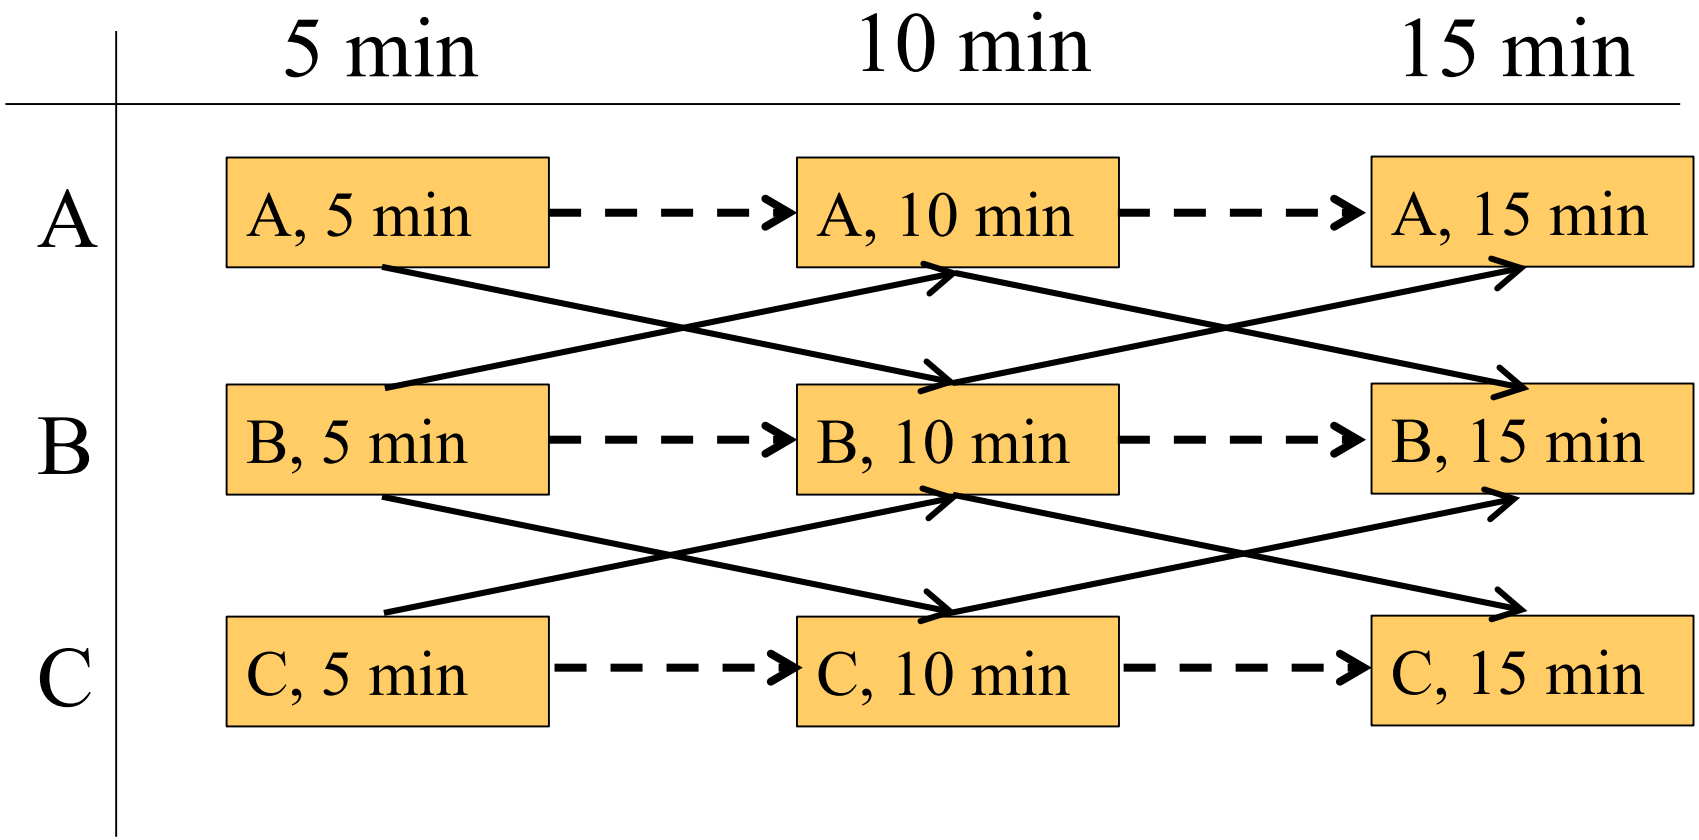

# Ferries: Scale-up with Mobile Resources & Moving Targets
# Transition Graph Representation

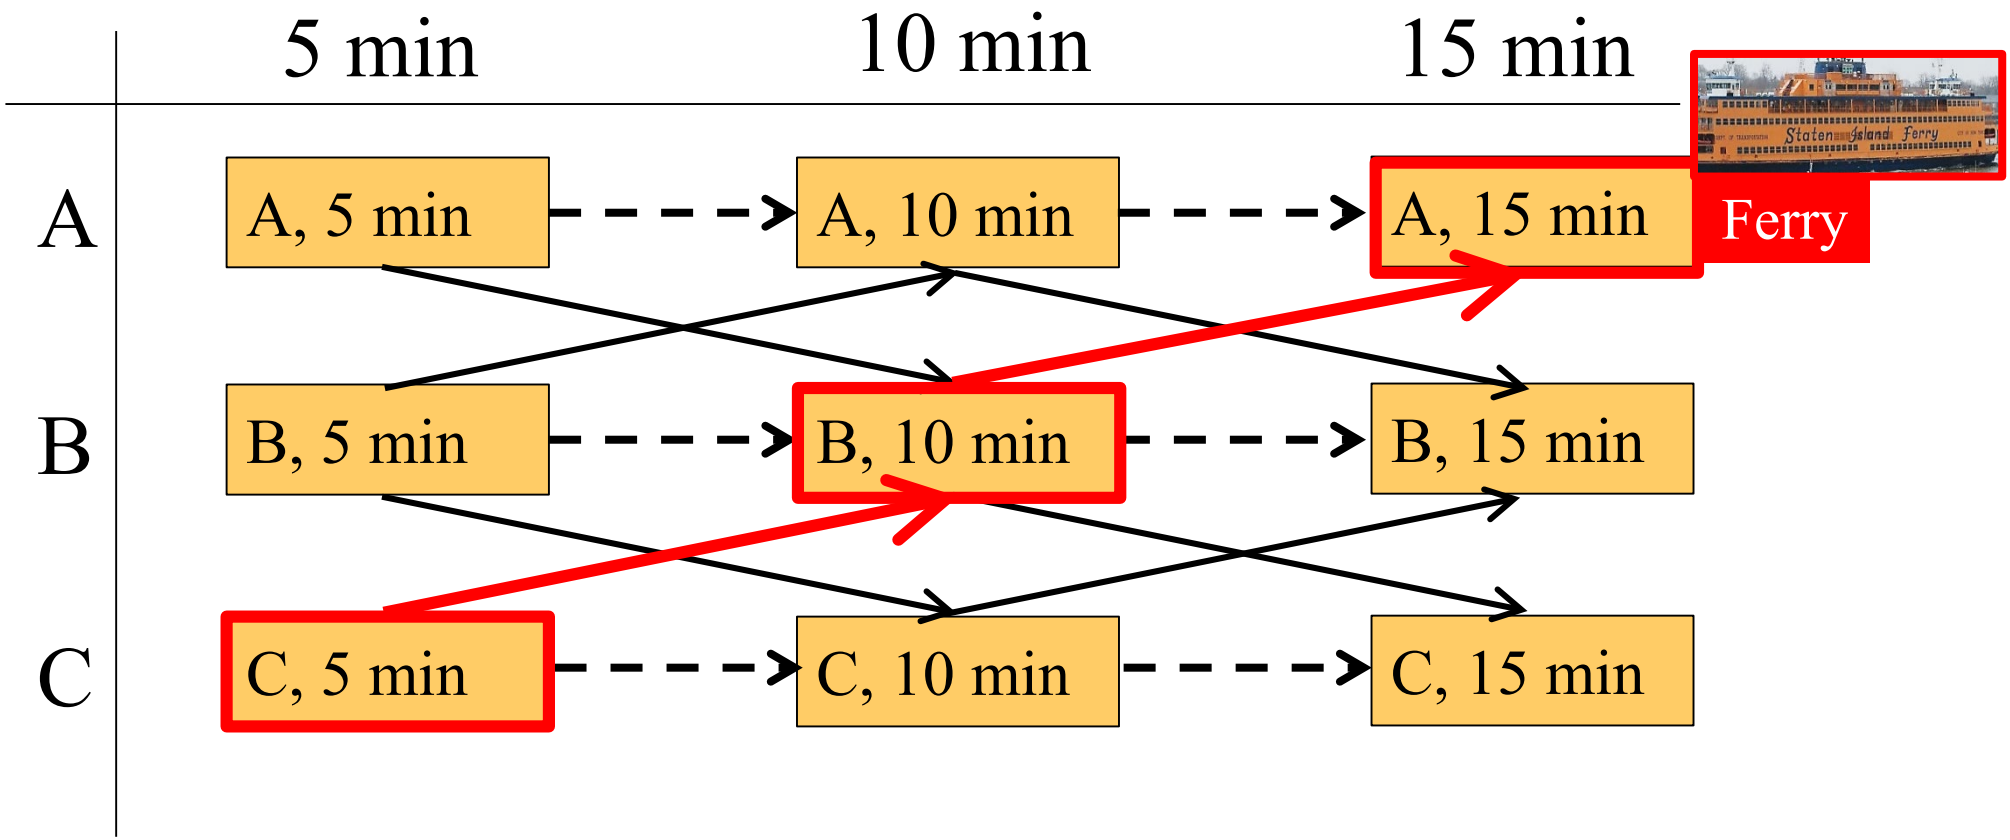

# Ferries: Patrol Routes
# Exponential Numbers of Patrol Routes

- Patrols protect nearby ferry location; Solve as done in ARMOR

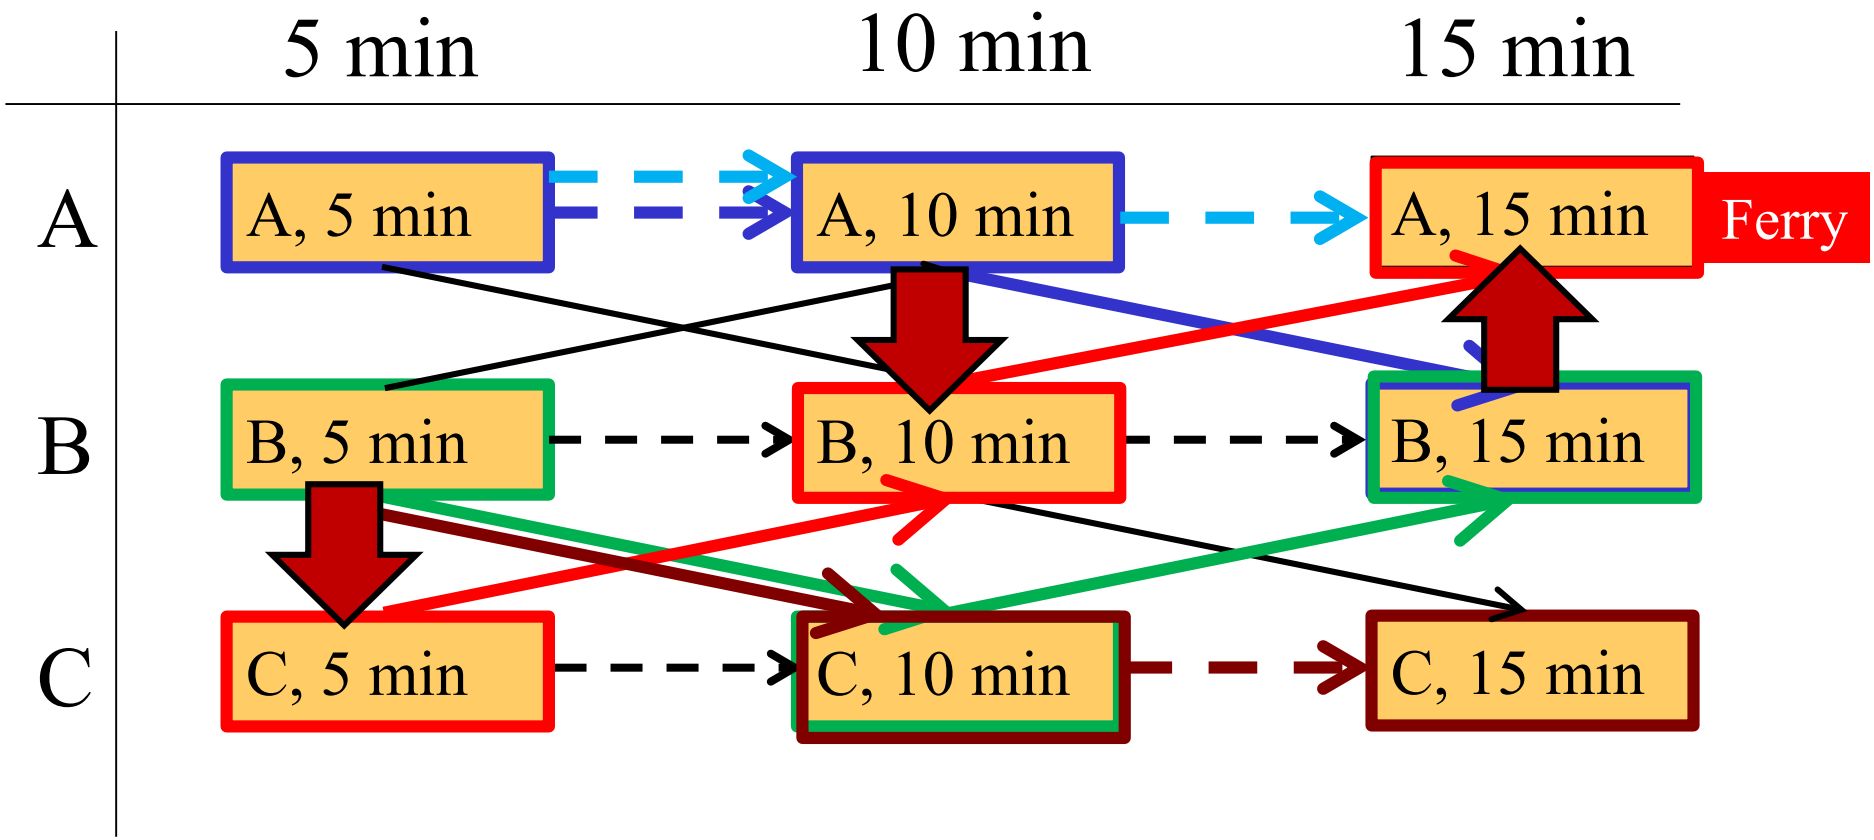

# Ferries: Patrol Routes
# Exponential Numbers of Patrol Routes

*Fang* *Jiang* *Kiekintveld*

- Patrols protect nearby ferry location; Solve as done in ARMOR
  - *Pr([(B,5), (C, 10), (C,15)]) = 0.17*
  - *Pr([(B,5), (C,10), (B,15)]) =0.13*
  - *Pr([(A,5), (A,10), (B,15)]) = 0.07*
  - *Pr([(A,5), (A,10), (A,15)]) = 0.03*

***$N^T$ variables***

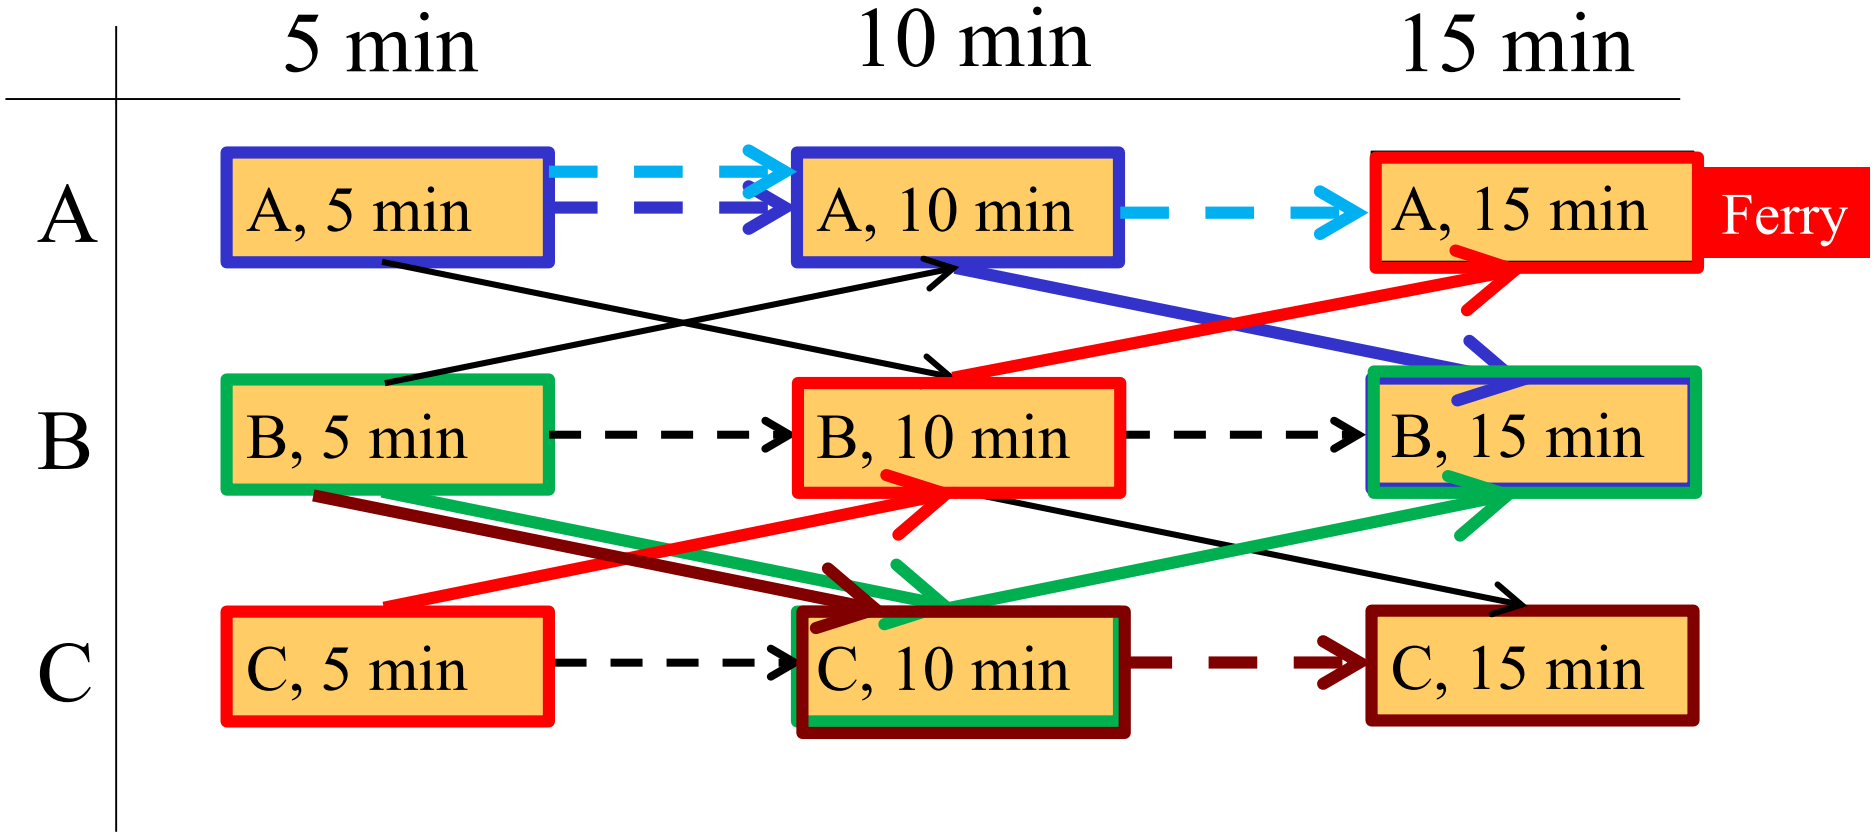

# Ferries: Scale-up
# Marginal Probabilities Over Segments

*Fang* *Jiang* *Kiekintveld*

- Variables: NOT routes, but probability flow over each segment

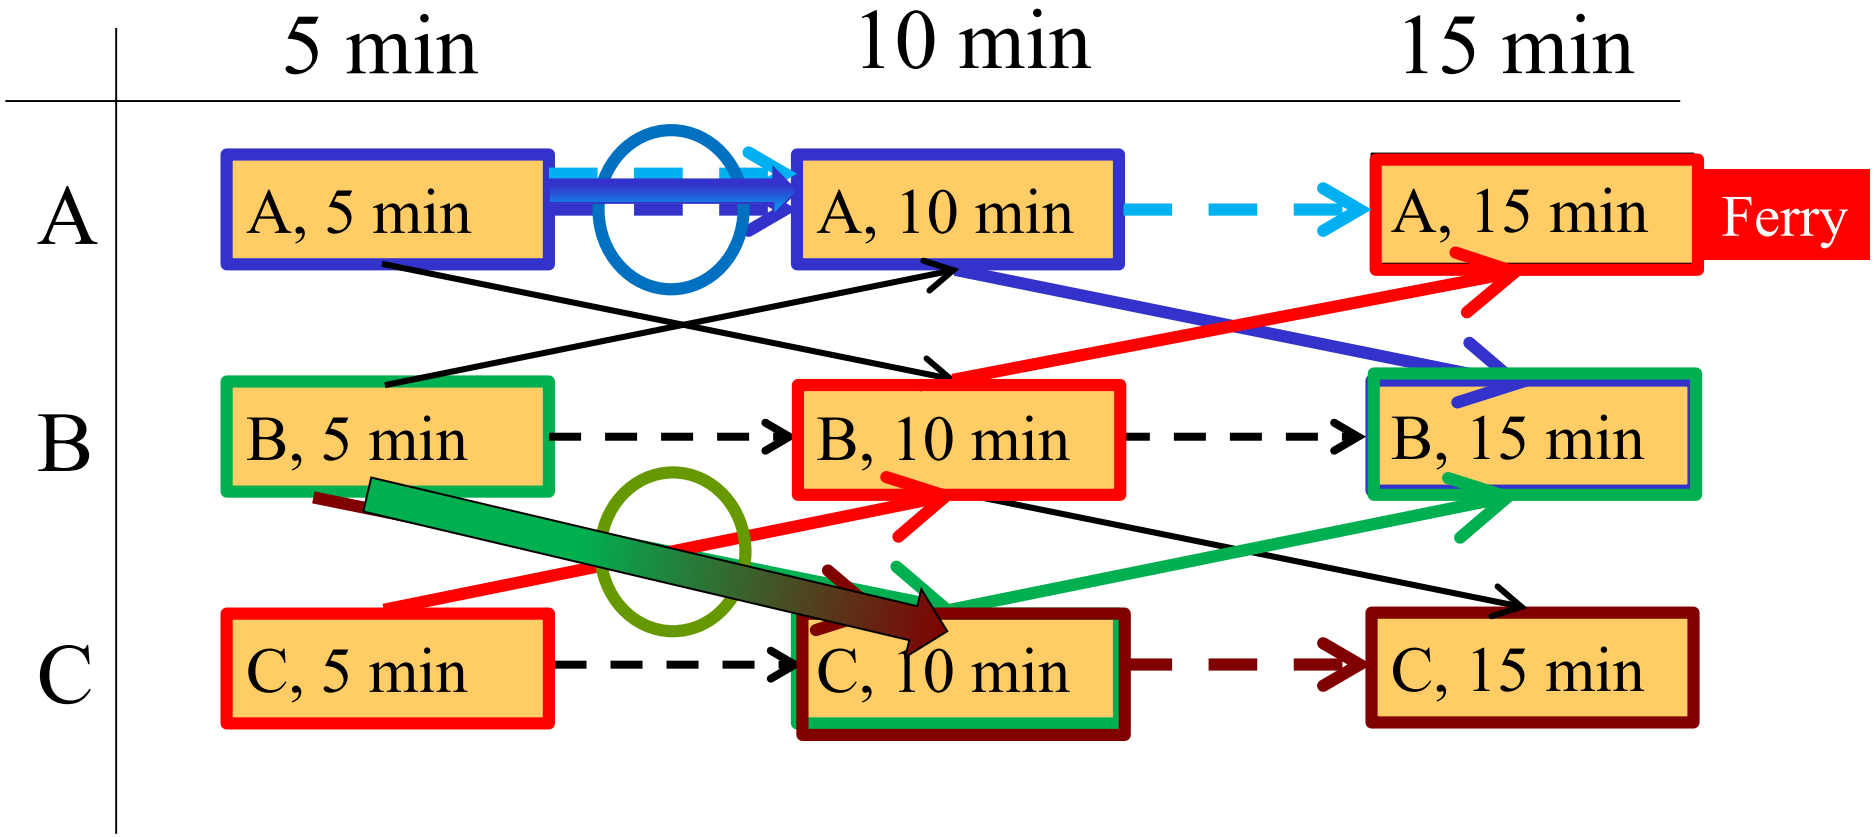

# Ferries: Scale-up with Marginals Over Separable Segments Significant Speedup

- *Obtain marginal probabilities over segments*

$N^2.T$ *variables*

$N^T$ *variables*

Extract: [illegible](B,5), (C, 10), (C,15)]) = 0.17

[illegible](B,[illegible]), (C,10), (B,15)]) =0.13

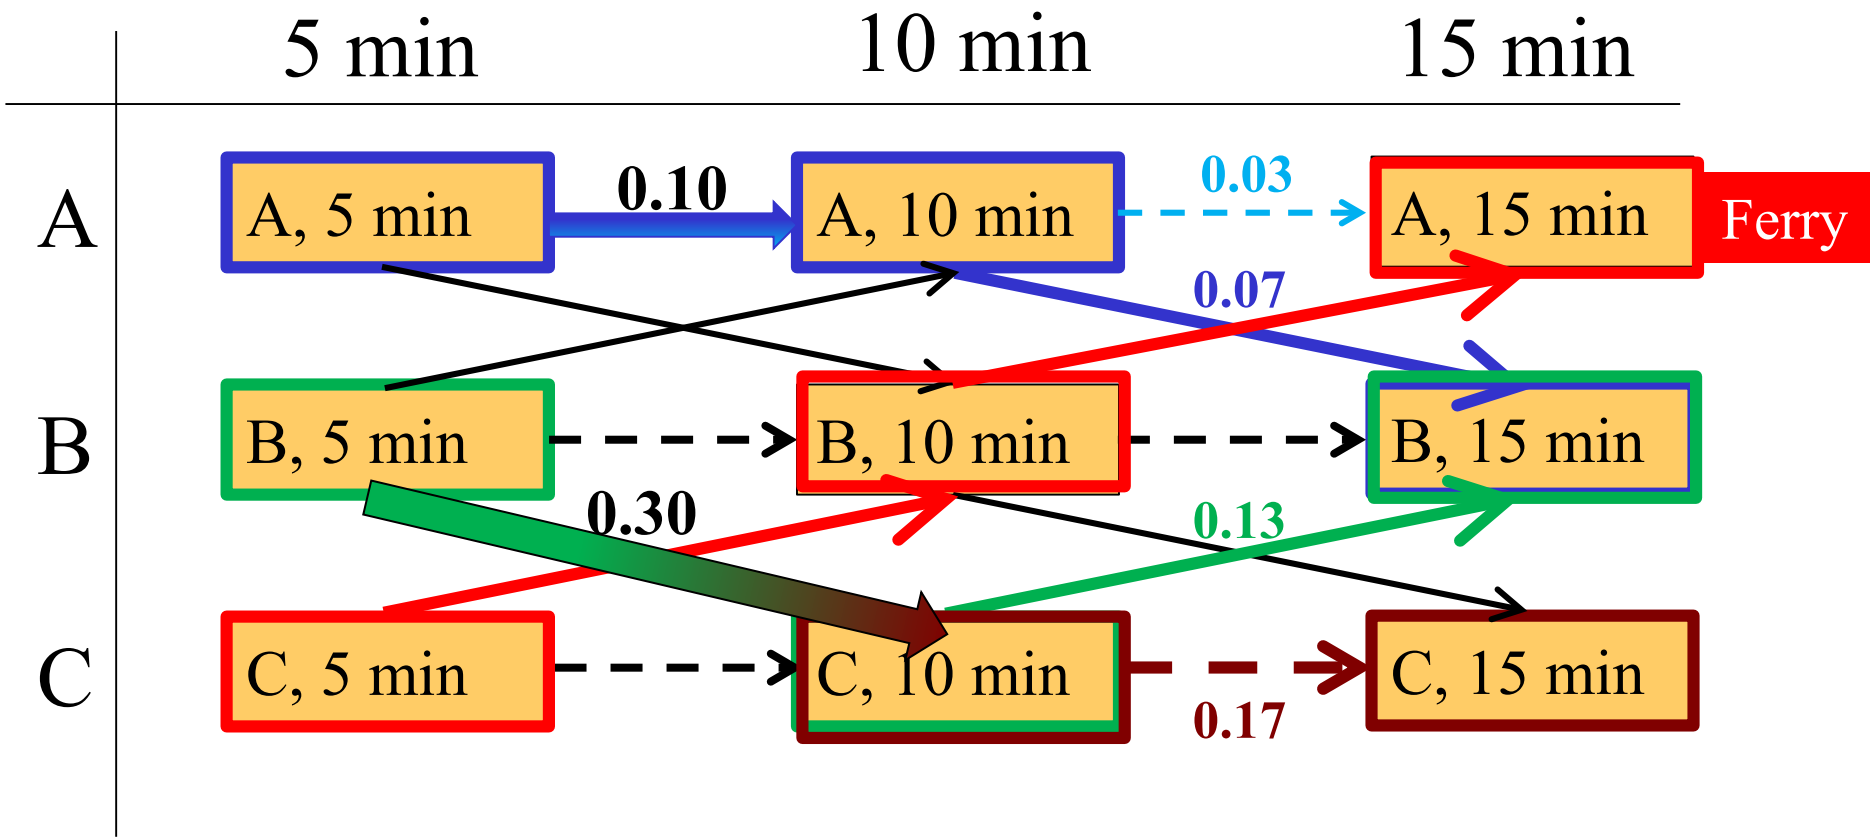

NOT VETTED BY CNN

3 of 5

g+1 Tweet Share Favorite

99 VIEWS | 0 COMMENTS | 6 SHARES

# U.S. Coast Guard protects the Staten Island Ferry: I feel safe!

By shortysmom | Posted September 8, 2013 | Staten Island, New York

Share on Facebook 0

**About this iReport**

- Not vetted for CNN

Posted September 8, 2013 by
shortysmom | Follow
Location
Staten Island, New York

I ride the Staten Island Ferry on a daily basis to and from work. We ferry riders have our own personal protectors in the form of the U.S. Coast Guard. The

# Outline: "Security Games" Research (2007-Now)

**Research Ch llenges:**
**Scale-up, U**

Airports

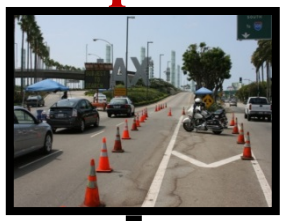

Trains

Environment

2007

2013-

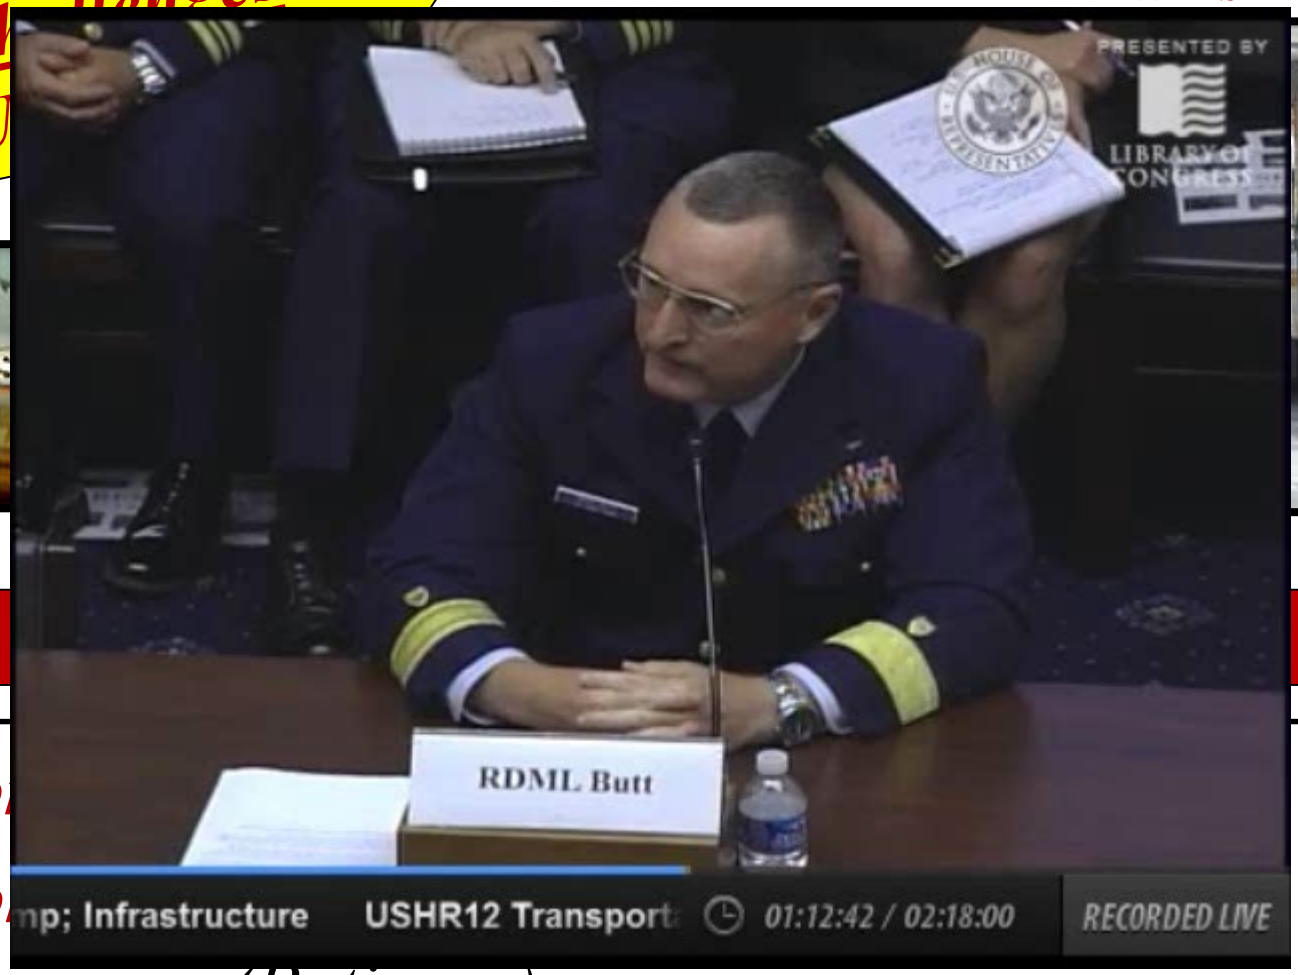


*Evaluatio*

*Evaluatio*

*(Patience)*

# TRUSTS: Frequent adversary interaction games
# Patrols Against Fare Evaders

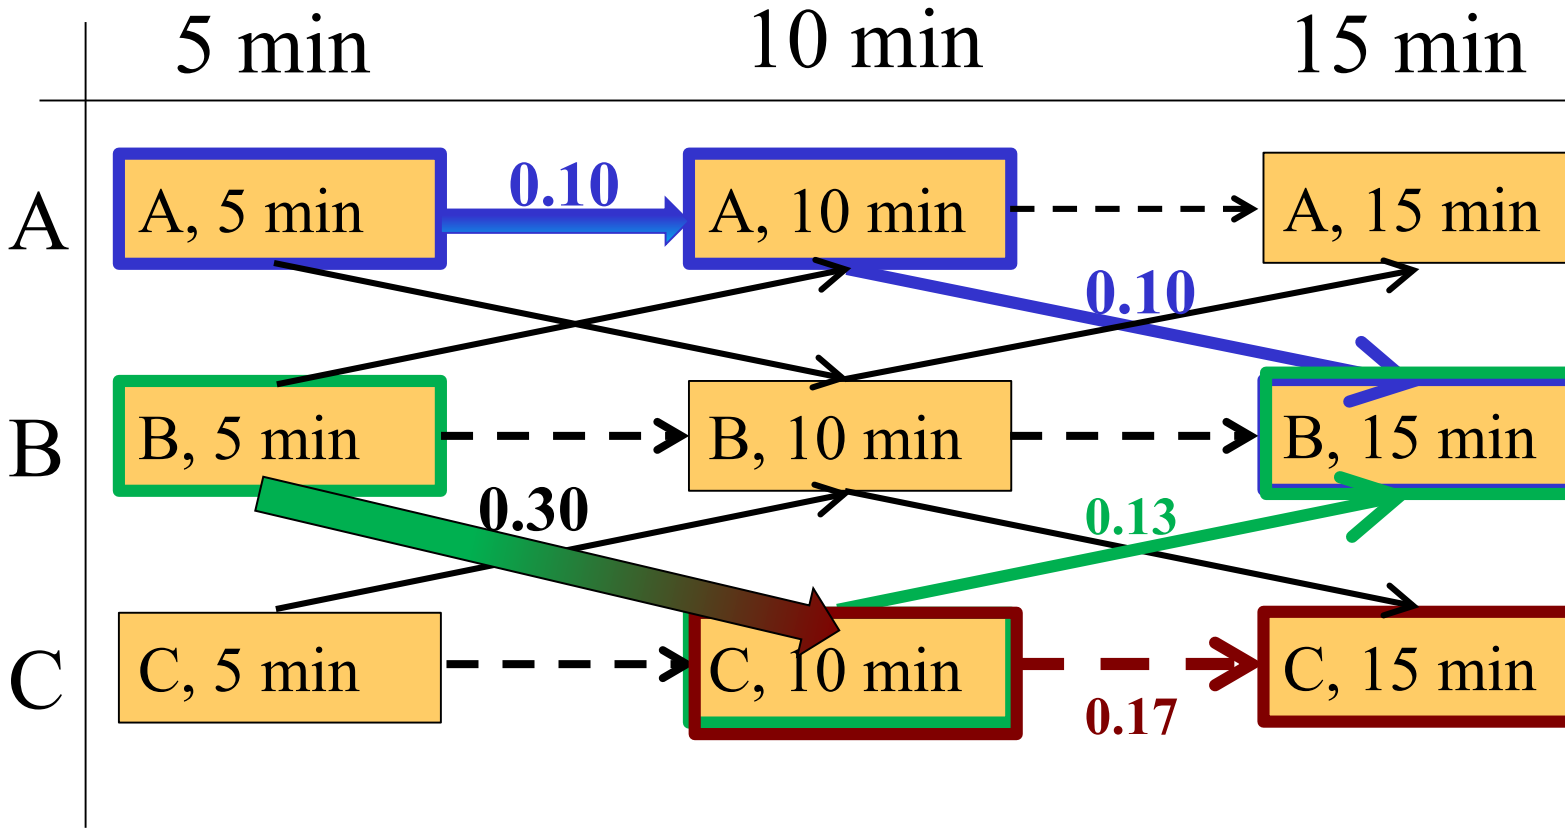

# TRUSTS: Frequent adversary interaction Uncertainty in Defender Action Execution

*Jiang* *Delle Fave*

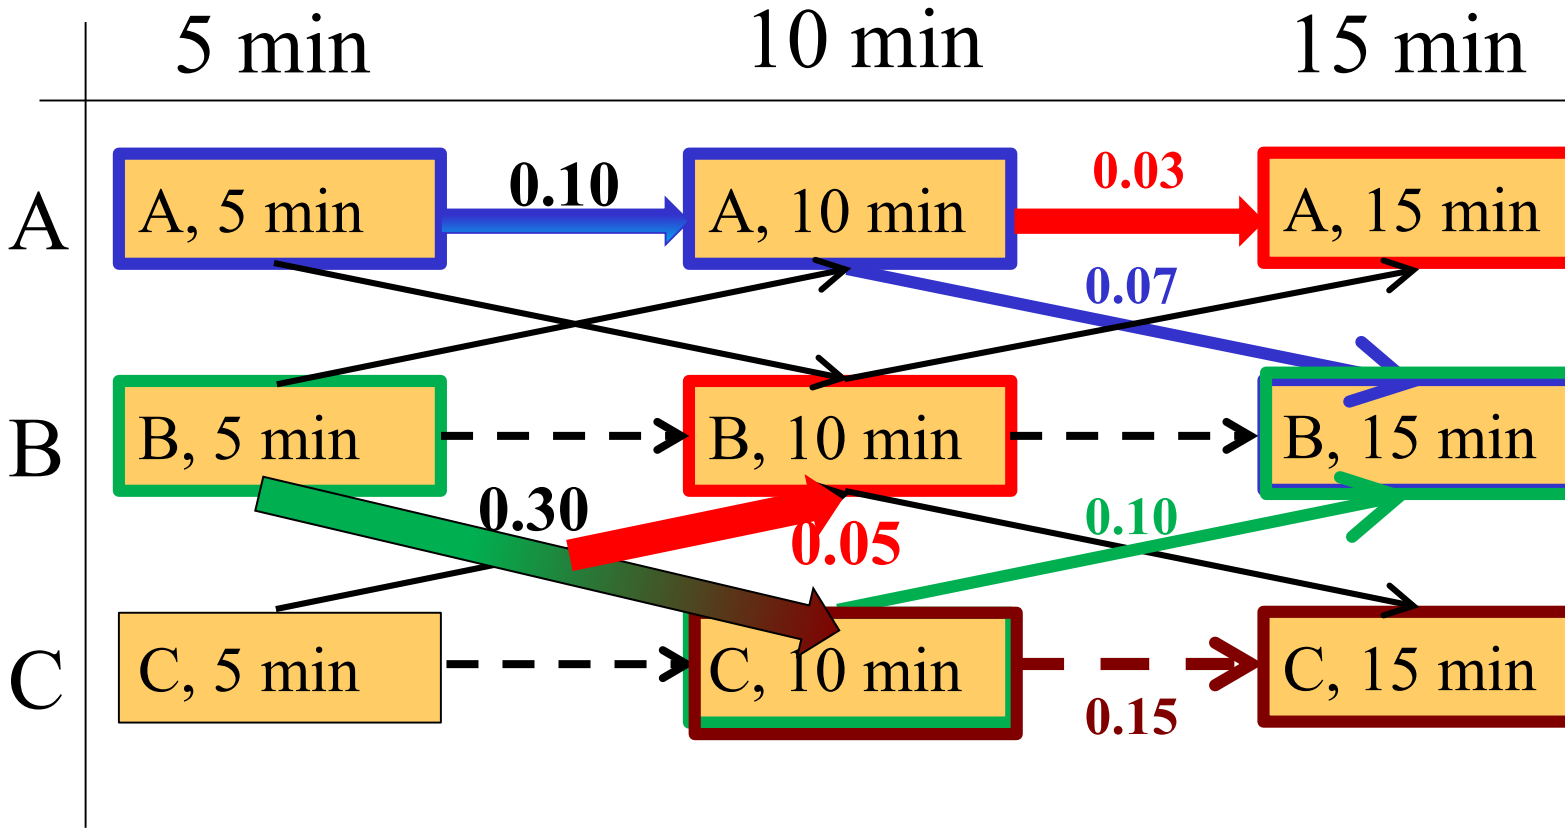

# TRUSTS: Frequent adversary interaction Uncertainty in Defender Action Execution

*Jiang* *Delle Fave*

- Markov Decision Problems *in Security games*

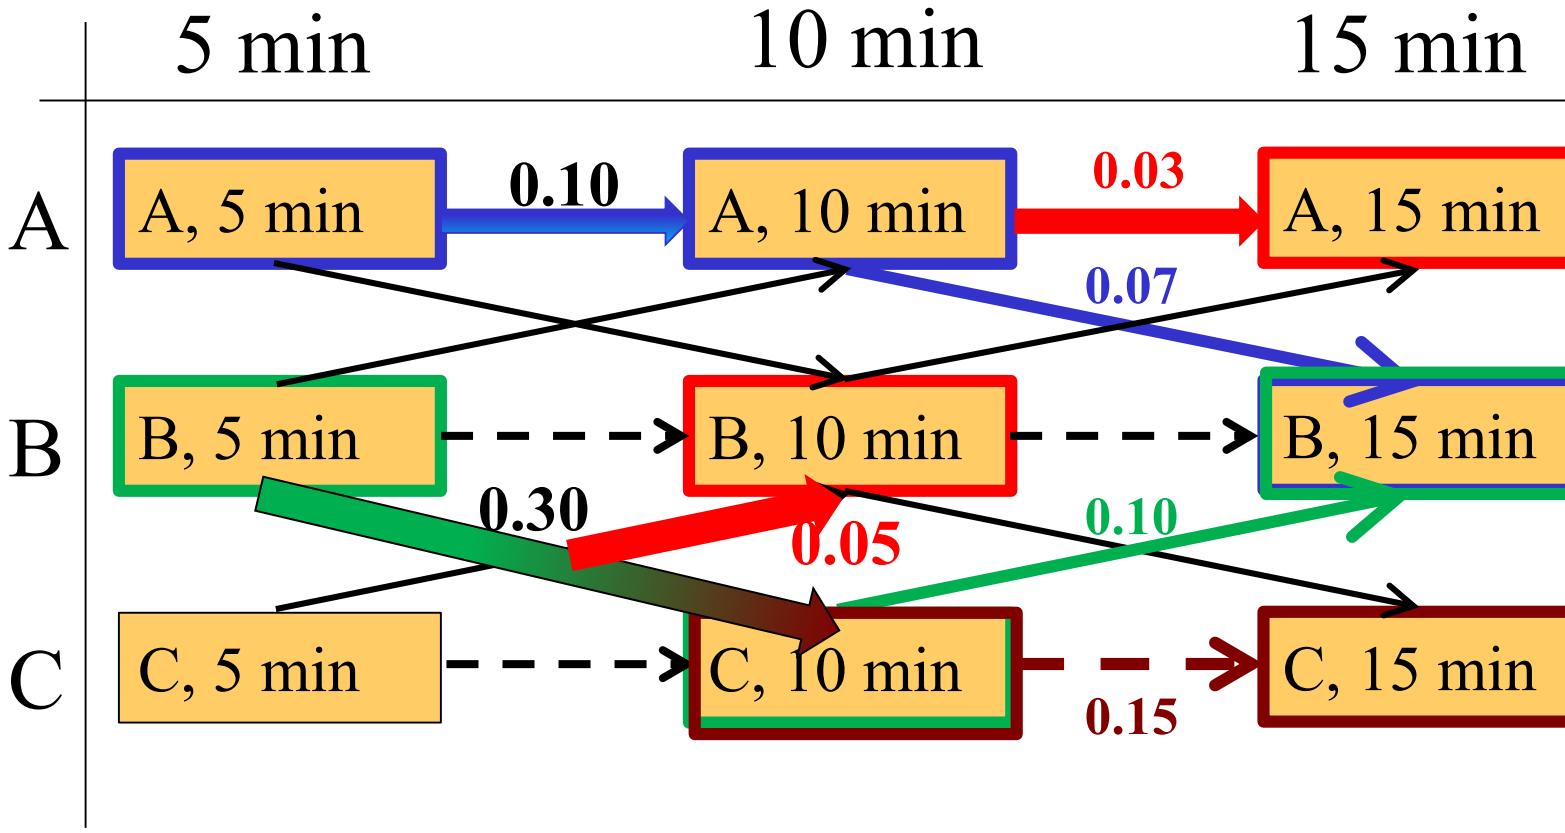

# Urban Transportation Security: MDPs & DEC-MDPs in Security Games

Shieh

Zhang

## *COPS*: LA Metro System
(Against Opportunistic Crime)

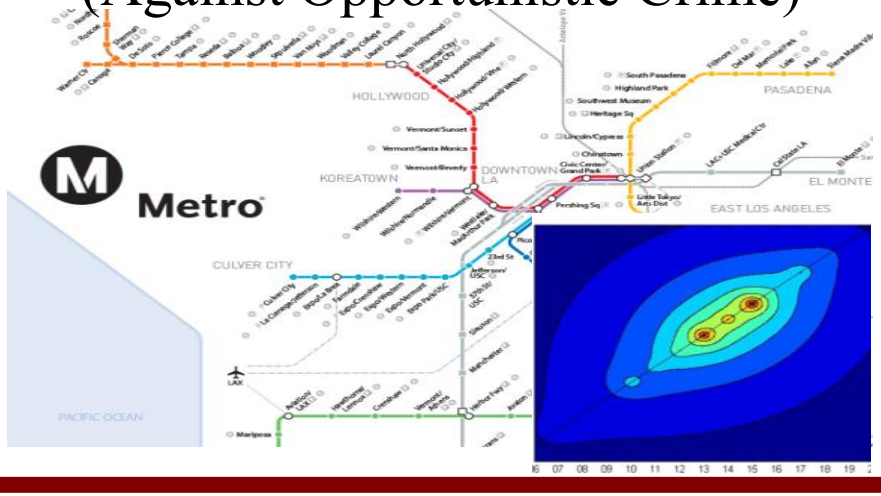


## *STREETS*: Singapore Roads
(Against Reckless Driving)

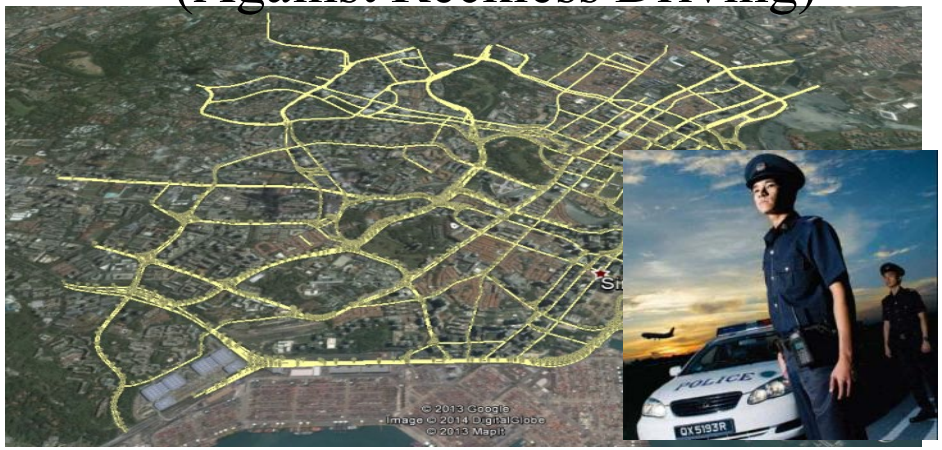


## *DEC-MDPs in Security Games*: Teamwork + Uncertainty

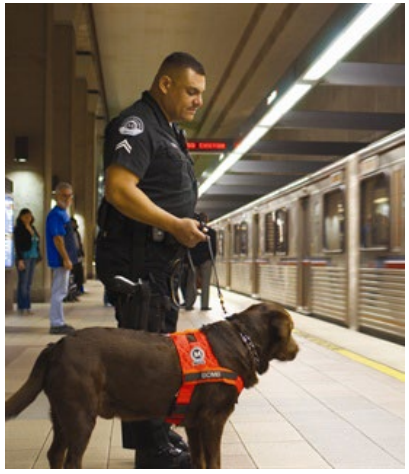

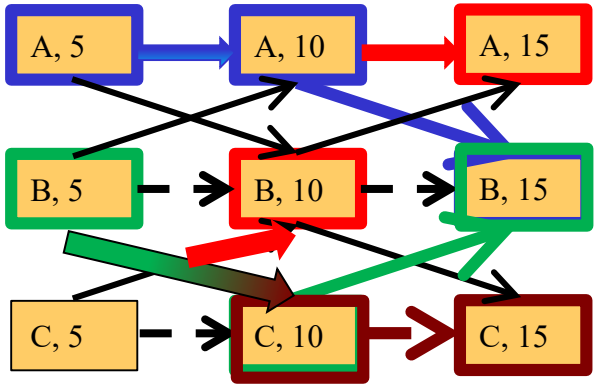

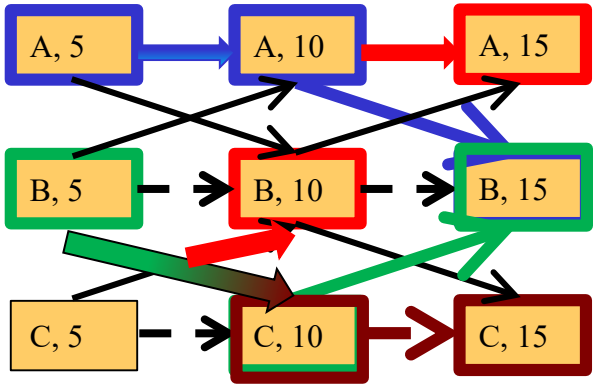

# Uncertainty Space Algorithms: Bayesian and Robust Approaches

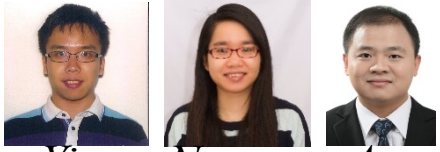


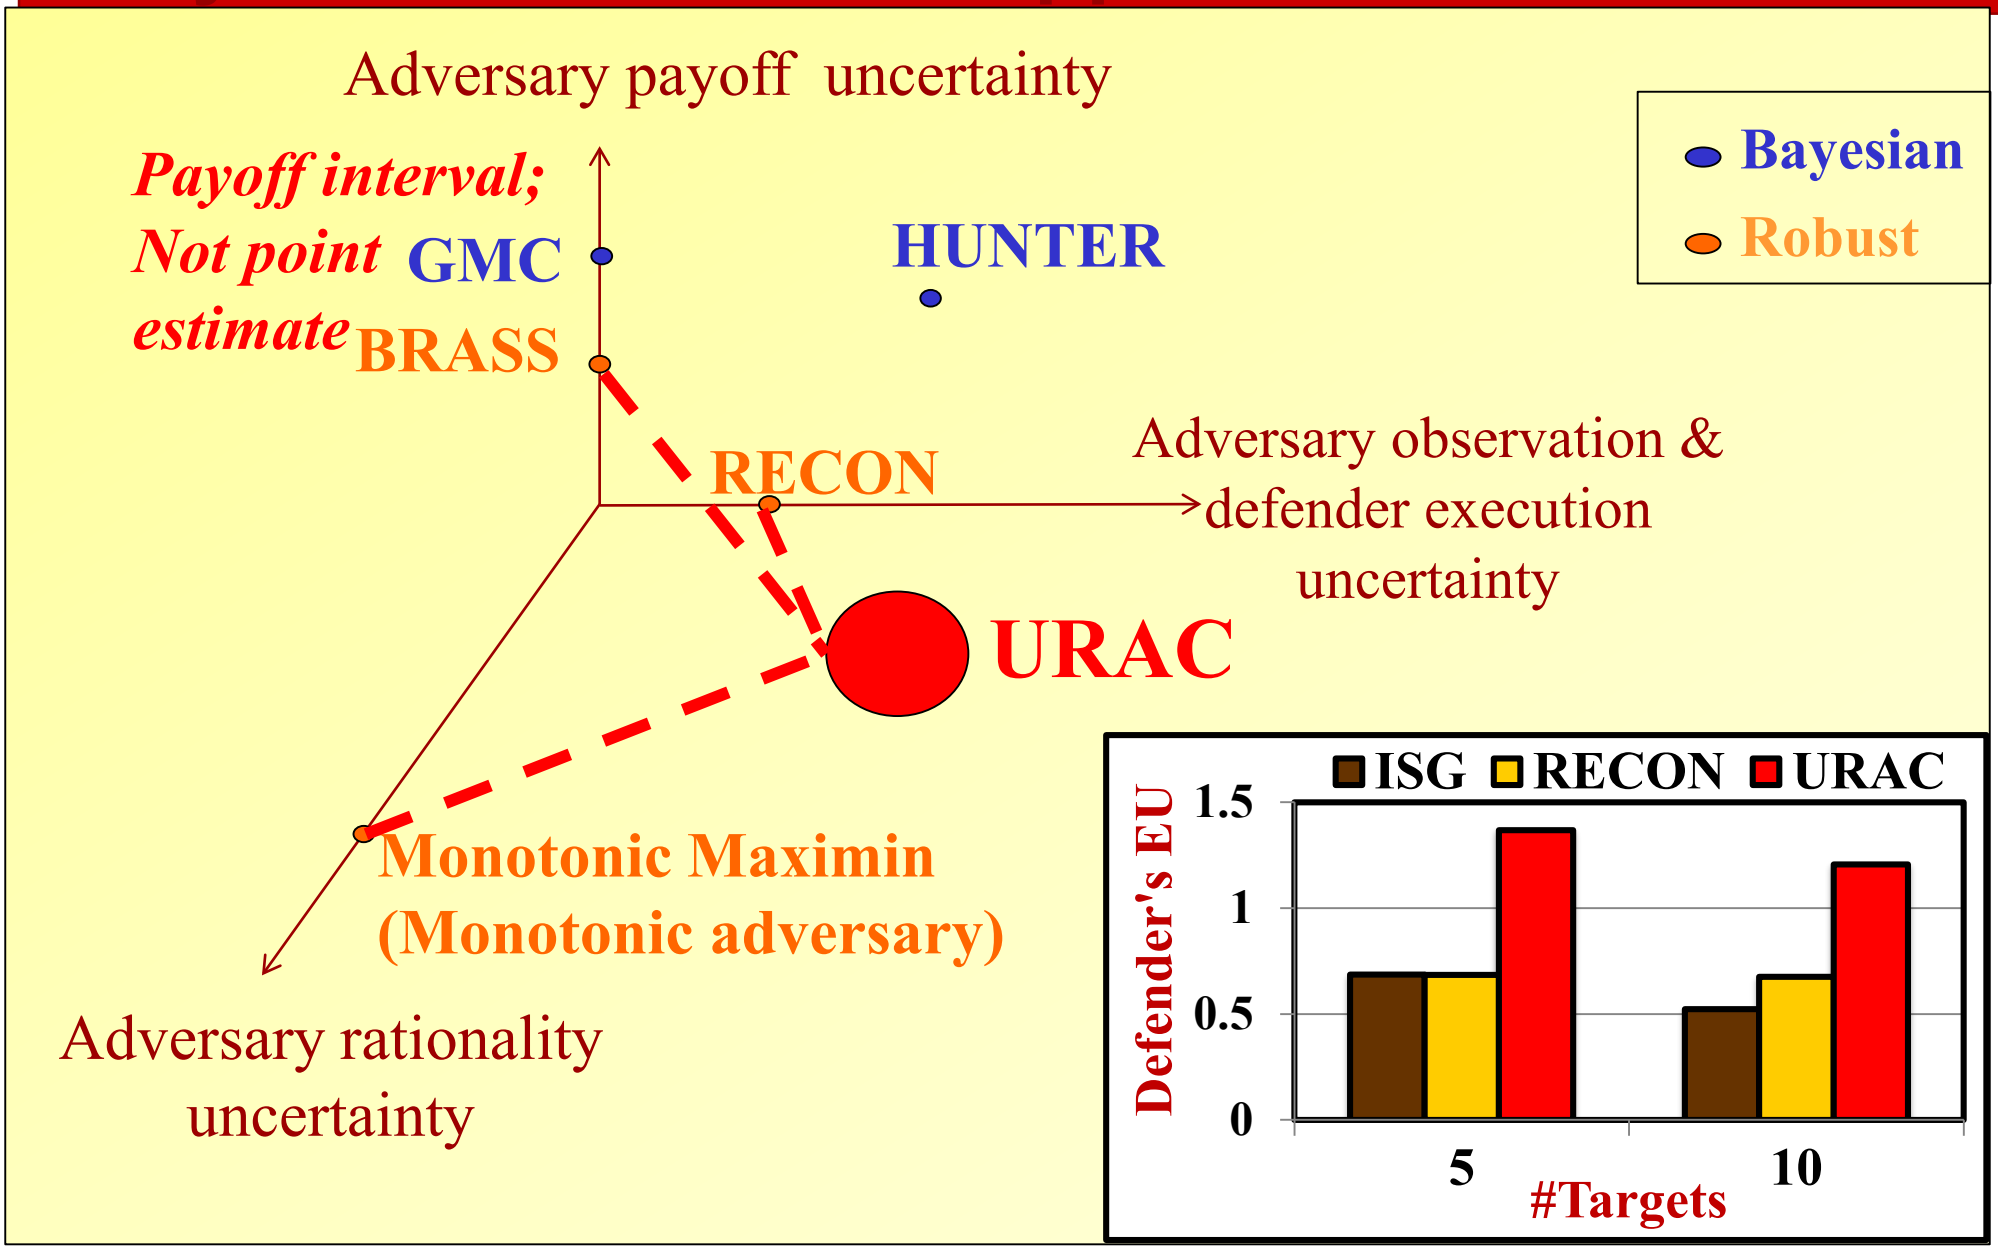

# Security Games, Environmental Crime & Bounded Rationality

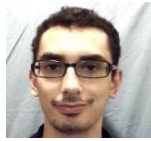

*Yang* *Ford* *Brown*

## Wildlife Queen Elizabeth National Park Uganda

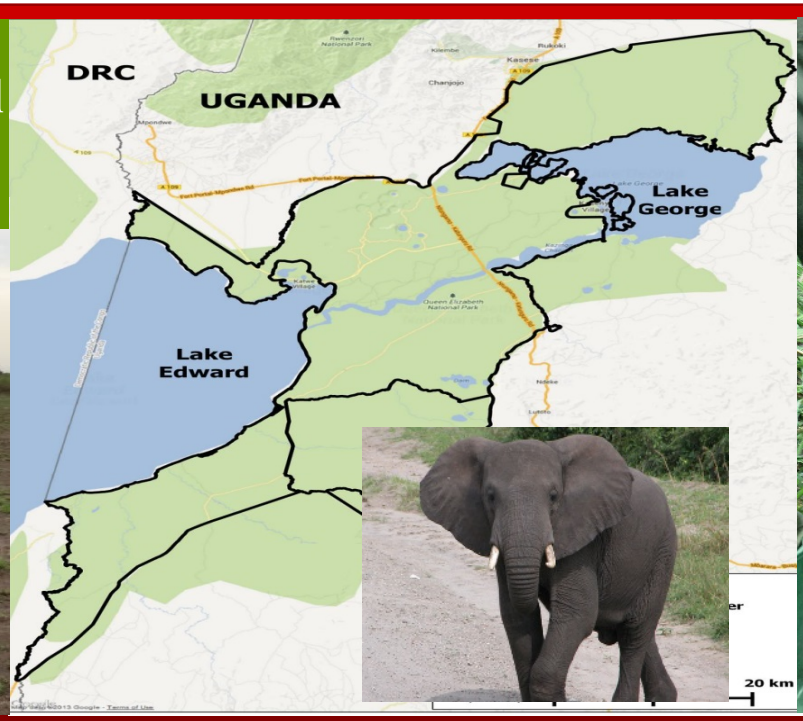

## Nakai Nam Theun Forest Area, Laos

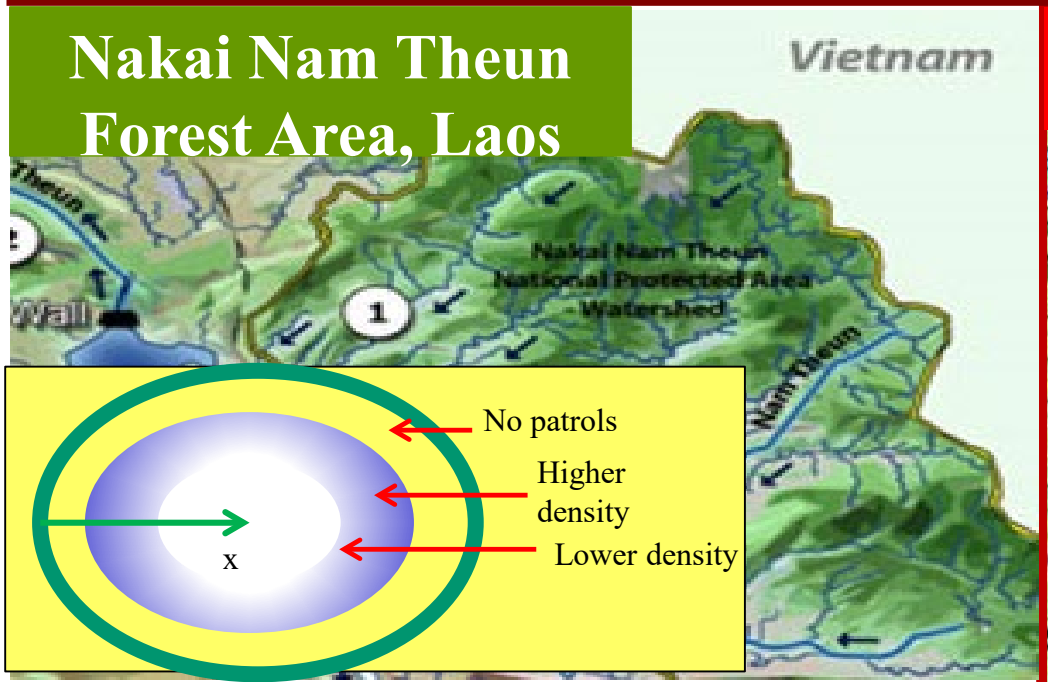


## Fishery Gulf of Mexico

# Uncertainty in Adversary Decision: Bounded Rationality
# Human Subjects as Poachers

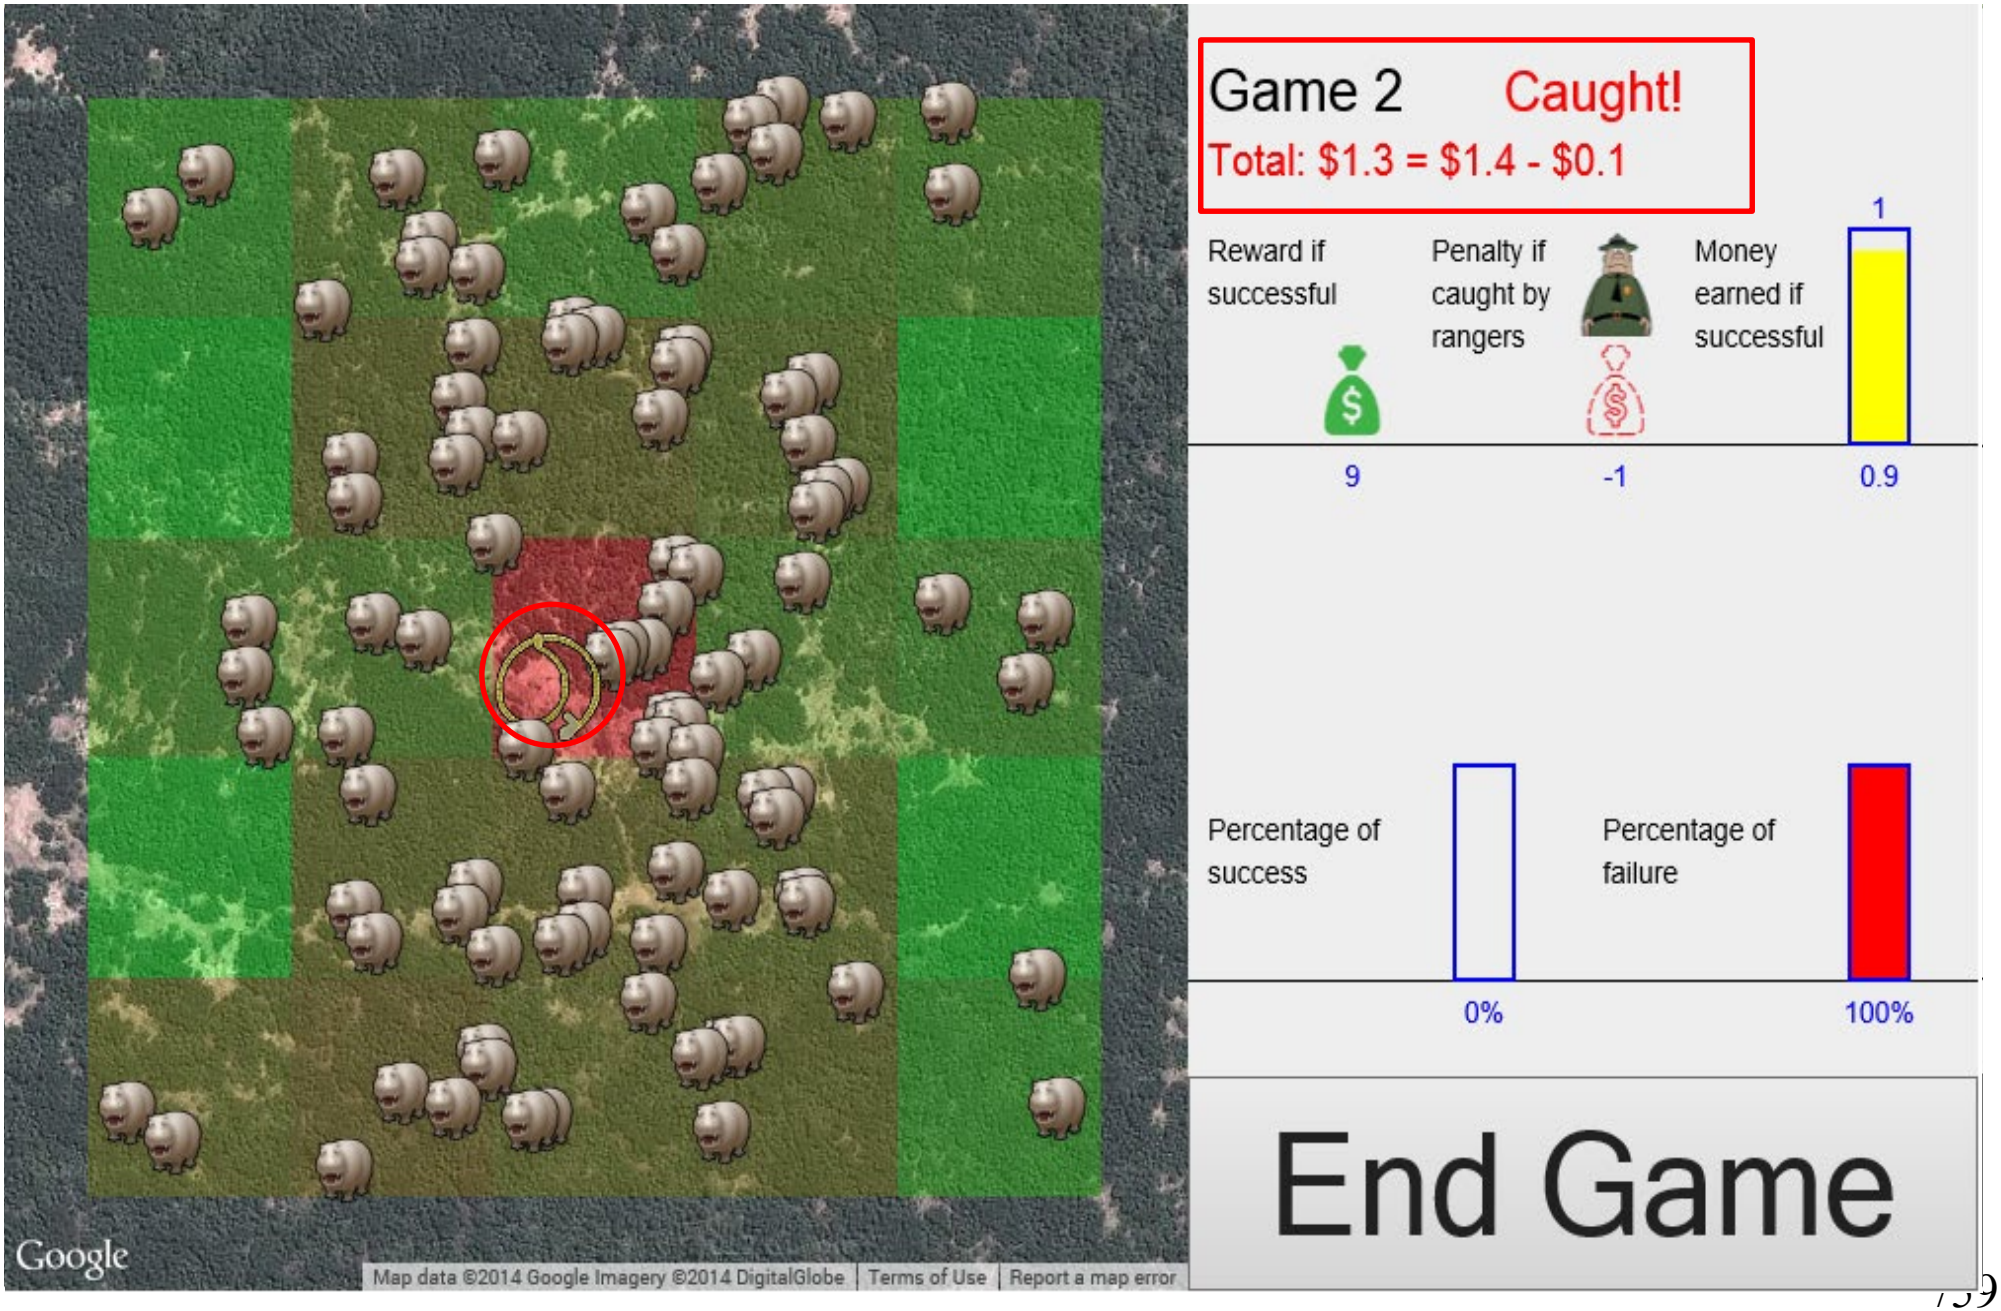

# Uncertainty in Adversary Decision[2009] Human subjects: Anchoring, ε-Optimality*

*Pita*

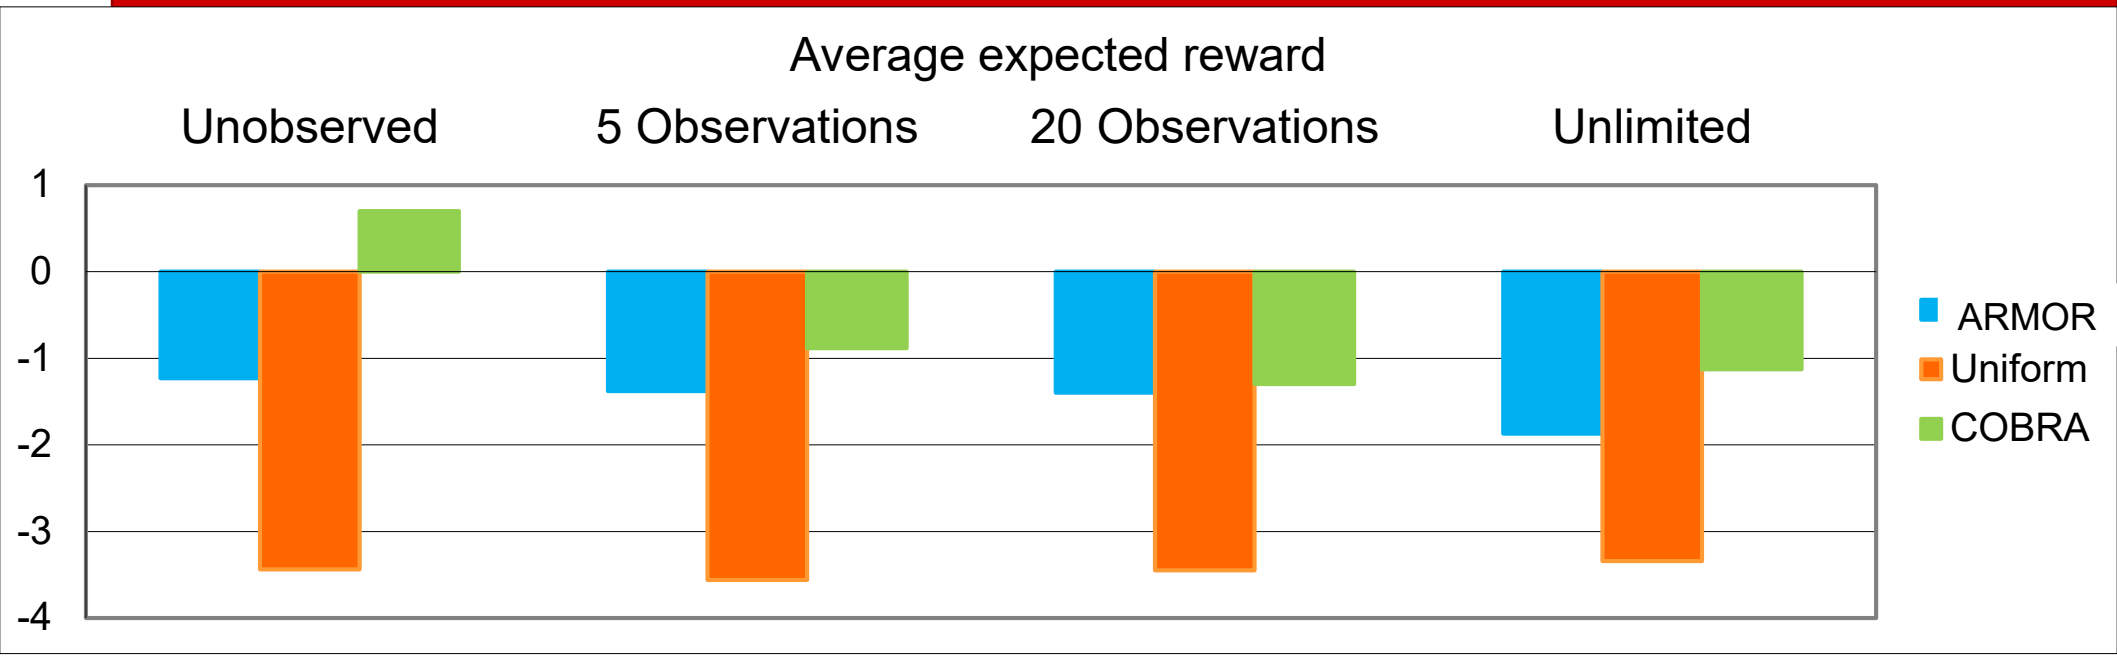


- **ARMOR:** Outperforms uniform random
- **COBRA:** $\max_{x,q} \sum_{i \in X} \sum_{j \in Q} R_{ij} x_i q_j$

Anchoring

$$s.t. \quad x' = (1-\alpha)x + \alpha(1/|X|)$$

ε-optimality

$$\varepsilon(1-q_j) \le (a - \sum_{i \in X} C_{ij} x'_i) \le \varepsilon + (1-q_j)M$$

*With Sarit Kraus

# Quantal Response Model of Adversary [2011]

## *Not Maximize Expected Utility*

Yang

***Quantal Response: Stochastic choice, better choice more likely***

$$\text{Adversary's probability of } \textbf{choosing target j} = \frac{e^{\lambda\cdot(EU^{adversary}(x,j))}}{\sum_{j'=1}^{T} e^{\lambda\cdot(EU^{adversary}(x,j'))}}$$

$$\max_x \sum_{j=1}^{T} \frac{e^{\lambda.EU^{adversary}(x,j)}}{\sum_{j'=1}^{T} e^{\lambda.EU^{adversary}(x,j')}} EU^{defend}(x,j)$$

$$s.t. \sum_t x_t = K; \quad 0 \le x_t \le 1$$

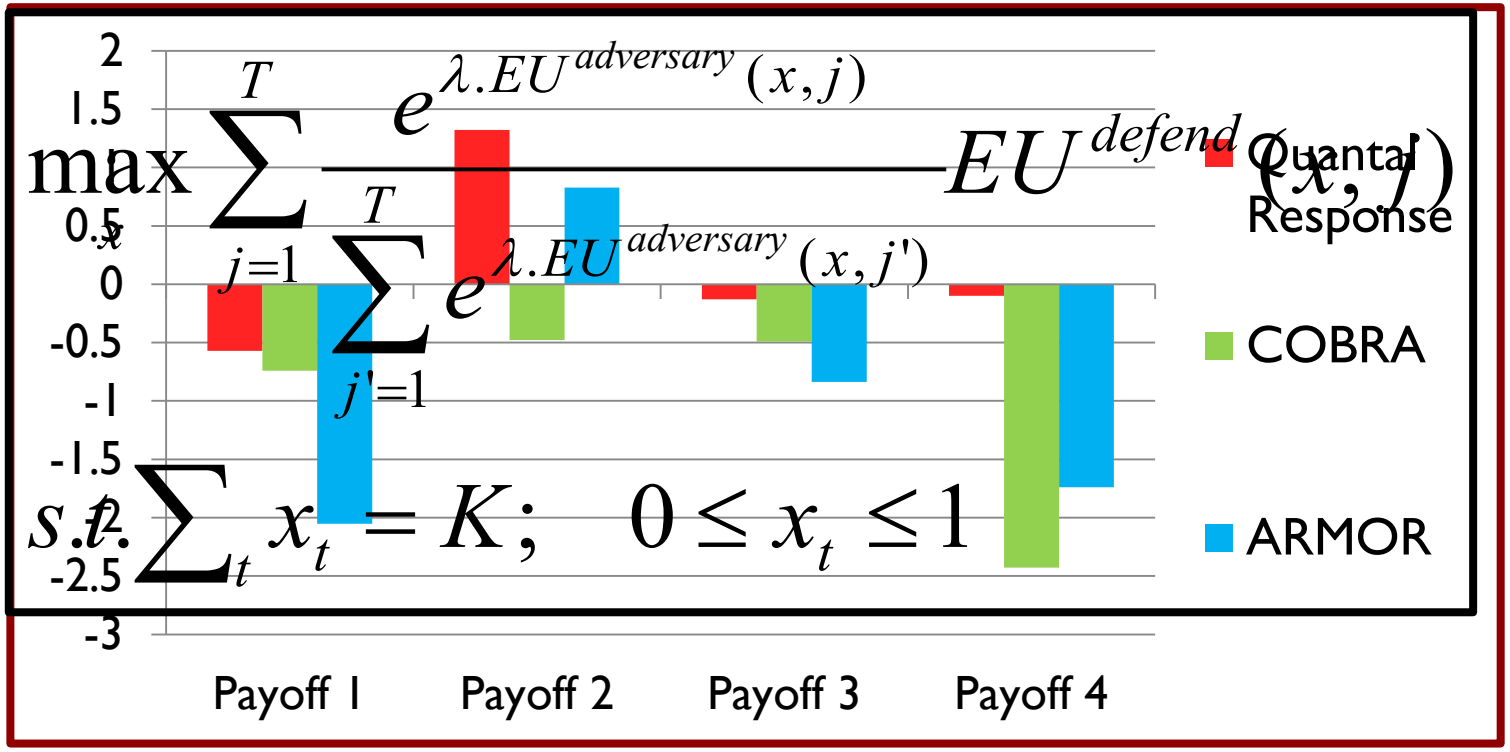

# Uncertainty in Adversary Decision [2012]
# Robust vs Modeling Adversaries*

*Robustness: Bound loss to defender; Not model attacker via QR*

β * (Adversary's utility loss if deviates from optimal) >= (Defender's utility loss due to adversary deviation)

**Results on 100 games**

| Robust wins | Draw | QR wins |
|---|---|---|
| 42 | 52 | 6 |

*Defeating Robust: Learned subjective utility*

$$q_j = \frac{e^{\lambda \cdot SEU^{adversary}(x,j)}}{\sum_{j'=1}^{M} e^{\lambda \cdot SEU^{adversary}(x,j')}}$$

$$SEU^a(j) = w_1 \times capture\ prob$$
$$+ w_2 \times attack\ reward$$
$$+ w_3 \times attack\ penalty$$

**Results on 22 games**

| SU-QR wins | Draw | Robust |
|---|---|---|
| 13 | 8 | 1 |

**Results against security experts**

| SU-QR wins | Draw | Robust |
|---|---|---|
| 6 | 13 | 3 |

**With Sarit Kraus*

# PAWS: Protection Assistant for Wildlife Security[2014]
# Repeated Stackelberg Game

*Bayesian SUQR: Heterogeneous Poachers*

$$\frac{e^{SEU_i(W1,W2,W3)}}{\sum_k e^{SEU_k(W1,W2,W3)}}$$

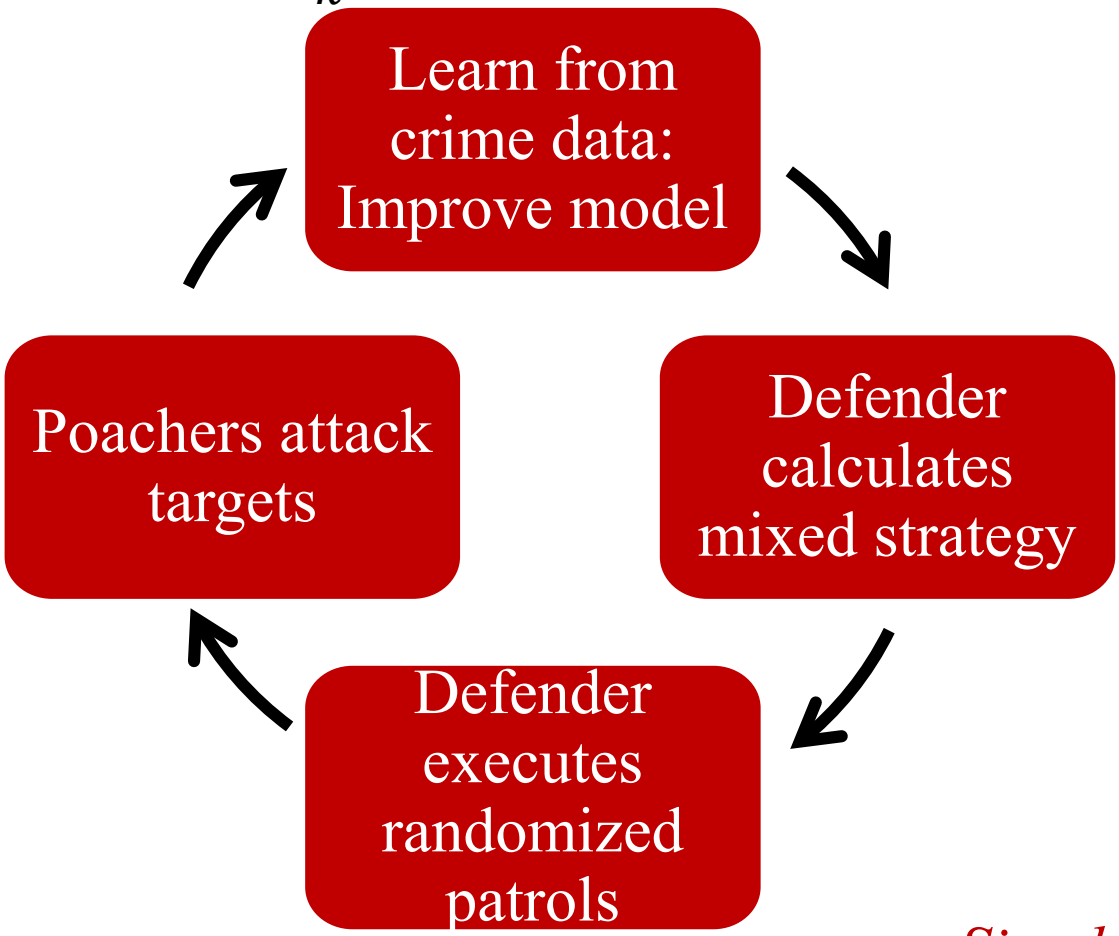


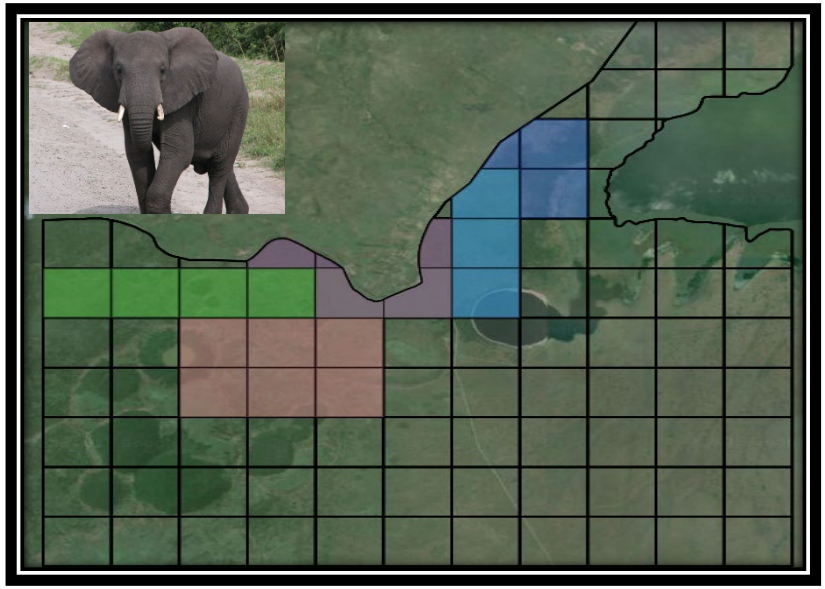

*Simulation*

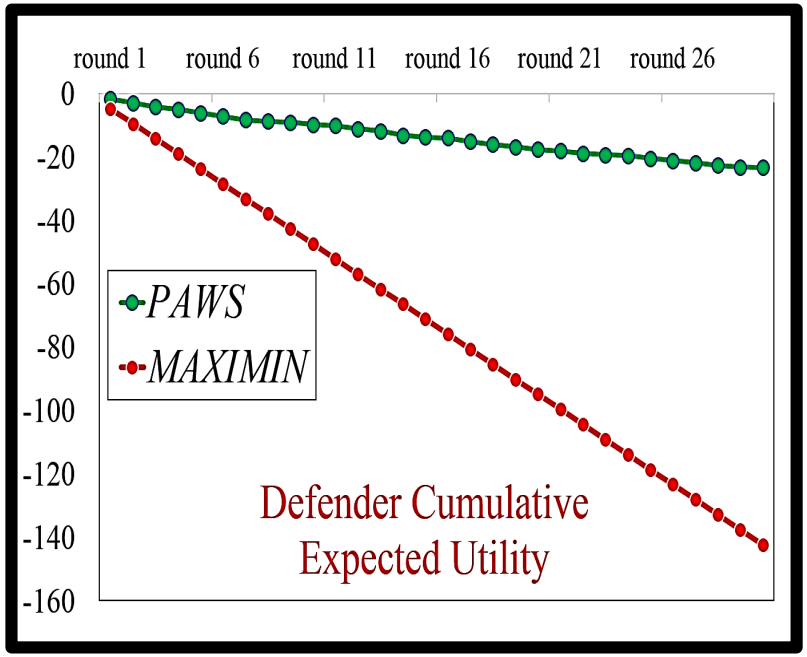

# PAWS Test: April 2014
# Trials in Queen Elizabeth National Park

Andrew Lemieux with rangers on PAWS patrol in Uganda

# Security Resource Optimization: Evaluating *Deployed Security Systems* Not Easy

- Game theory: Improvement over previous approaches
  - *Previous: Human schedulers or "simple random"*

| Lab Evaluation | Field Evaluation: Patrol quality Unpredictable? Cover? | Field Evaluation: Tests against adversaries |
|---|---|---|
| Simulated adversary | Compare real schedules | "Mock attackers" |
| Human subject adversaries | Scheduling competition | Capture rates of real adversaries |
| | Expert evaluation | |

# Why Does Game Theory Perform Better? Weaknesses of Previous Methods

- Human schedulers:
  - *Predictable patterns, e.g., US Coast Guard*
  - *Scheduling effort & cognitive burden*
- Simple random (e.g., dice roll):
  - *Wrong weights/coverage, e.g. officers to sparsely crowded terminals*
  - *No adversary reactions*
- Multiple deployments over multiple years: without us forcing them

# Lab Evaluation via Simulations: Example from IRIS (FAMS)

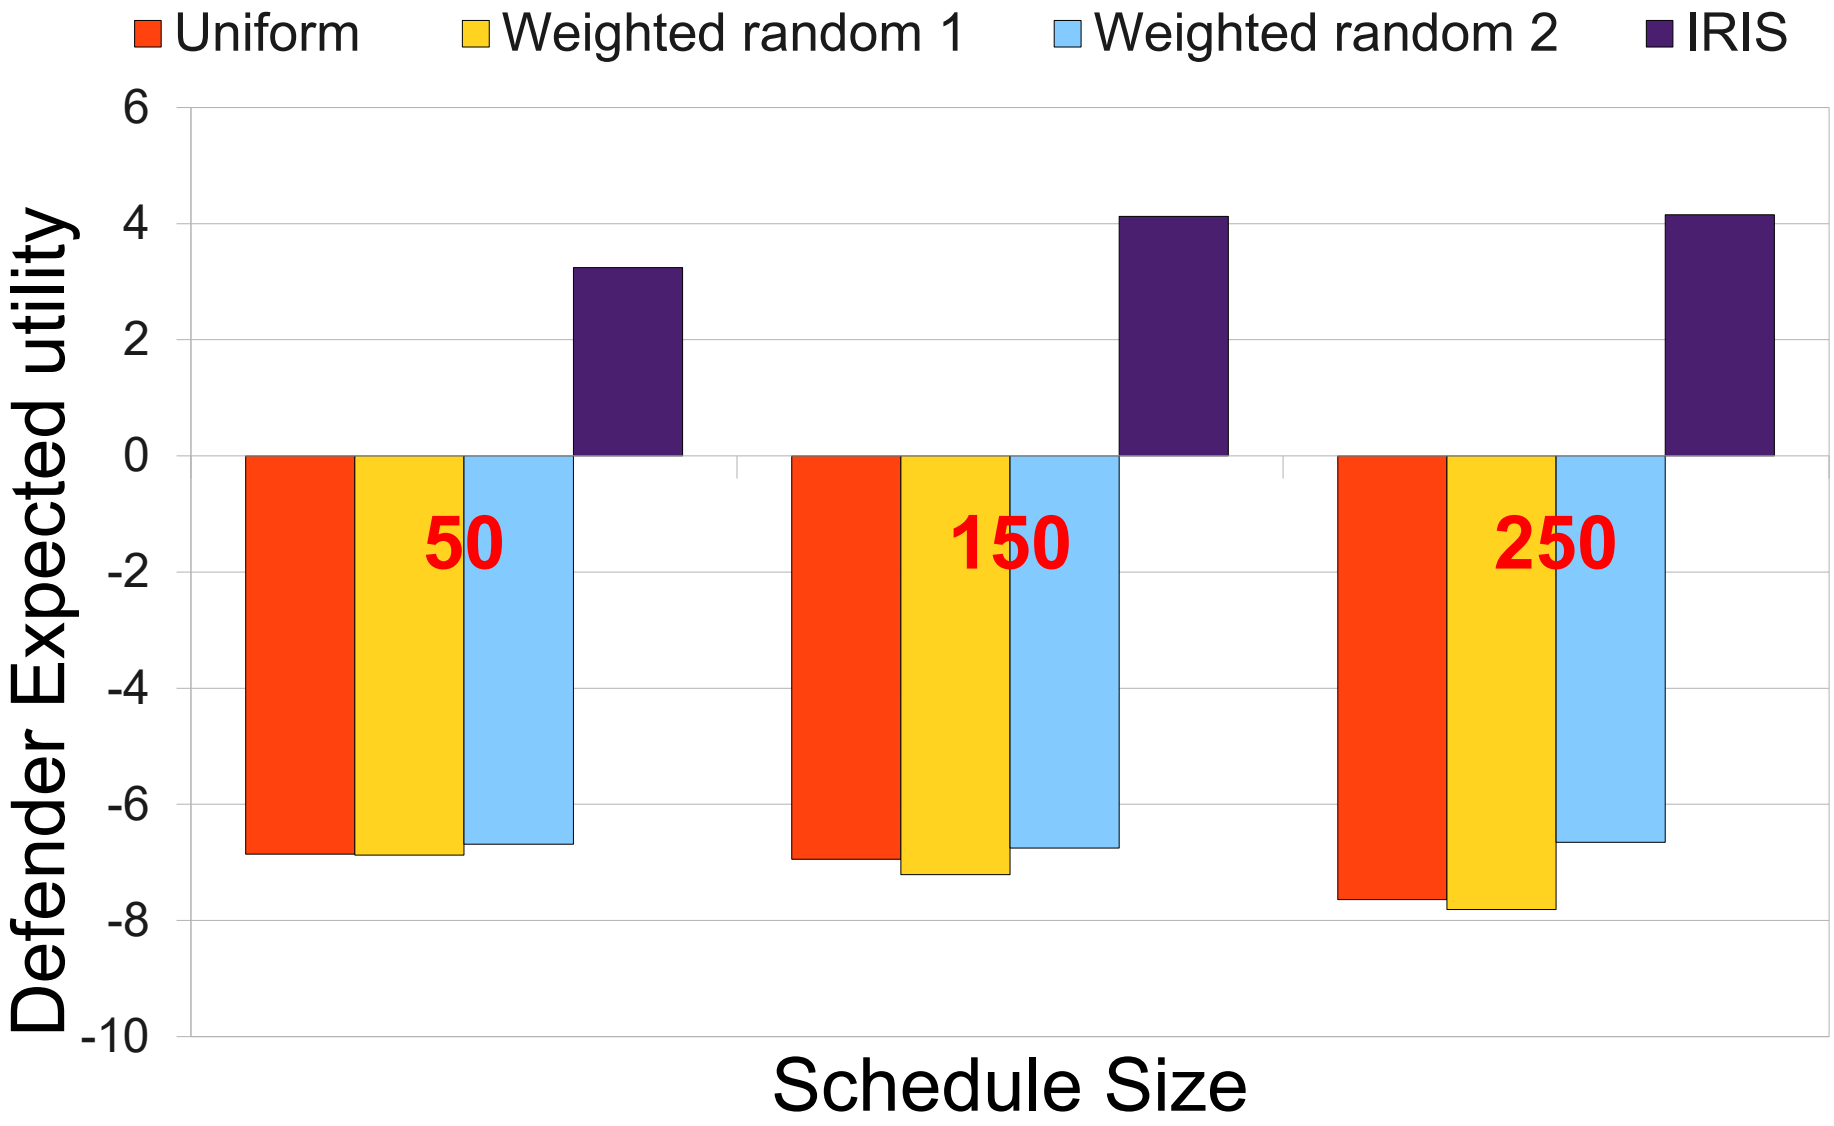

# Field Evaluation of Schedule Quality:

## Improved Patrol Unpredictability & Coverage

PROTECT (Coast Guard): *350% increase defender expected utility*

### Patrols Before PROTECT: Boston

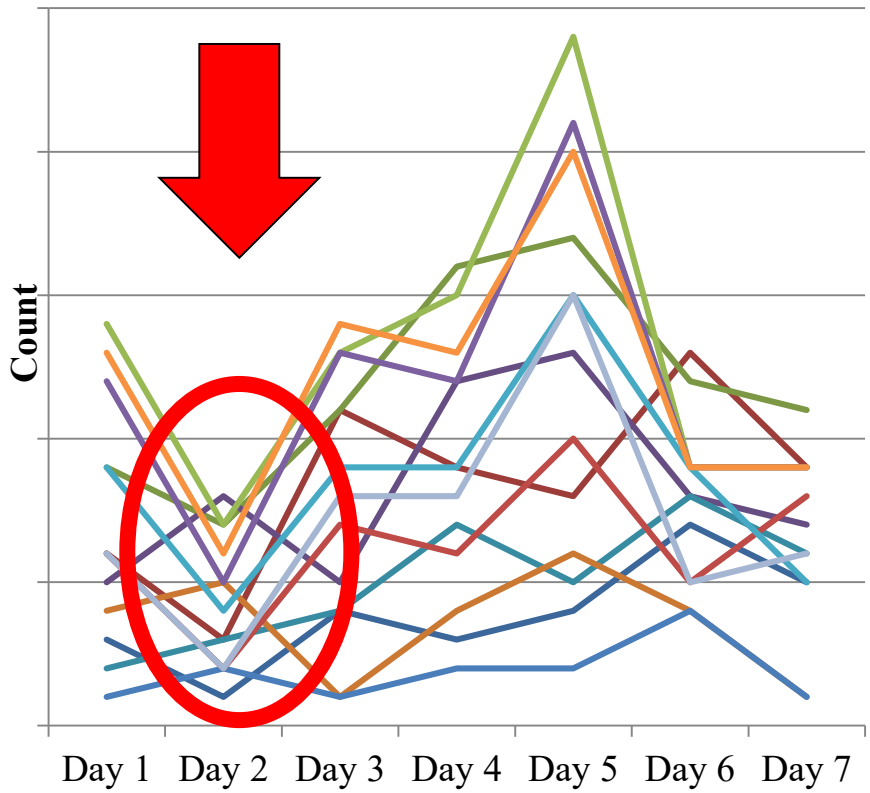


### Patrols After PROTECT: Boston

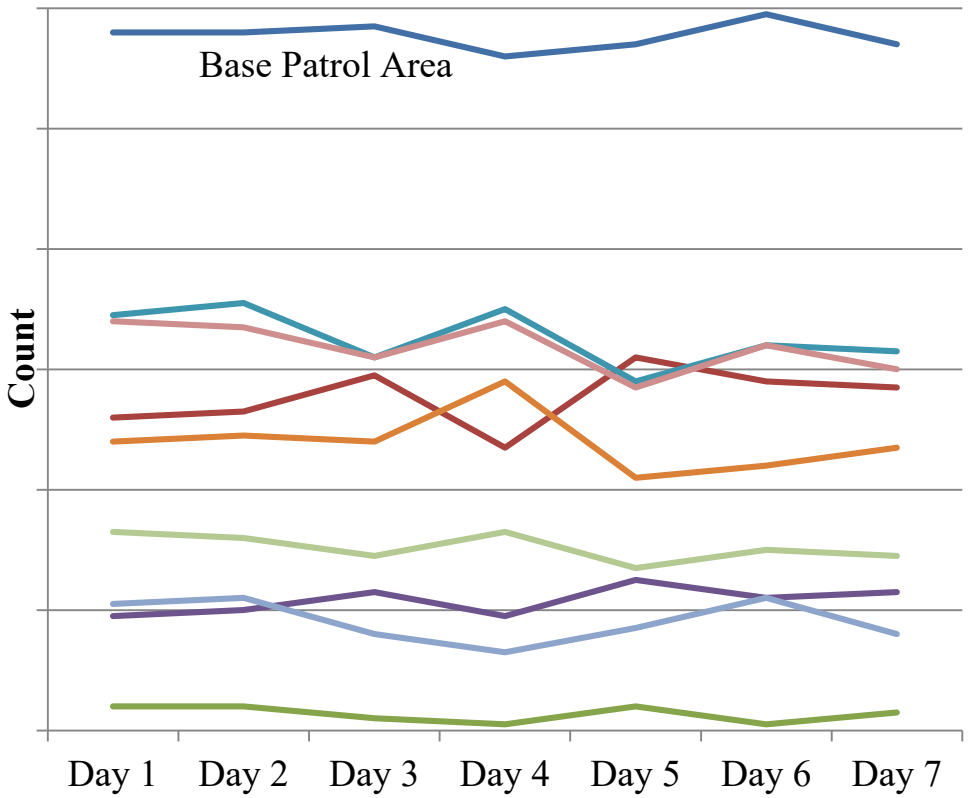

# Field Evaluation of Schedule Quality:

## Improved Patrol Unpredictability & Coverage for Less Effort

***IRIS for FAMS***: Outperformed expert human over six months

*Report:GAO-09-903T*

***ARMOR at LAX***: Savings of up to an hour a day in scheduling

# Field Evaluation: Human vs Game Theory Competition
## Counter-terrorism Patrol Scheduling

- 90 officers on LA Metro Trains

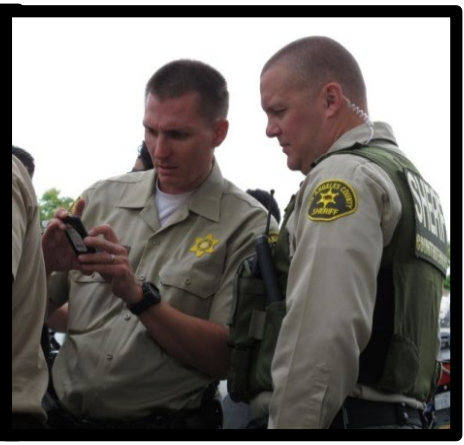

- Humans required significant effort
  - *Worse schedules than game theory*

- Observer's report on questions:

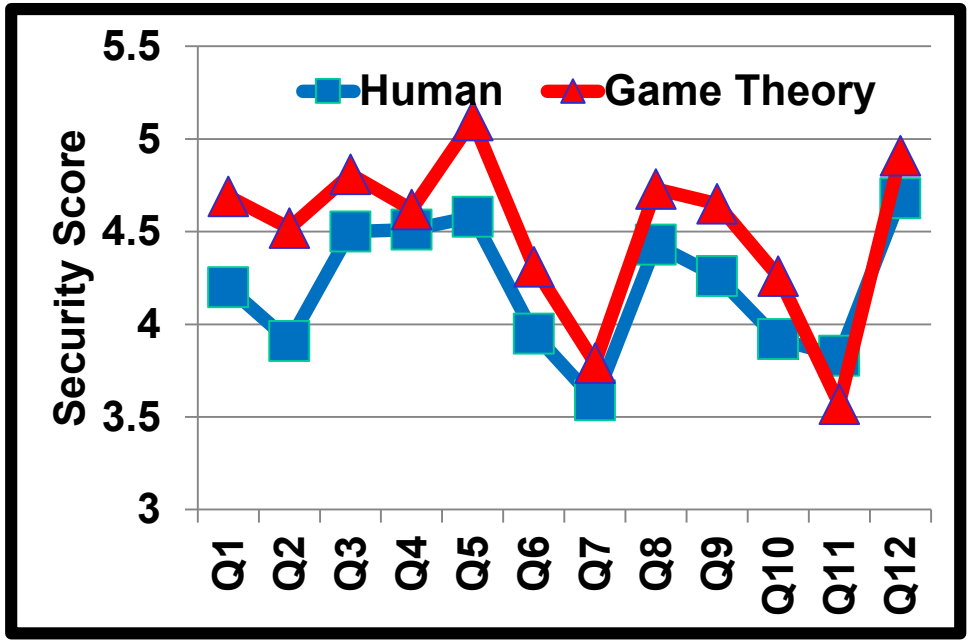

# Field Test Against Adversaries: Mock Attackers Example from PROTECT

- "Mock attacker" team deployed in Boston
  - *Comparing PRE- to POST-PROTECT: "deterrence" improved*

- Additional real-world indicators from Boston:
  - *Boston boaters questions:*
    - "..has the Coast Guard recently acquired more boats"
  - *POST-PROTECT: Actual reports of illegal activity*

# Field Tests Against Adversaries
# Computational Game Theory in the Field

## *Controlled*

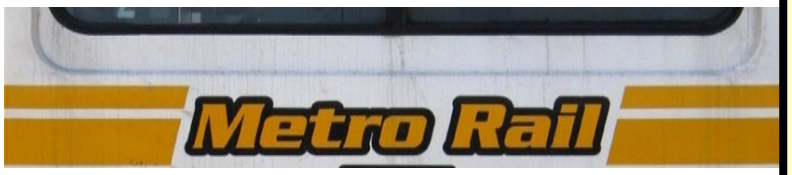


- *Game theory vs Random*
- *21 days of patrol*
- *Identical conditions*
- *Random + Human*

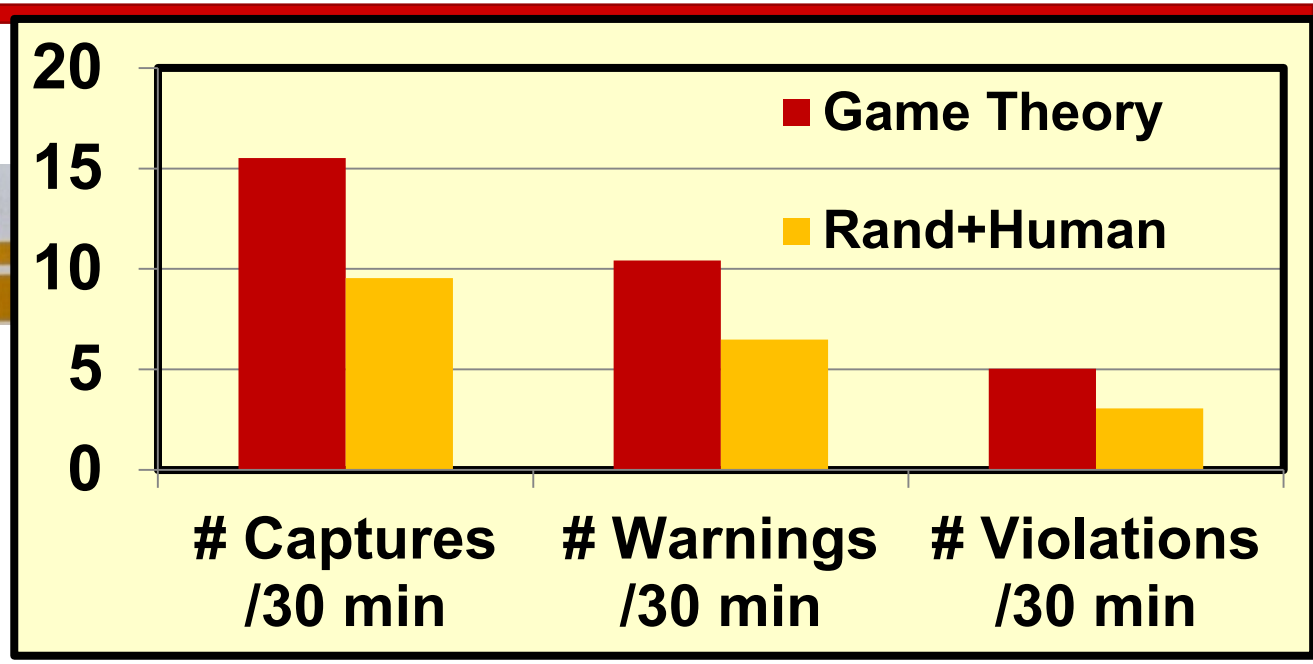


## *Not controlled*

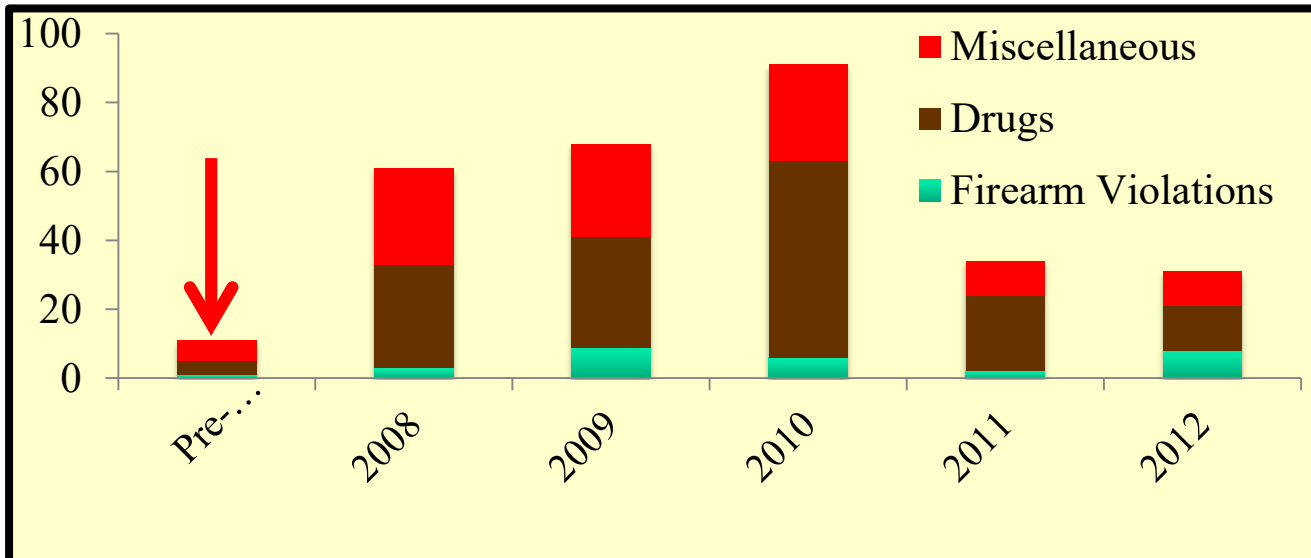

# Expert Evaluation
# Example from ARMOR, IRIS & PROTECT

## June 2013: Meritorious Team Commendation from Commandant (US Coast Guard)

## July 2011: Operational Excellence Award (US Coast Guard, Boston)

## September 2011: Certificate of Appreciation (Federal Air Marshals)

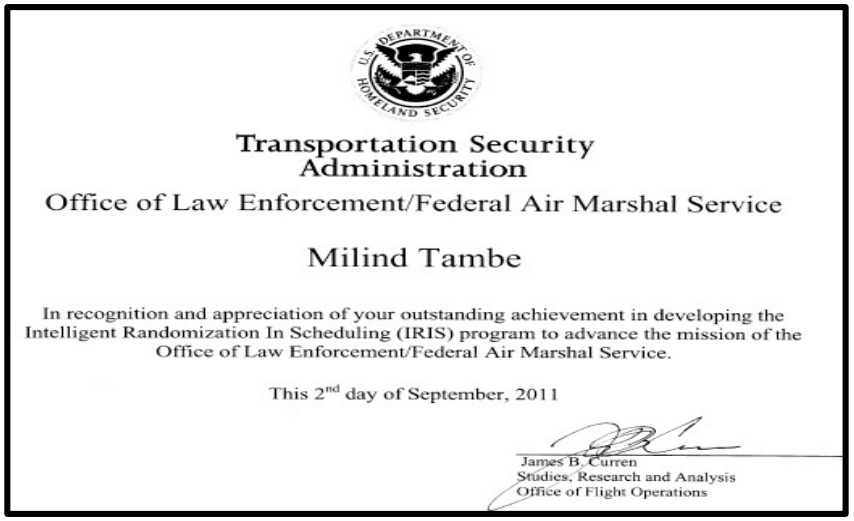

U.S. DEPARTMENT OF HOMELAND SECURITY

Transportation Security Administration

Office of Law Enforcement/Federal Air Marshal Service

Milind Tambe

In recognition and appreciation of your outstanding achievement in developing the Intelligent Randomization In Scheduling (IRIS) program to advance the mission of the Office of Law Enforcement/Federal Air Marshal Service.

This 2nd day of September, 2011

James B. Curren
Studies, Research and Analysis
Office of Flight Operations

## February 2009: Commendations LAX Police (City of Los Angeles)

# Summary: Security Games

Decision aids based on computational game theory in daily use

- Optimize limited security resources against adversaries

Applications yield research challenges: Science of security games

- Scale-up: Incremental strategy generation & Marginals
- Uncertainty: Integrate MDPs, Robustness, Quantal response

Current applications (wildlife security): AI challenge problems

- Global challenges: Merge planning/learning & security games

# Just the Beginning of "Security Games"….

- Follow on research & applications:

*Privacy audits*
*(Sinha IJCAI'13)*

*Software testing*
*(Kukreja ASE'13)*

*Sport event security*
*(Yin AAAI'14)*

*Singapore train*
*(Varakantham IAAI'13)*

*Exam questions*
*(Li IJCAI'13)*

- Our next steps:

***Startup:ARMORWAY***

***Game theory in the field***

**Thank you to sponsors:**

CREATE HOMELAND SECURITY CENTER

Transportation Security Administration

Army Research Office

# THANK YOU

tambe@usc.edu
http://teamcore.usc.edu/security