Towards a Science of Security Games:

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, Sara McCarthy, Thanh Nguyen, Praveen Paruchuri, Jonathan Pearce,James Pita, Yundi Qian, Aaron Schlenker, Eric Shieh, Jason Tsai, Pradeep Varakantham, Haifeng Xu, Amulya Yadav, Rong Yang, Zhengyu Yin, Chao Zhang

TEAMCORE
USC

Other collaborators:

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) Barney Long(WWF), Arnault Lyet (WWF), Kevin Leyton-Brown (UBC), F. Ordonez (U Chile), M. Pechoucek (CTU, Czech R), Ariel Procaccia (CMU), Martin Short (GATech), P. Stone (Univ of Teaxas), Y. Vorobeychik (Vanderbilt), ….

Global Challenges for Security:
Game Theory for Security Resource Optimization

Example Model: Stackelberg Security Games

Security allocation:

- Targets have weights
- Adversary surveillance

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:

- Targets have weights
- Adversary surveillance

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:

- Targets have weights
- Adversary surveillance

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 strategy:
  - *Increase cost/uncertainty to attackers*
- Stackelberg game:
  - *Defender commits to mixed strategy*
  - *Adversary conducts surveillance; responds*
- Stackelberg Equilibrium: Optimal random?

Adversary

Defender

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

Security Games: Research & Applications

Massive scale games, Plan/reason under uncertainty

Human adversary behavior modeling, Repeated games, Learn from field data

Infrastructure Security Games

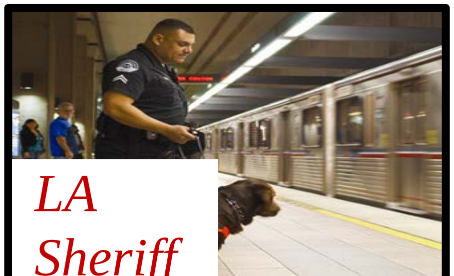

Green Security Games

Opportunistic Crime Security Games

ARMORWAY

Global Presence of Security Games Efforts

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
- Human behavior models: Quantal response

"Green security games": Global, interdisciplinary challenges

- Intersection of criminology, computation, conservation

Outline: Security Games Research (2007-)

Infrastructure Security Games

Green Security Games

Opportunistic Crime Security Games

Air travel

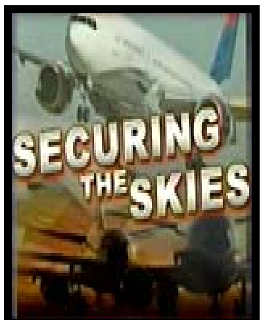

Ports

Trains

Fisheries

Wildlife

Urban Crime

2007 2011 2013 2014 2014 2015

Evaluation I: *AAAI, IJCAI, AAMAS papers…*

Evaluation II: *Real-world deployments (Patience)*

ARMOR Airport Security: LAX [2007]
First Stackelberg Security Game Application

Template for Security Game Operations [2007]
ARMOR Example


Pita


Paruchuri

	Target #1	Target #2	Target #3
Defender #1	2, -1	-3, 4	-3, 4
Defender #2	-3, 3	3, -2	….
Defender #3	….	….	….

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		

Security Game MIP [2007]

Generate Mixed Strategy for Defender in ARMOR


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

Security Game 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


Newsweek
ARMOR…throws a digital cloak of invisibility….

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

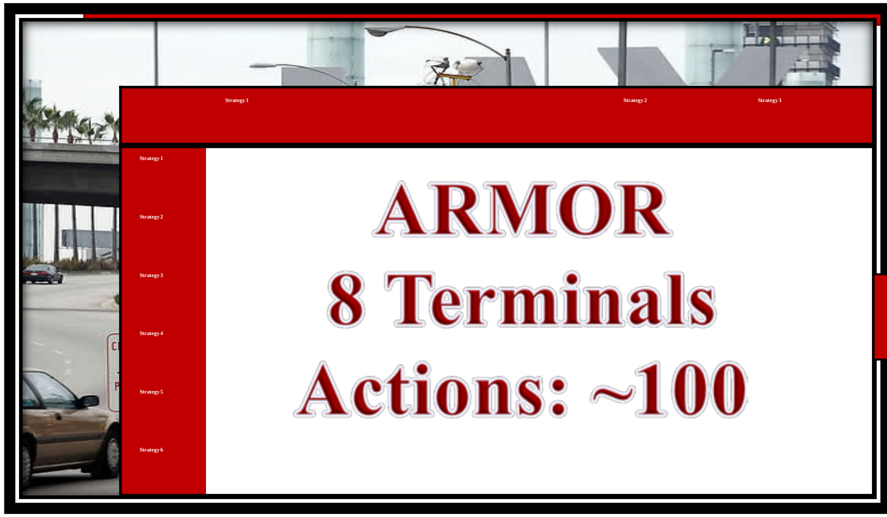

- <u>*1000 Flights, 20 air marshals: 10^{41} combinations*</u>
  - *ARMOR out of memory*
- <u>*Incremental strategy generation:*</u>
  - *Avoid enumerating all possible 10^{41} combinations*

IRIS: Scale Up [2009]
Small Support Set for Mixed Strategies


Jain Kiekintveld

Small support set size:

- Most xi variables zero

1000 flights, 20 air marshals:

10^{41} combinations

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

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

~~x123=0.0~~

x124=0.239

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

~~x135=0.0~~

x378=0.123

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

	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~~
…				
…				

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

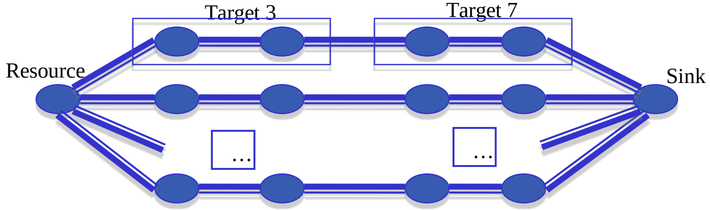

Converge:
GLOBAL
OPTIMAL

	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-)

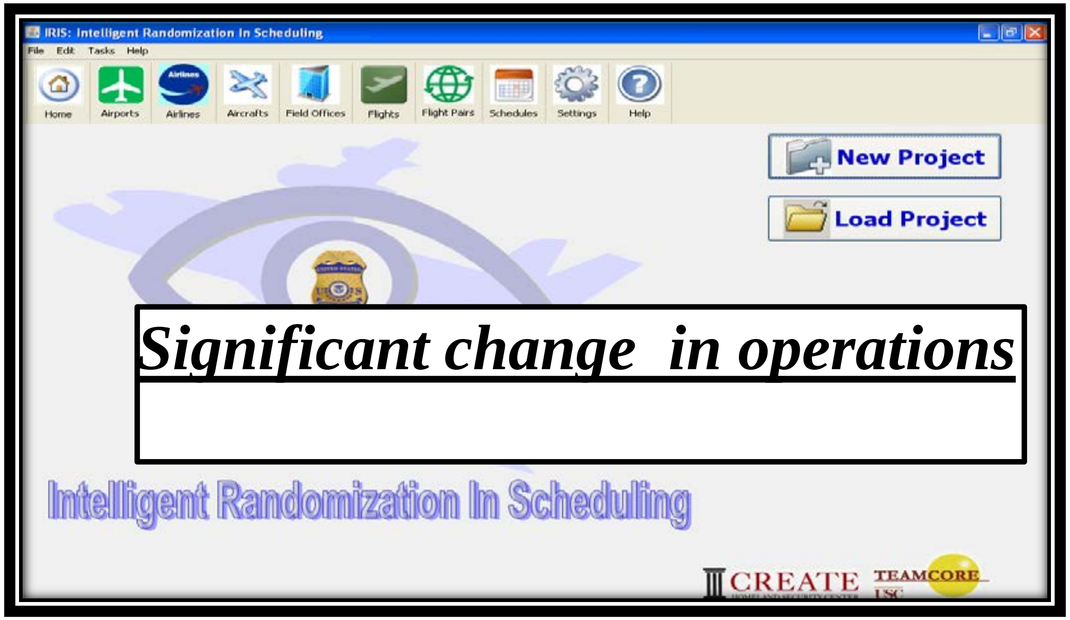

"…in 2011, the Military Operations Research Society selected a University of Southern California project with **FAMS on randomizing flight schedules for the prestigious Rist Award** *…"*

-R. S. Bray (TSA)

Transportation Security Subcommittee

US House of Representatives 2012

Road, Social Networks[2013]
Results of Scale-up

Jain

Road networks:
e.g., checkpoints

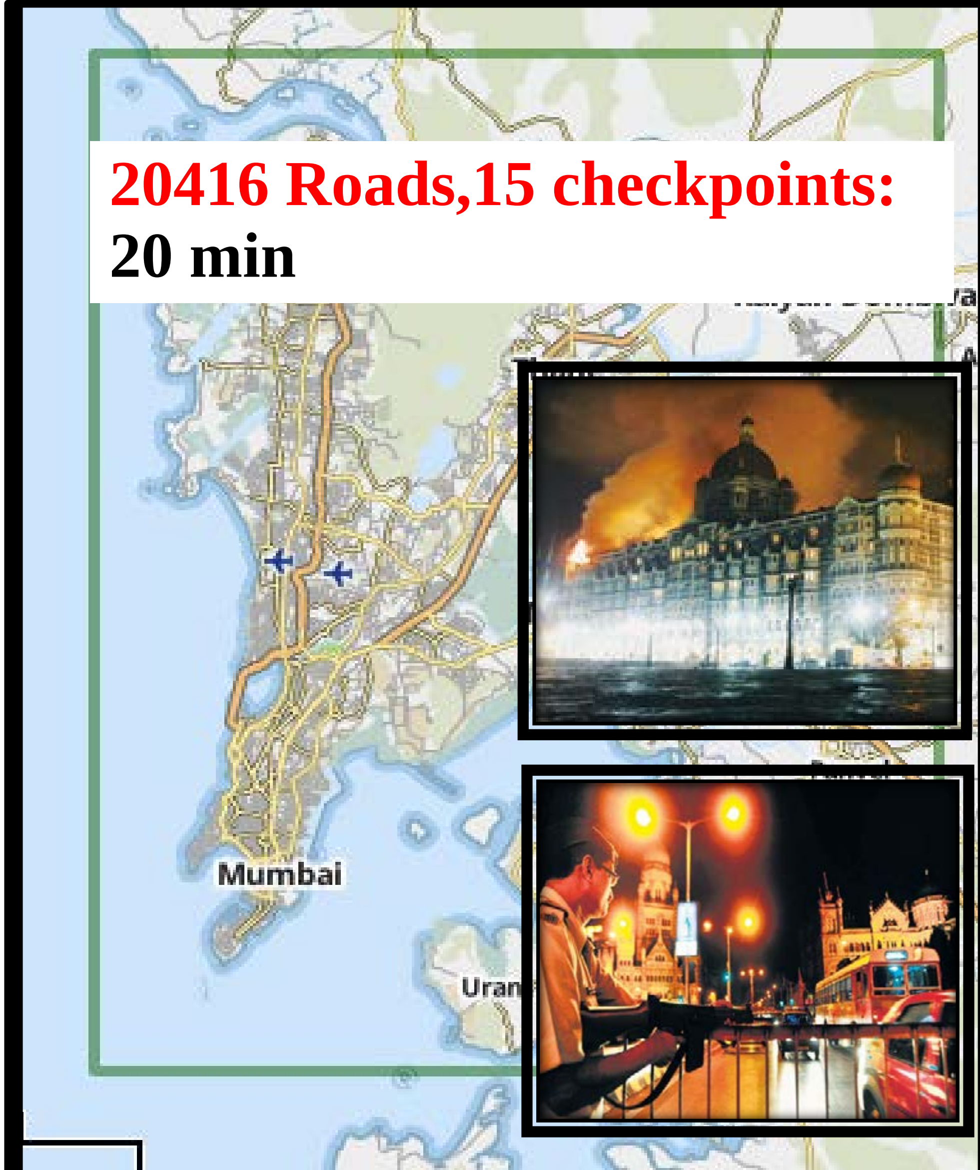

20416 Roads,15 checkpoints:
20 min

Outline: Security Games Research (2007-)

Air travel

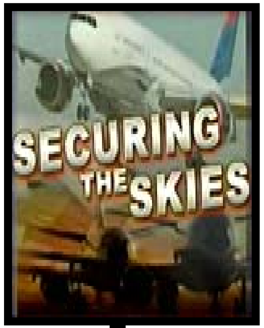

2007

Ports

2011

Trains

2013

Fisheries

2014

Wildlife

2014

Urban Crime

2015

Port Security Threat Scenarios

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)

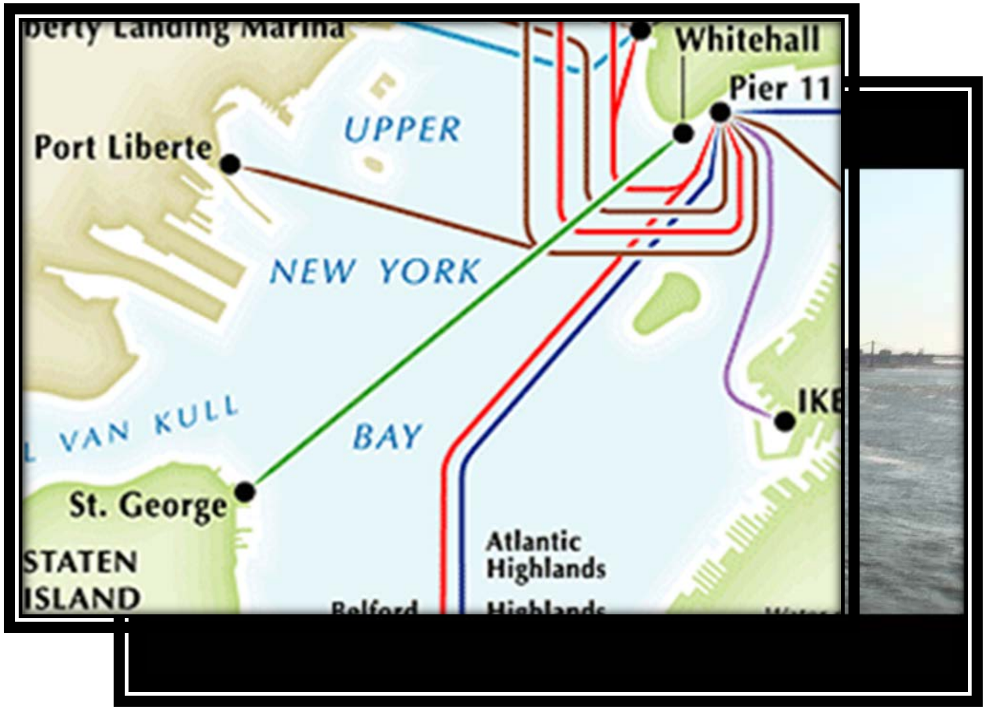

Ferries: Scale-up with Mobile Resources & Moving Targets
Transition Graph Representation

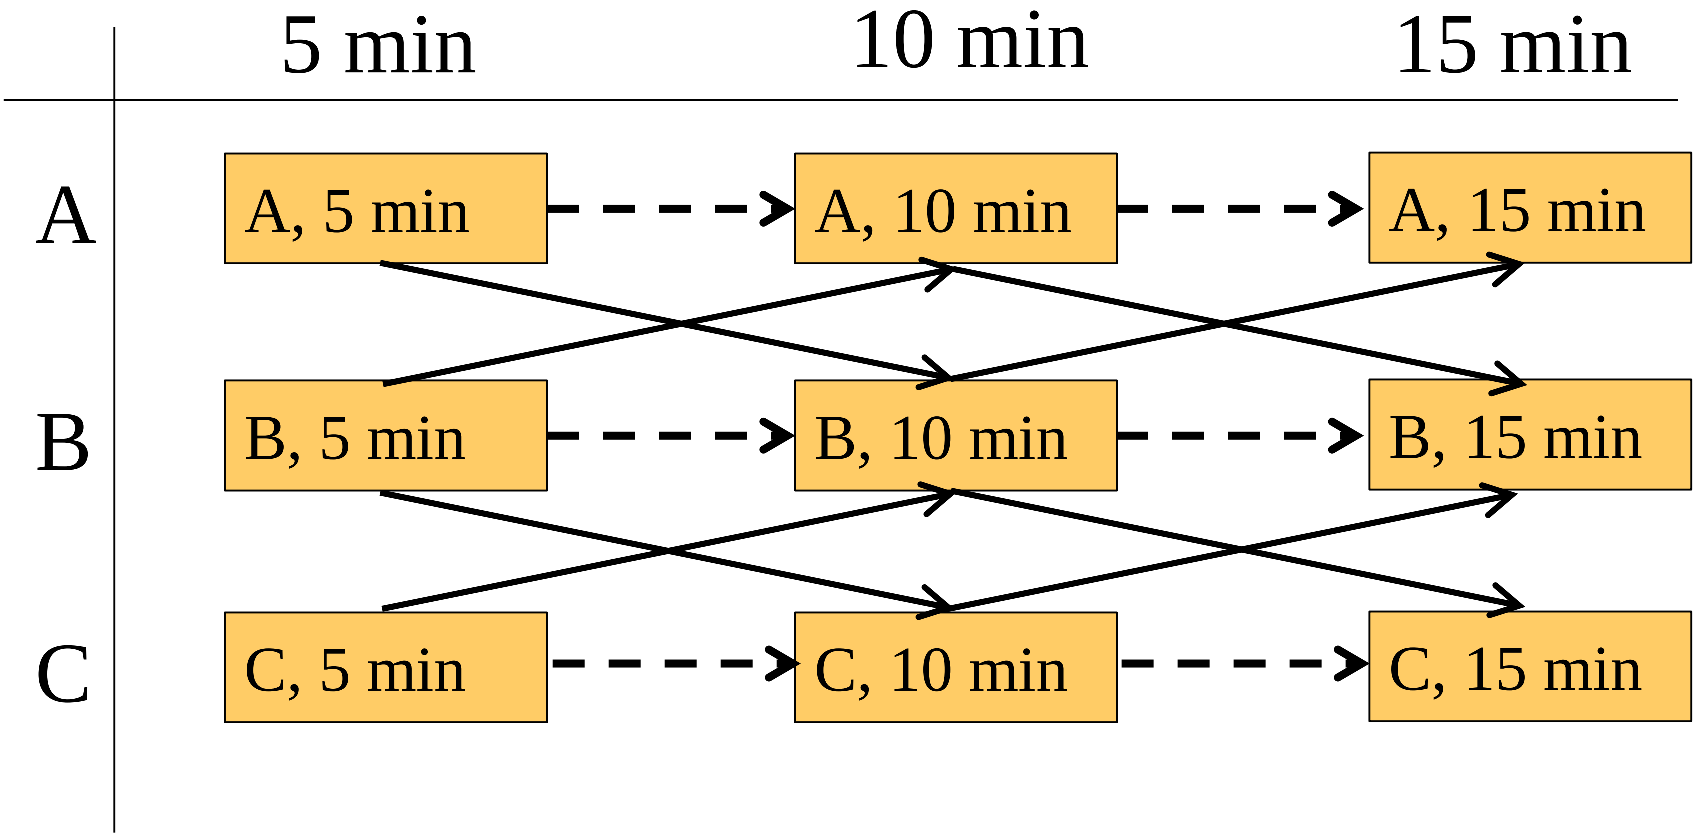

Ferries: Scale-up with Mobile Resources & Moving Targets
Transition Graph Representation

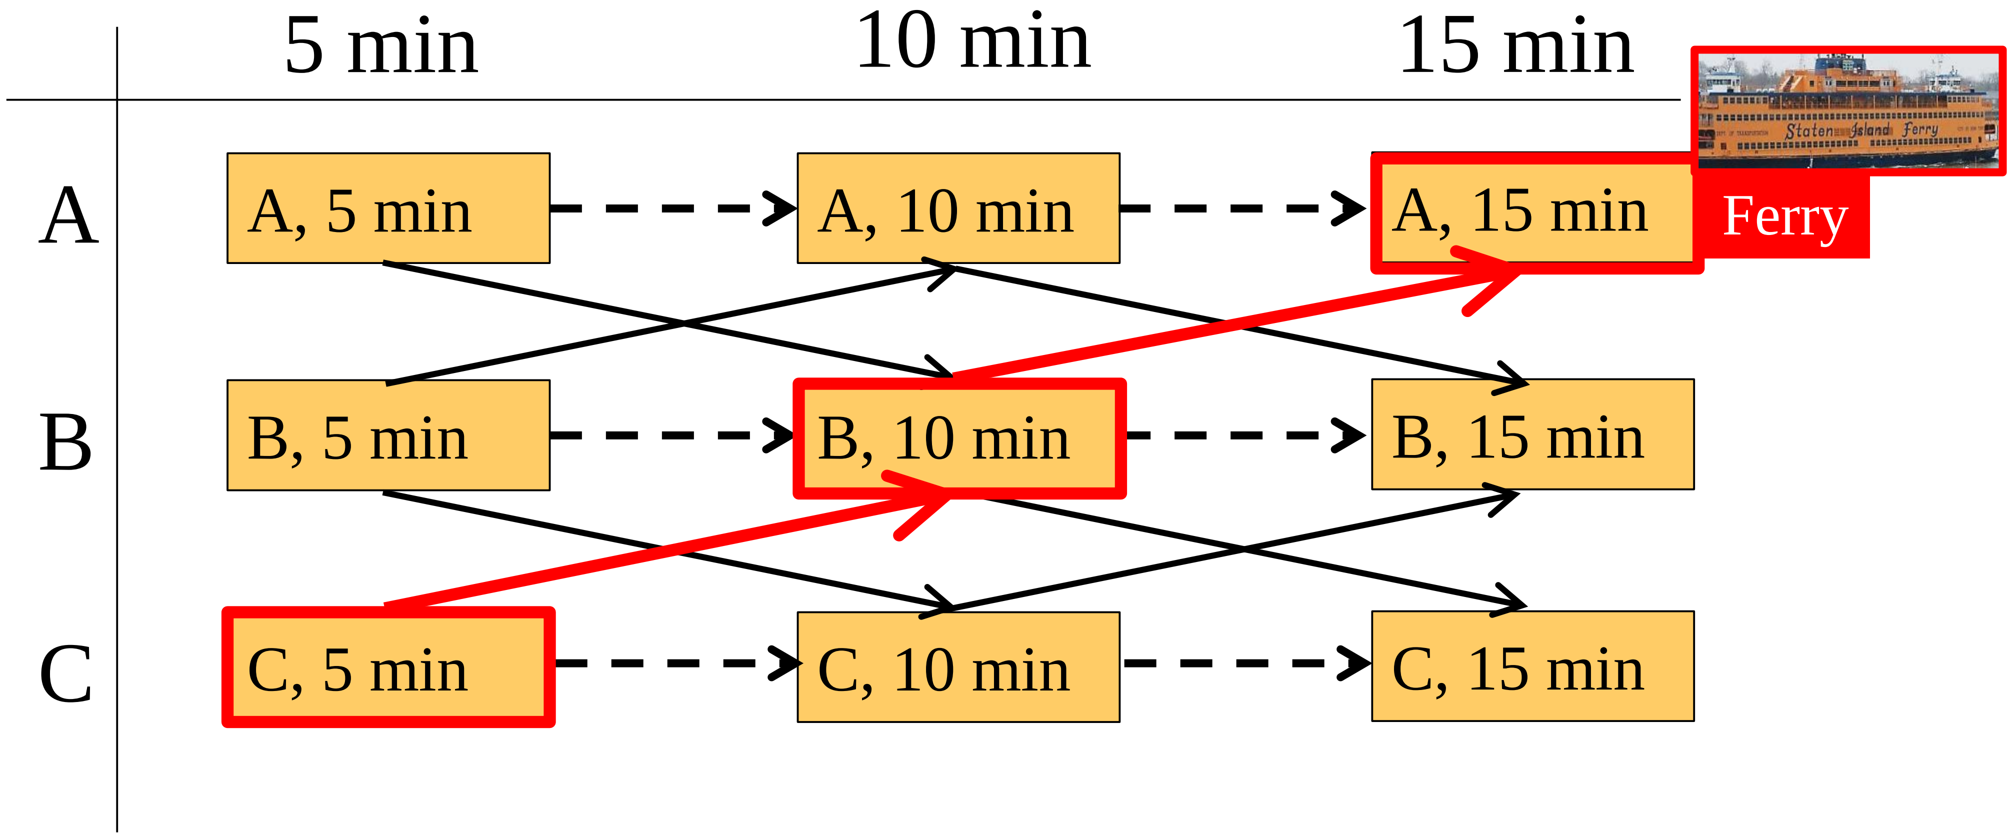

Ferries: Patrol Routes
Exponential Numbers of Patrol Routes

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

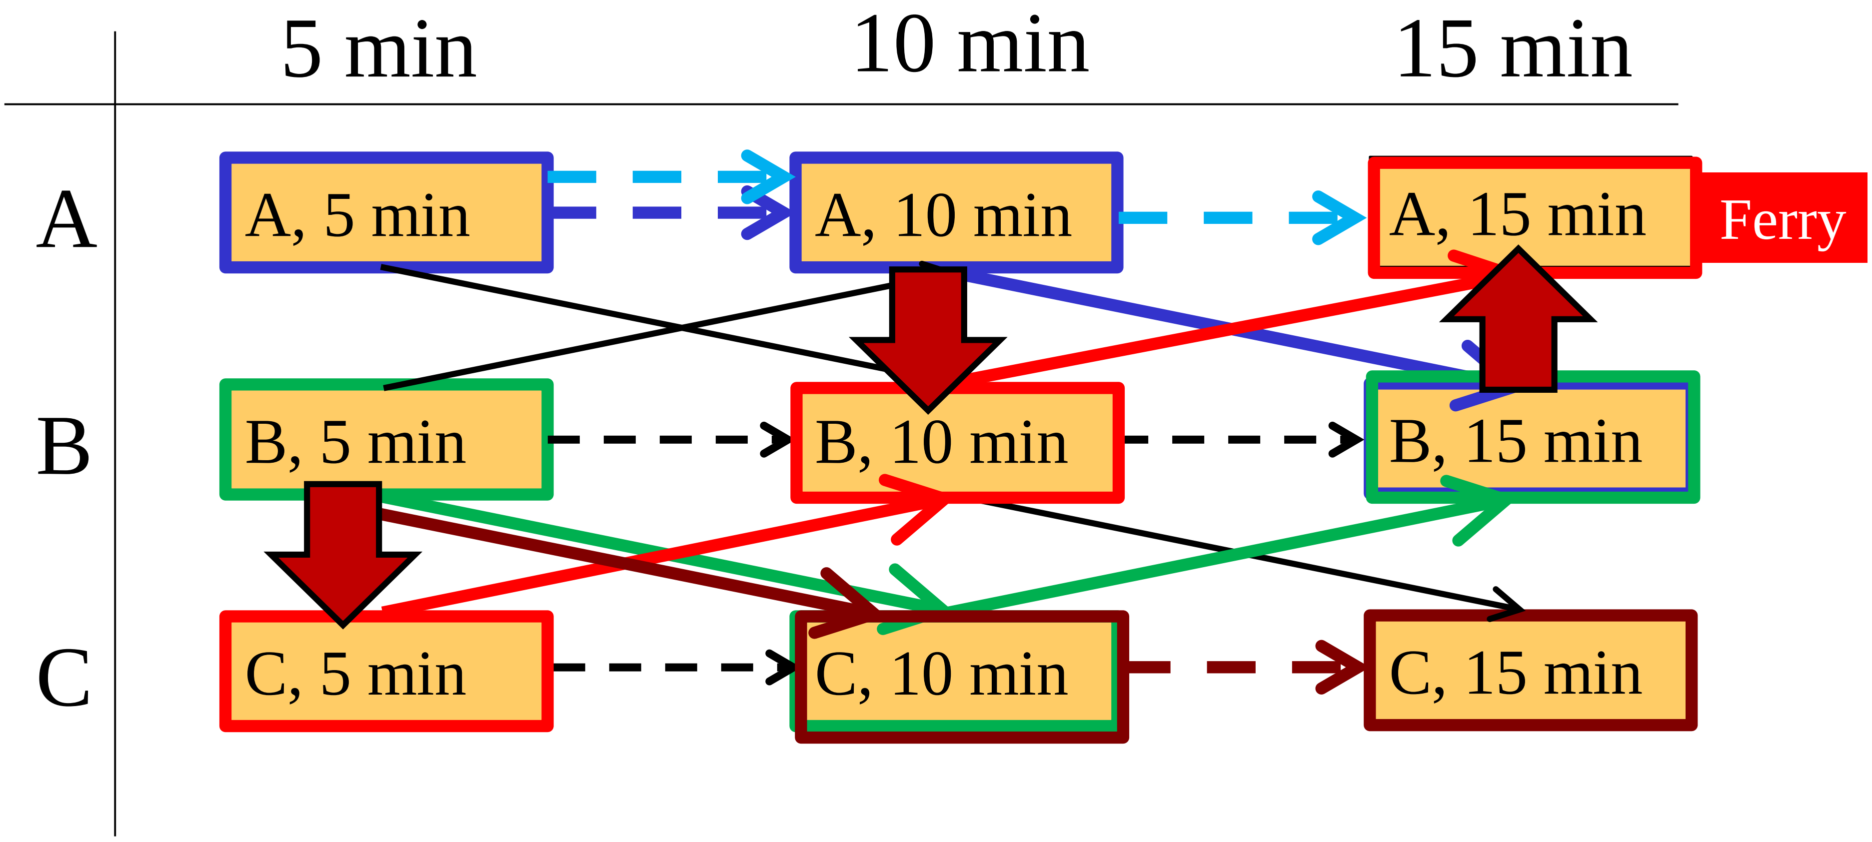

Ferries: Patrol Routes
Exponential Numbers of Patrol Routes

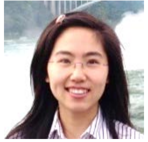
Fang


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

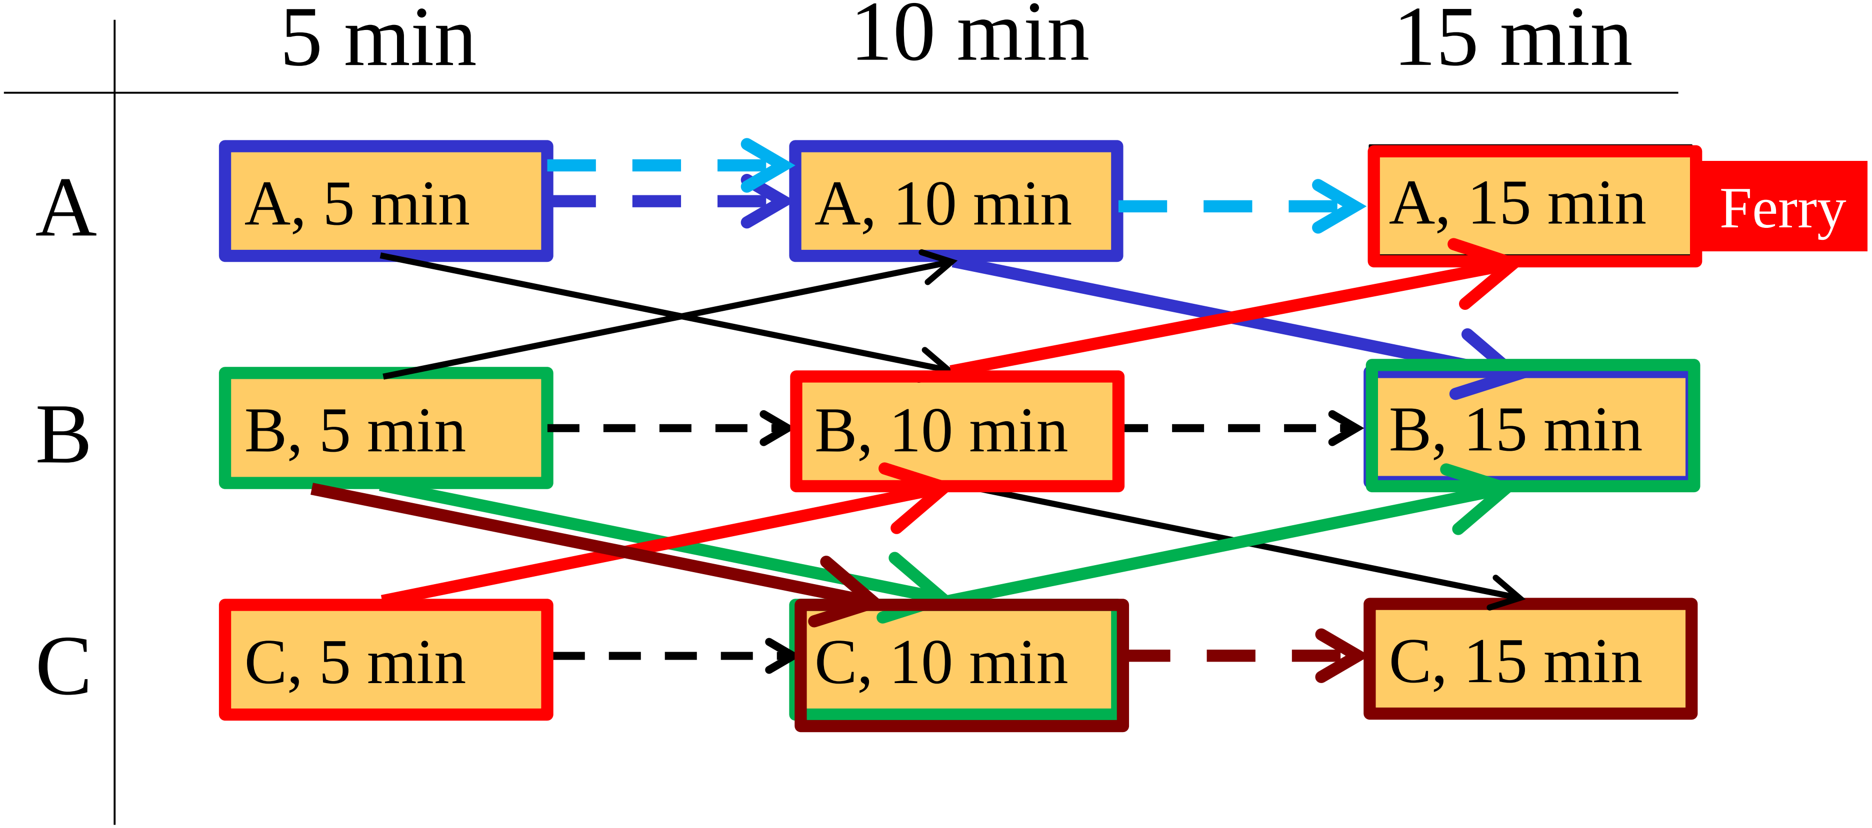

Ferries: Scale-up
Marginal Probabilities Over Segments

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

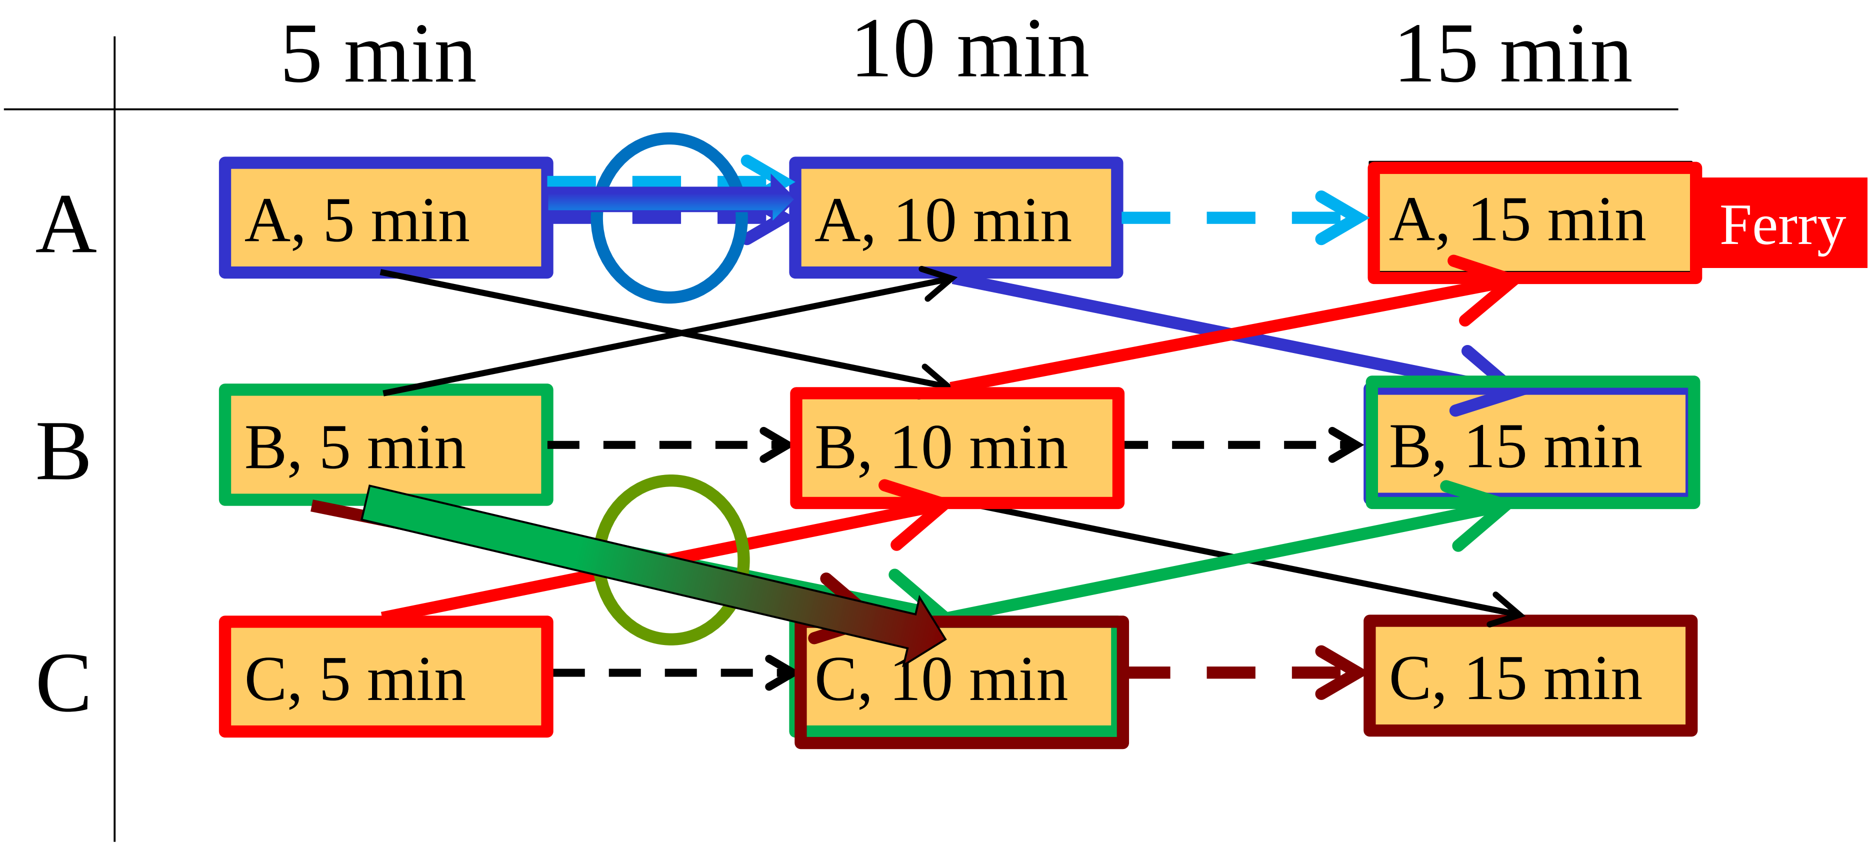

Ferries: Scale-up with Marginals Over Separable Segments
Significant Speedup

- *Obtain marginal probabilities over segments*

$N^2.T$ variables

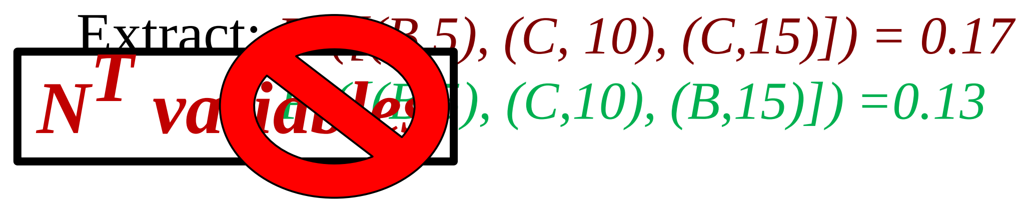

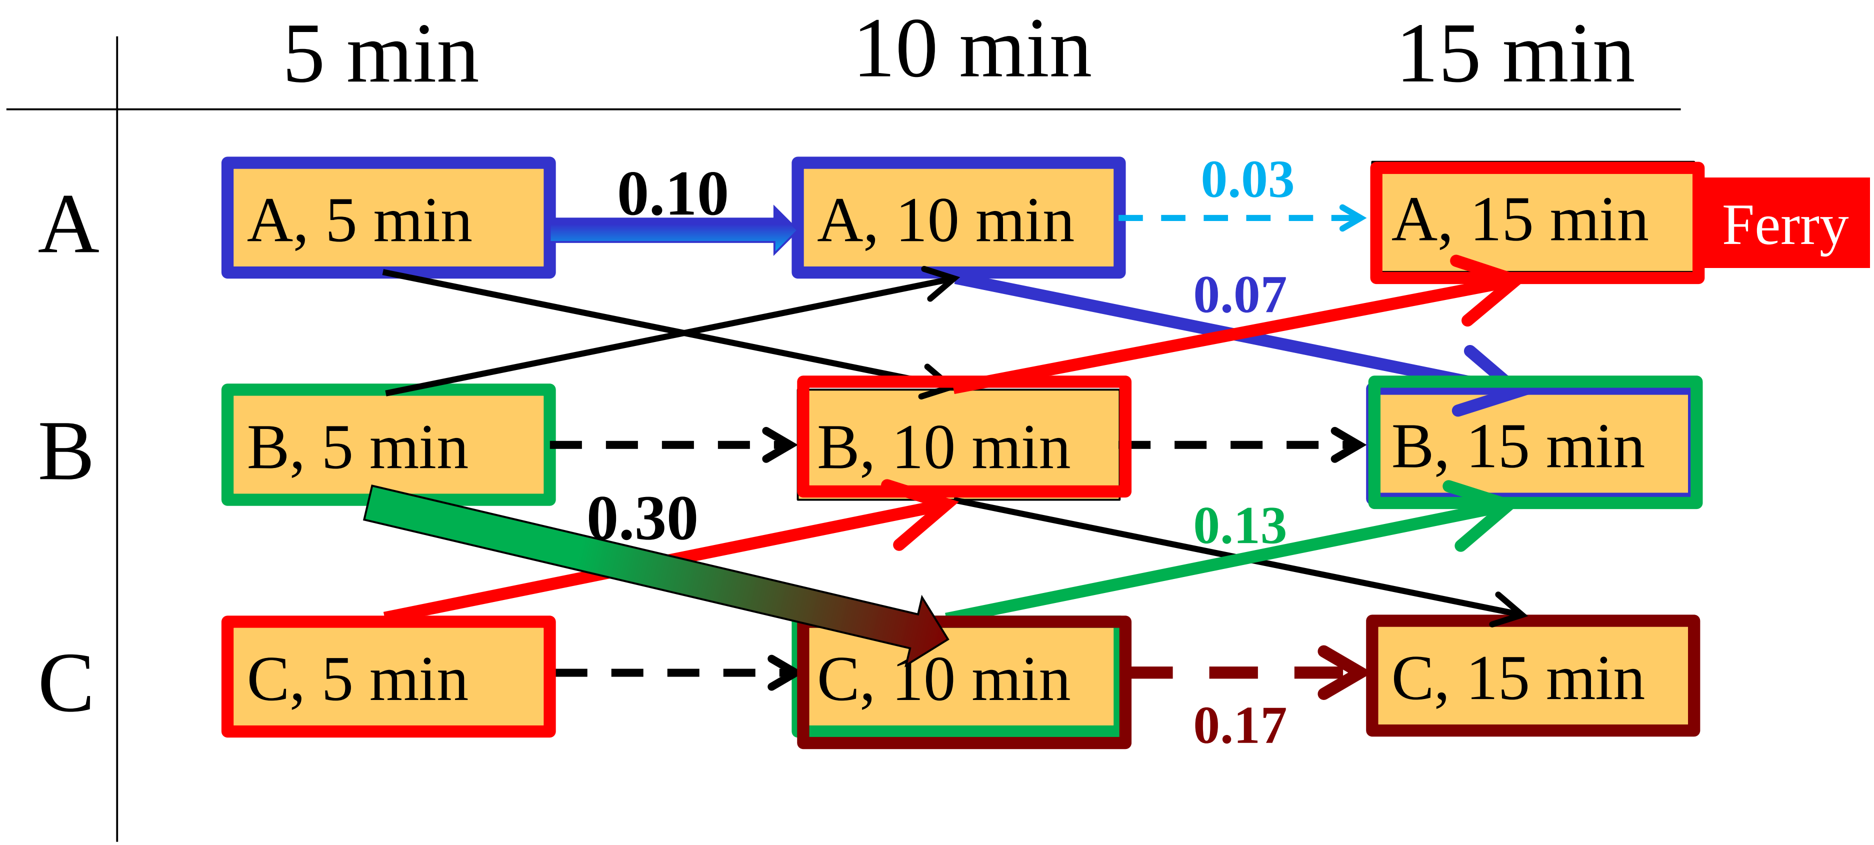

NOT VETTED BY CNN

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-)

Air travel

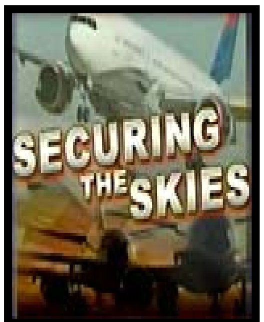

Urban Crime

2007

2015

TRUSTS: Frequent adversary interaction games
Patrols Against Fare Evaders

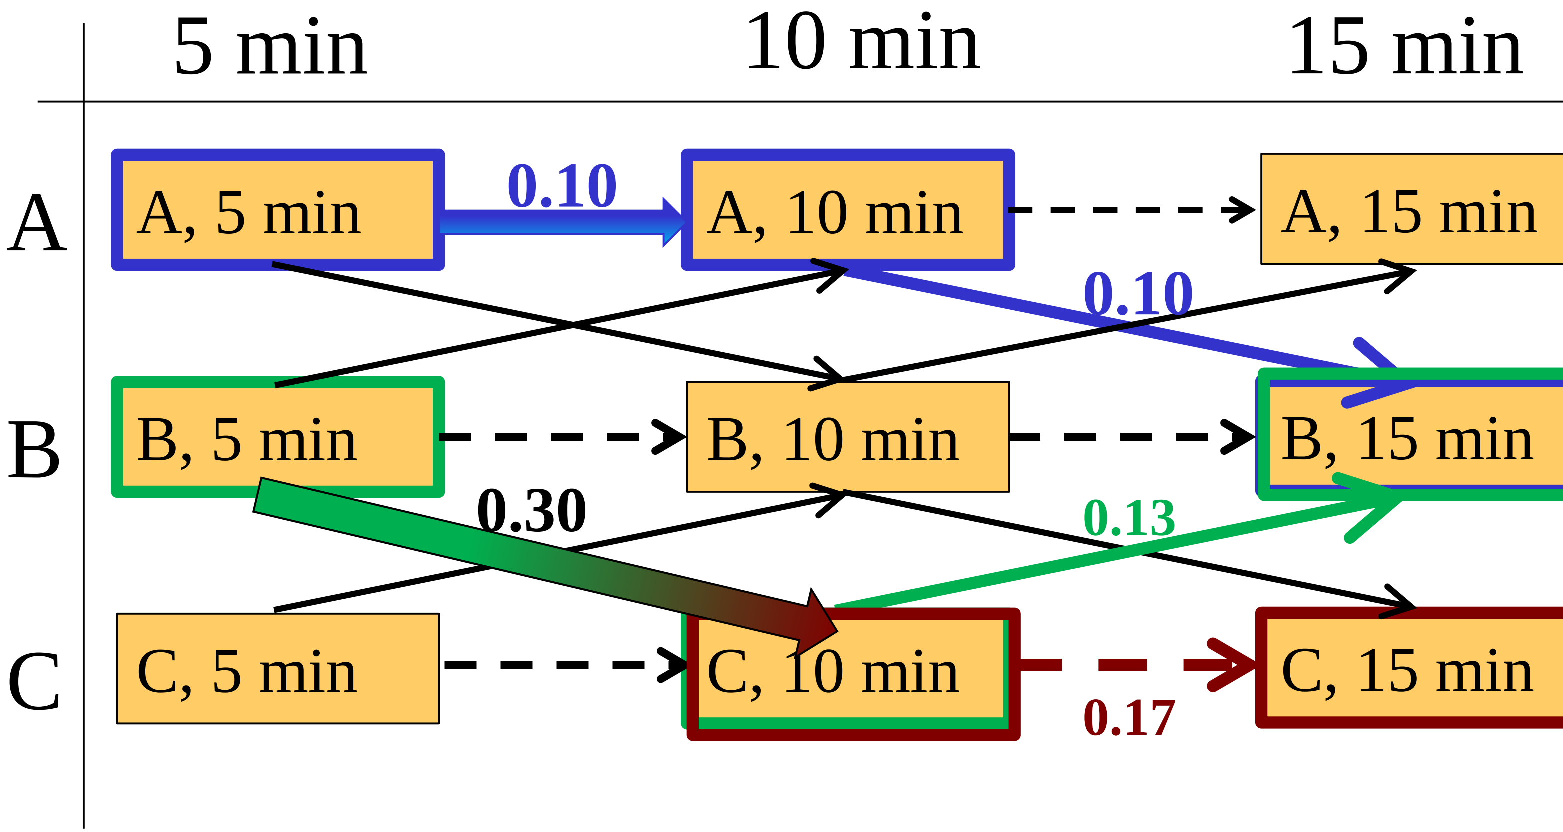

TRUSTS: Frequent adversary interaction
Uncertainty in Defender Action Execution


Jiang


Delle Fave

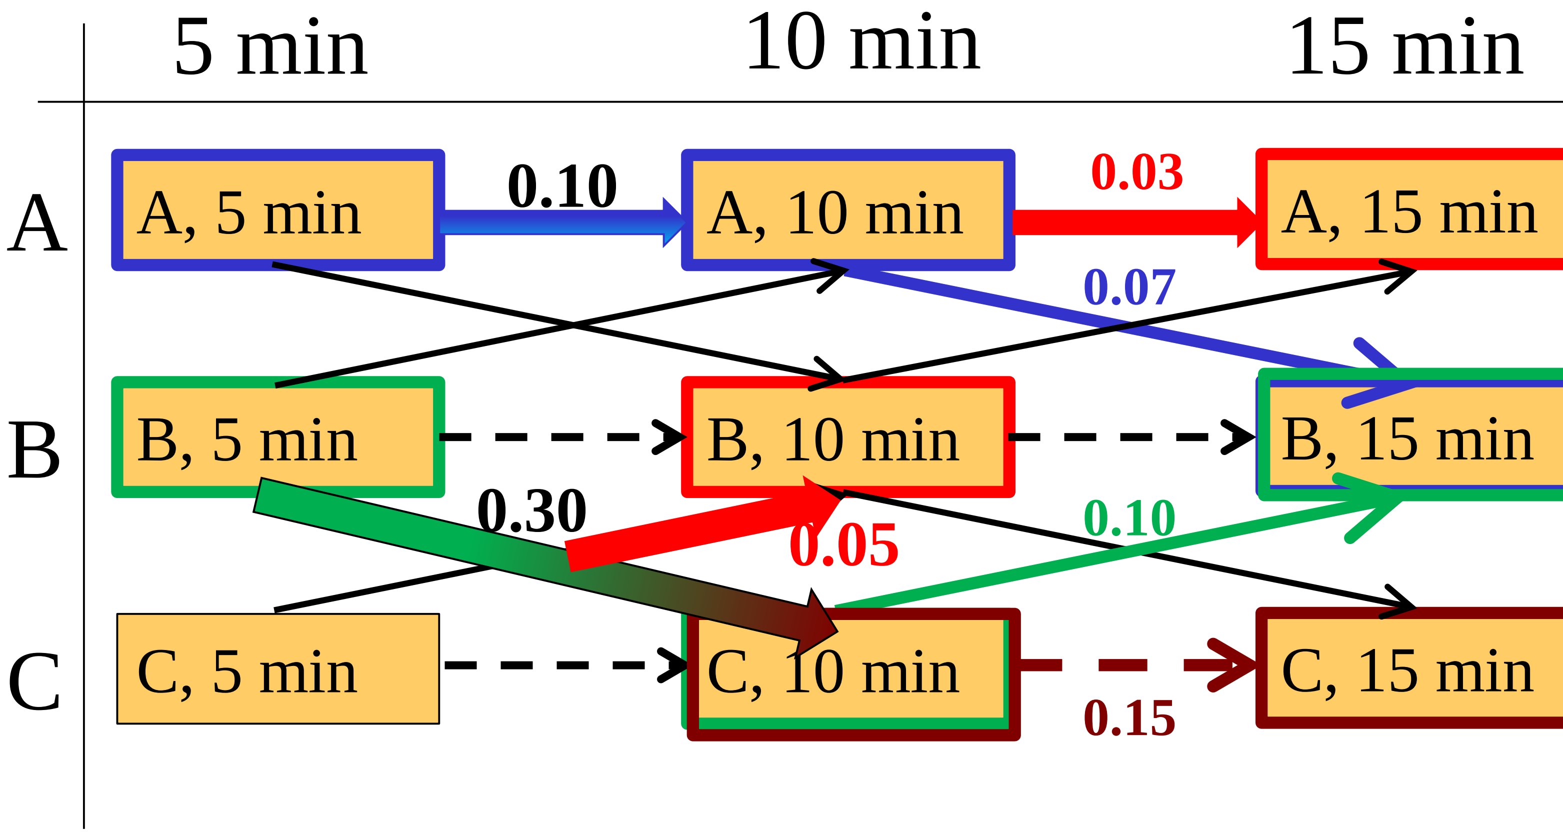

TRUSTS: Frequent adversary interaction Uncertainty in Defender Action Execution


Jiang


Delle Fave

- Markov Decision Problems *in Security games*

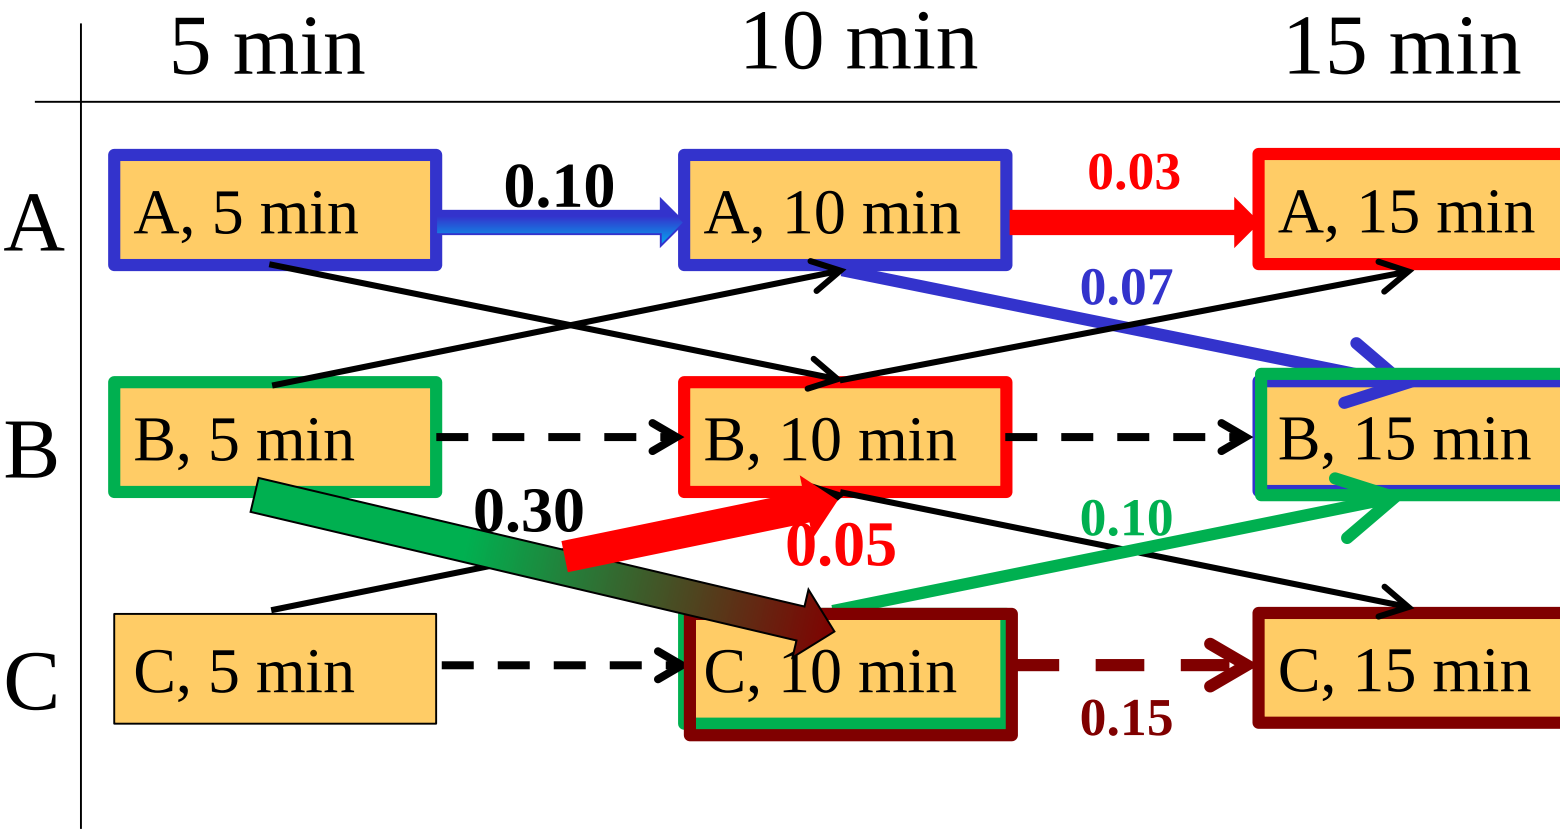

Uncertainty Space Algorithms: Bayesian and Robust Approaches

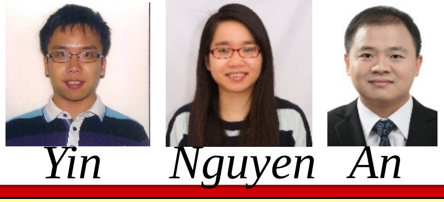

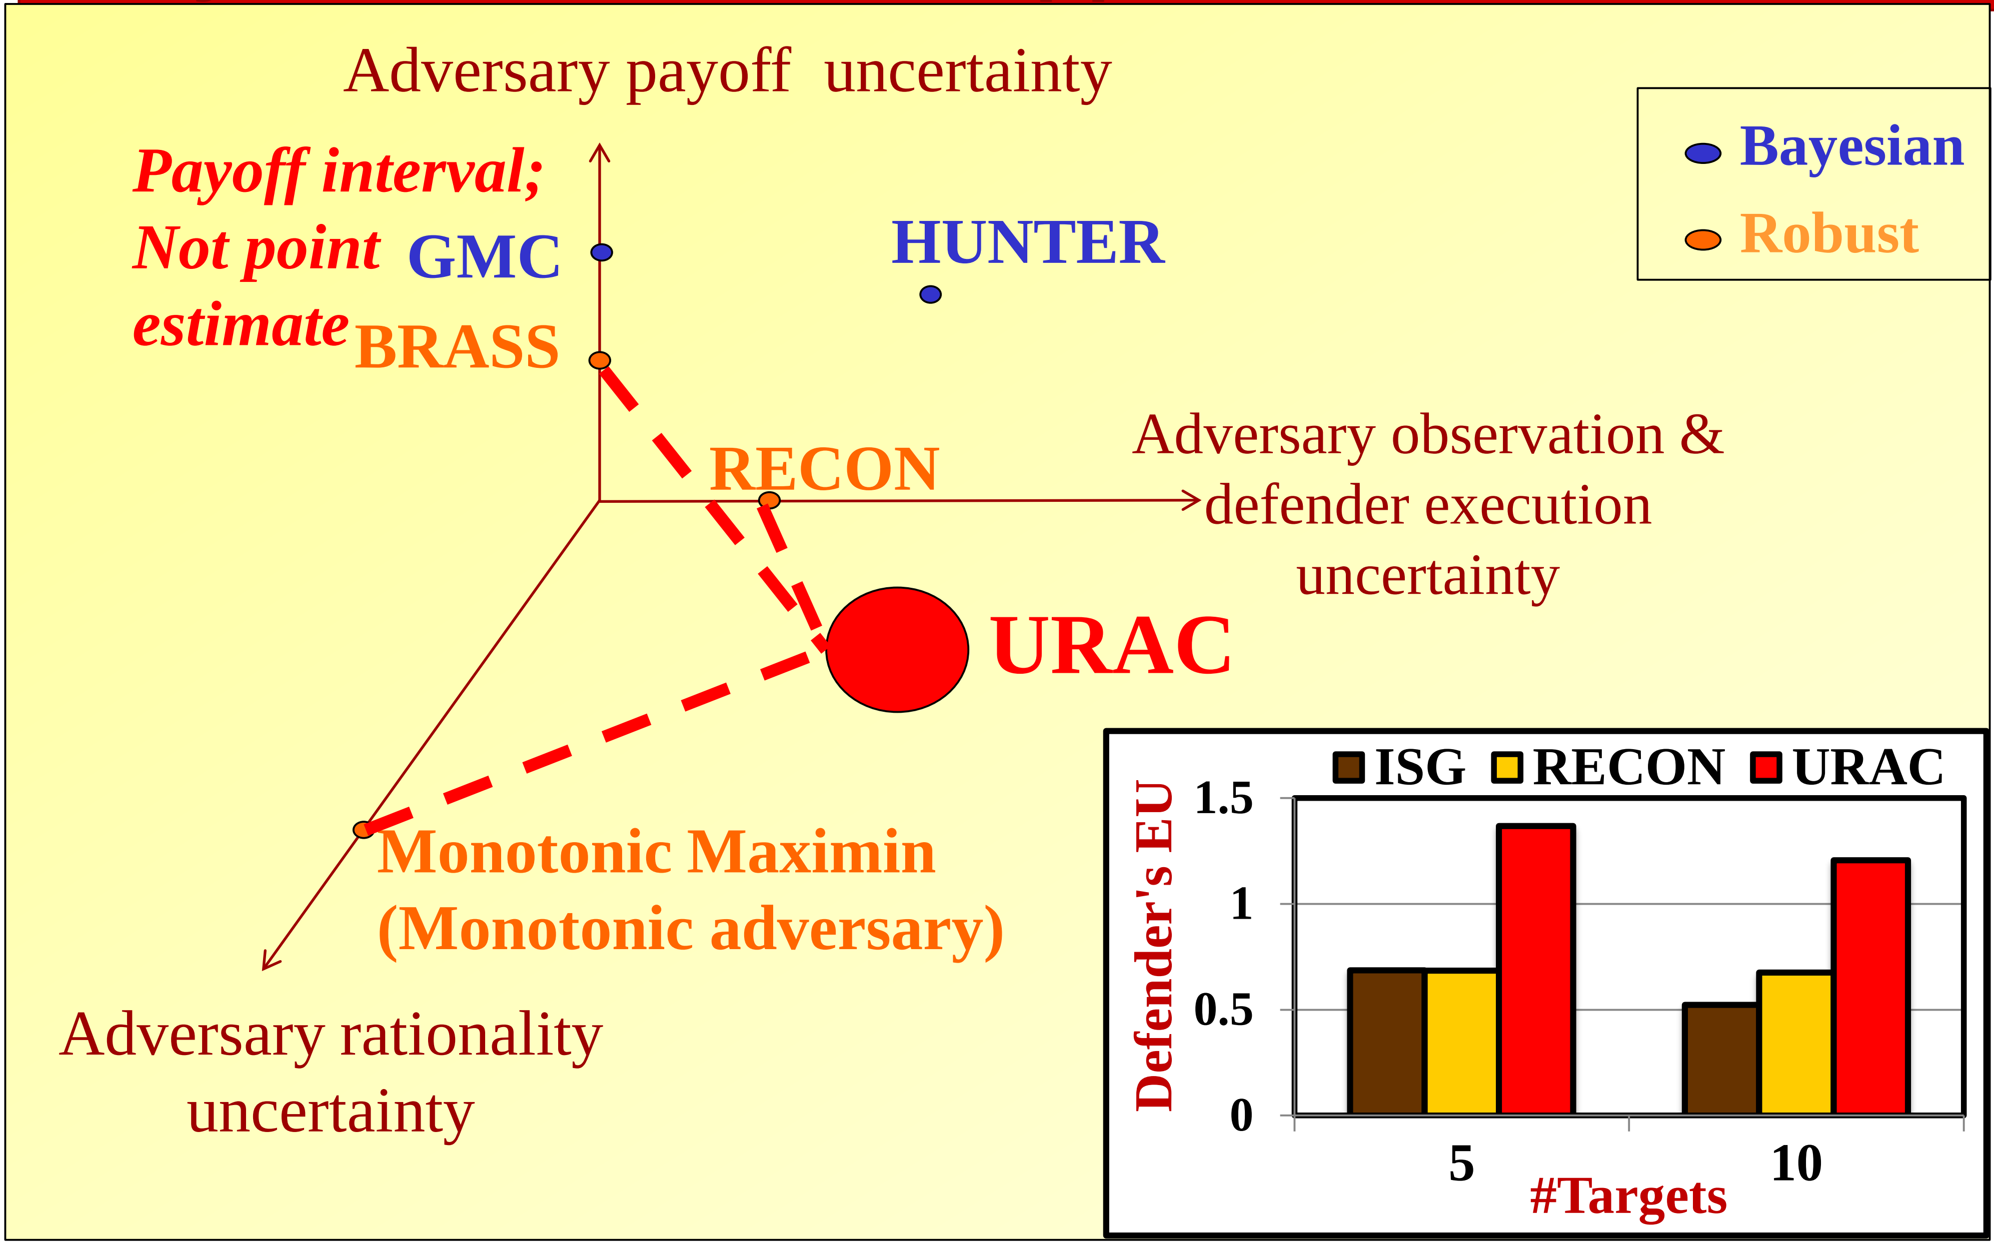

Outline: Security Games Research (2007-)

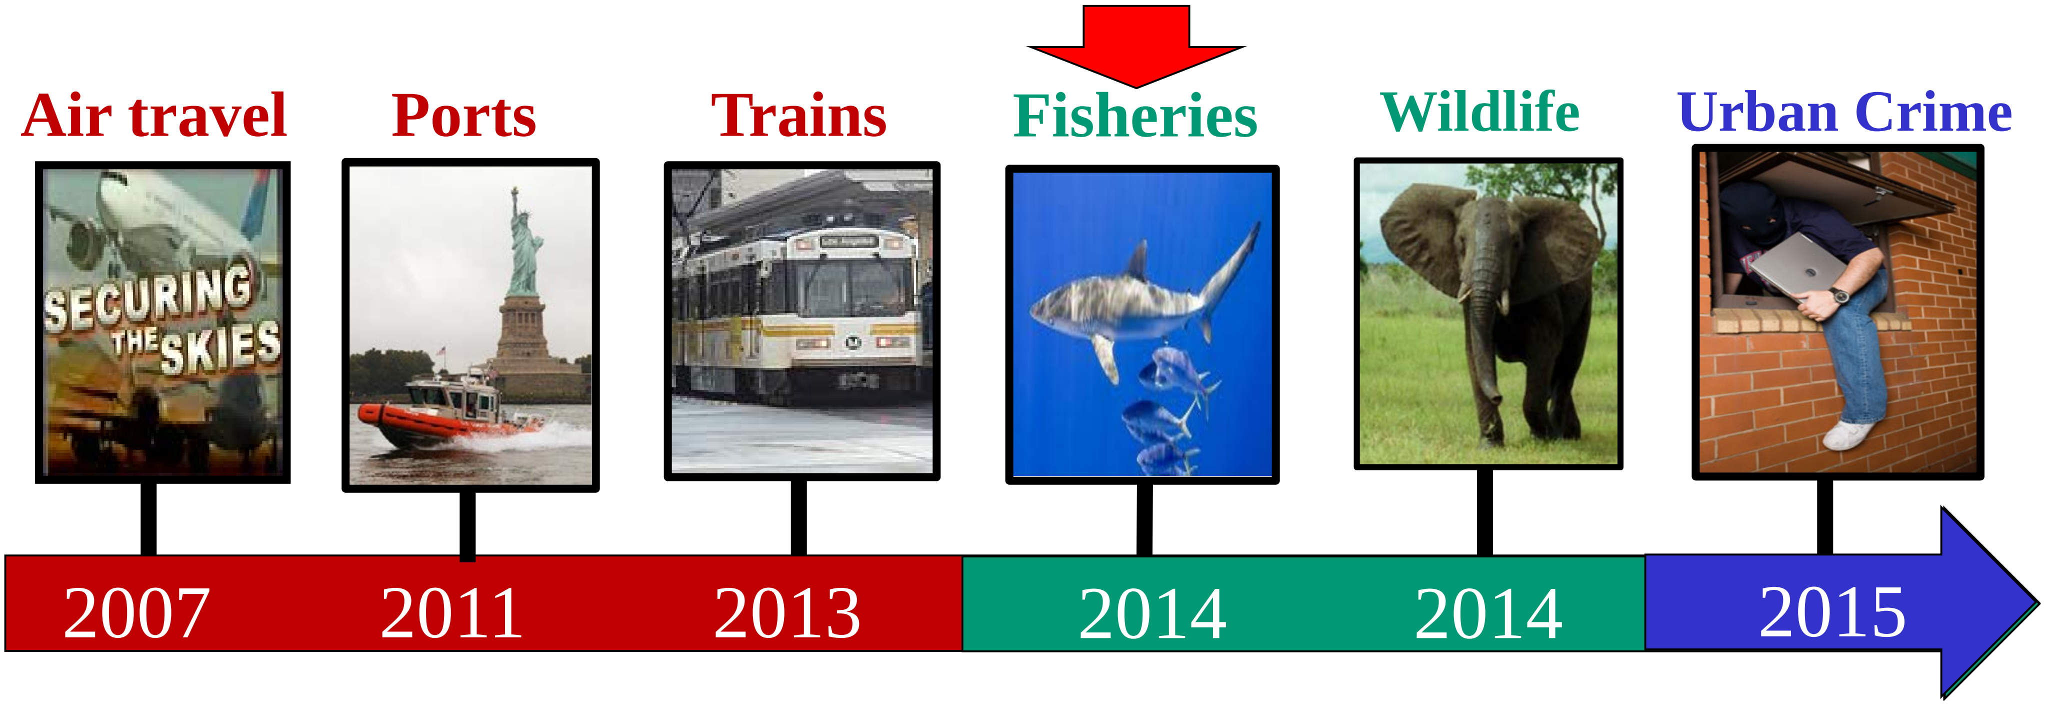

Protecting Forests, Fish and Wildlife: Green Security Games


Yang

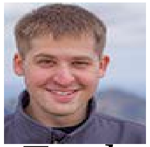
Ford


Brown

Wildlife Protection

Forest Protection

Fishery Protection

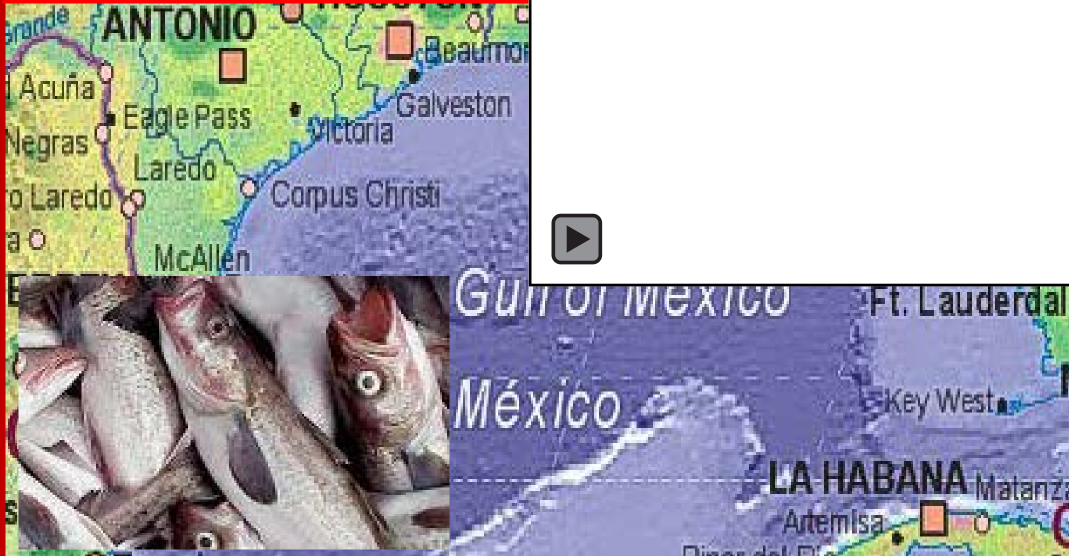

Murchison Falls National Park, Uganda






Green Security Games: Repeated Stackelberg Game

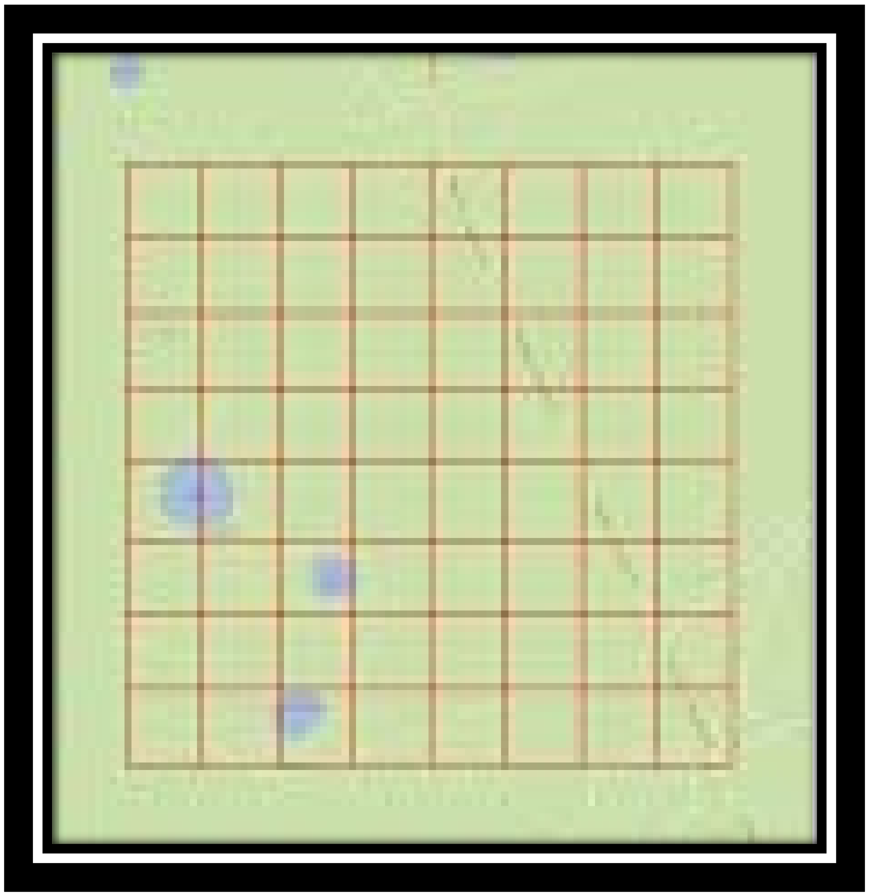

Bounded rationality model of poachers

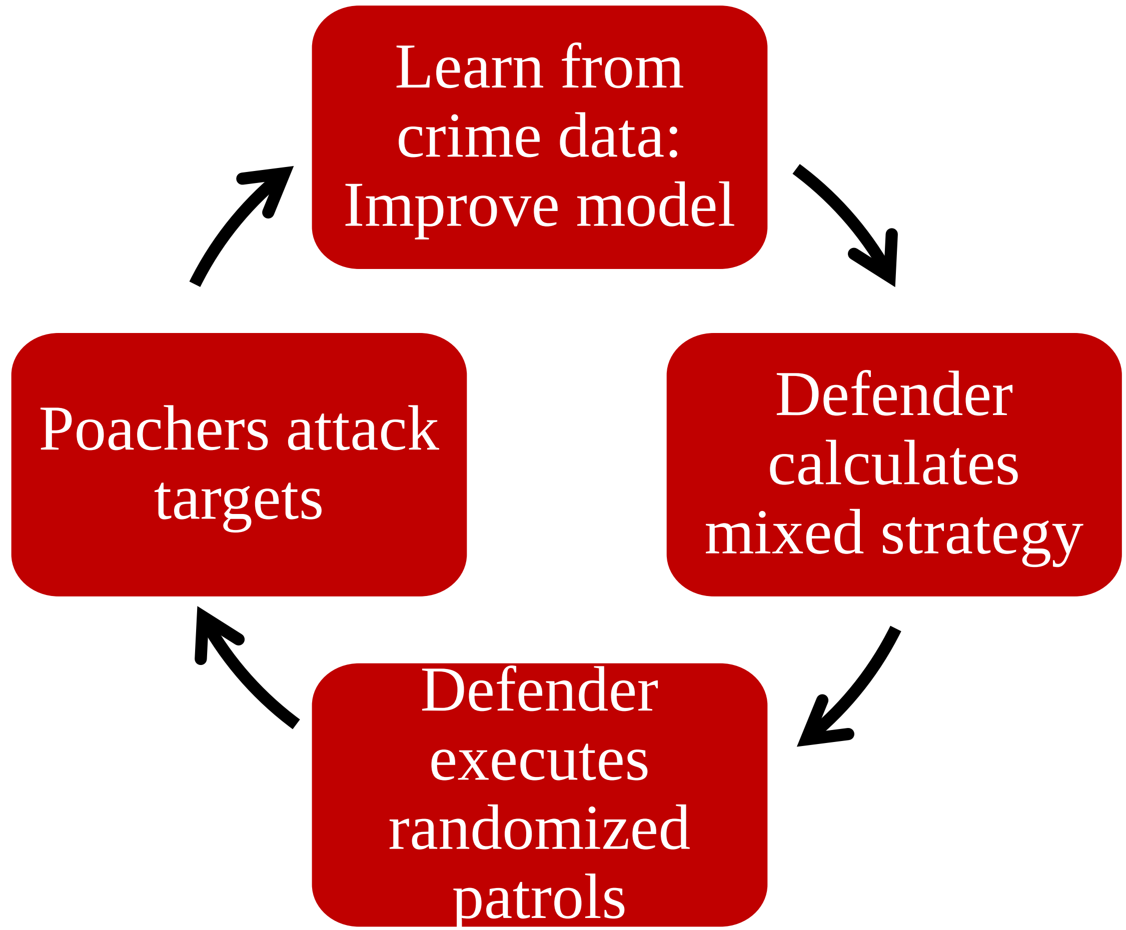

Uncertainty in Adversary Decision: Bounded Rationality
Human Subjects as Poachers

Kar

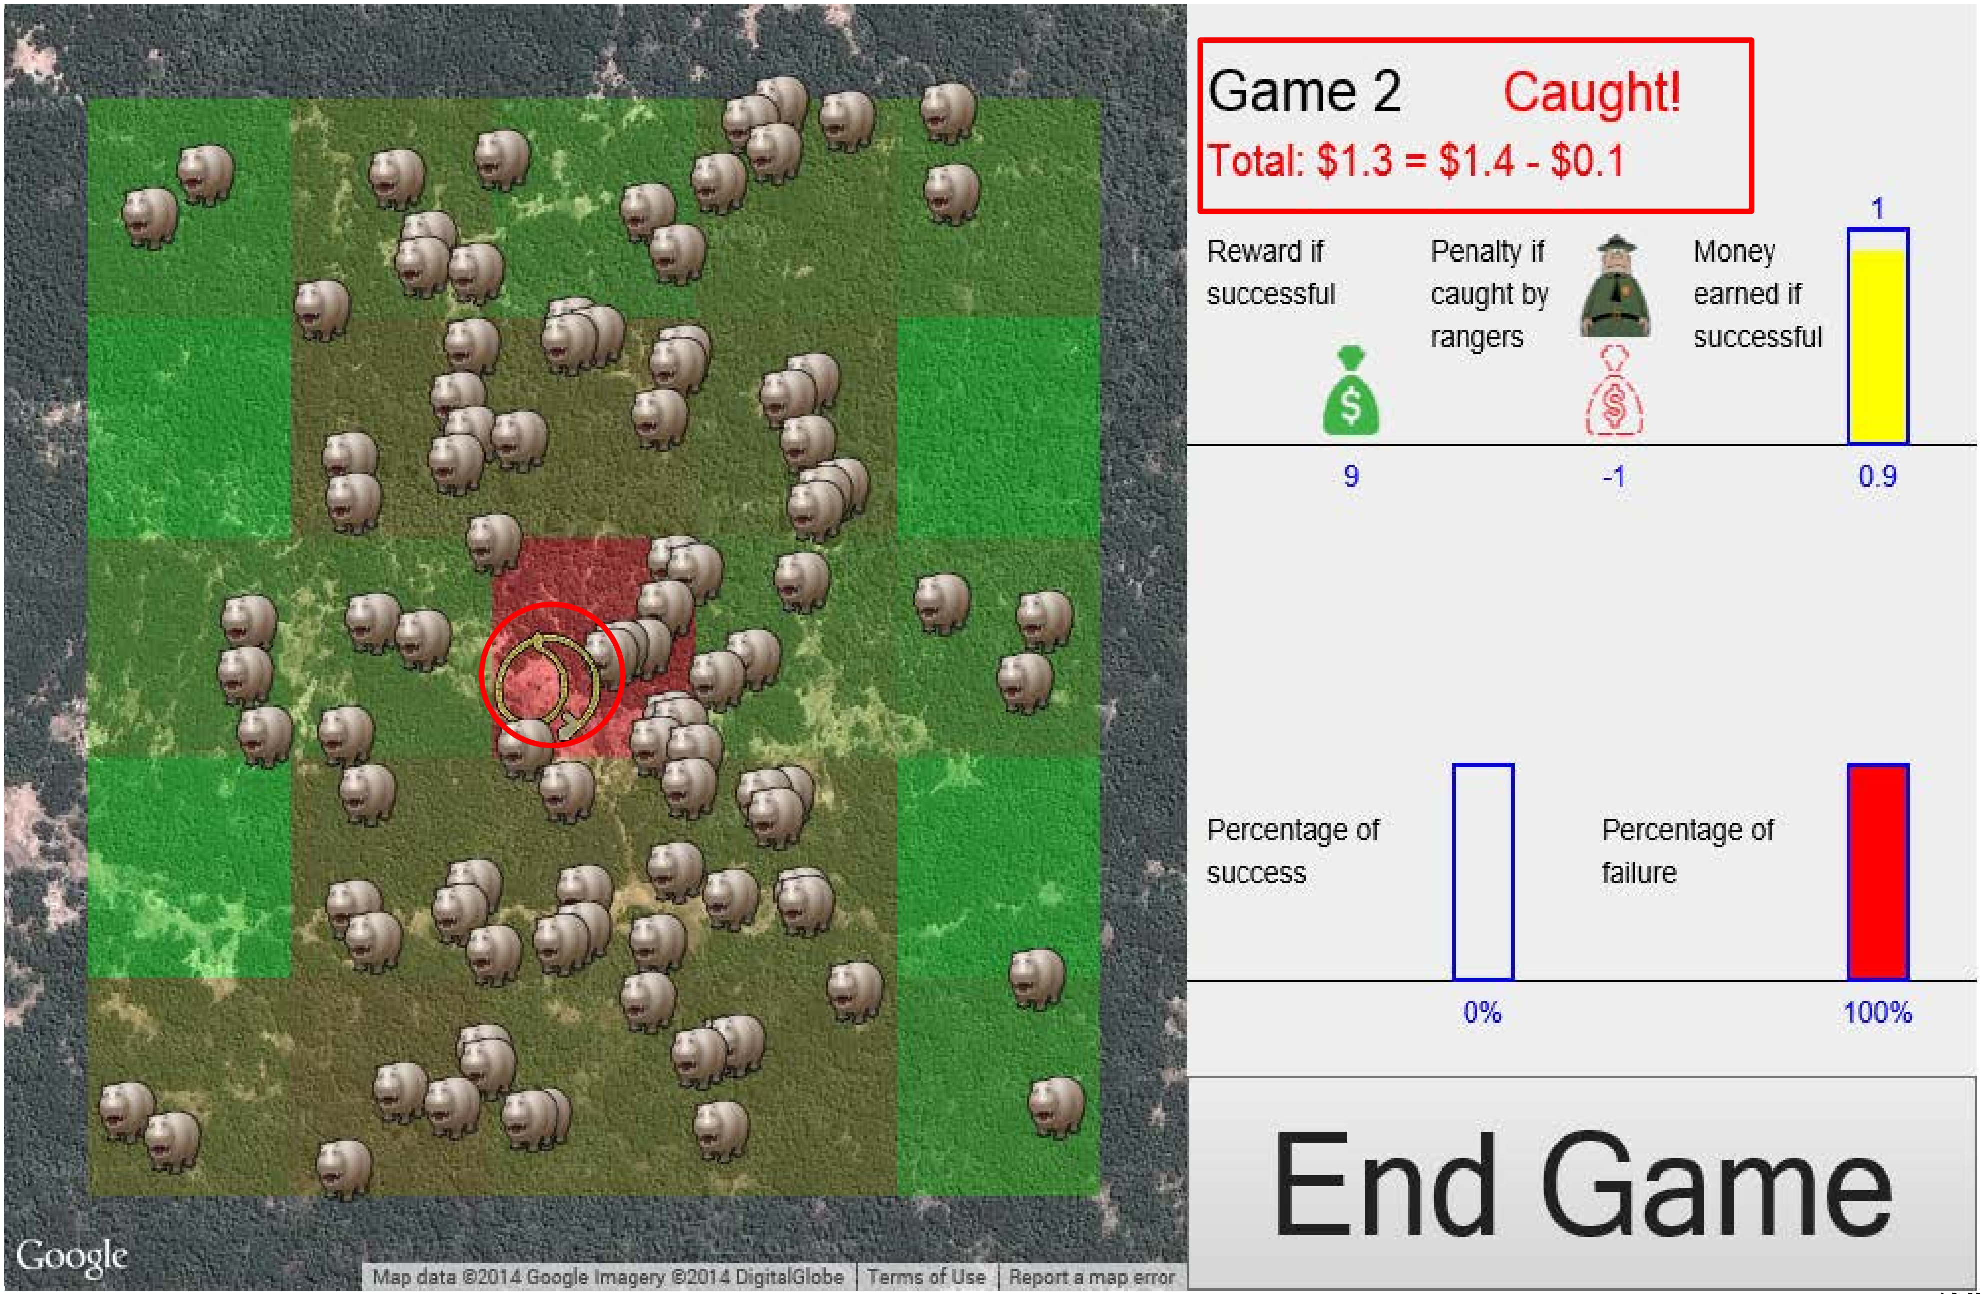

Lesson 1: Quantal Response [2011]: Models of Bounded Rationality

Pita

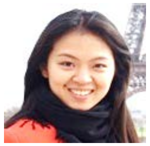

Yang

Perfect: $$EU^{adversary}(j) = Capture\text{Prob} \times \text{Penalty} + (1 - Capture\,\text{Prob}) \times \text{Reward}$$

Quantal Response(QR) ***[McFadden 73]:*** Stochastic Choice, Better Choice More likely

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

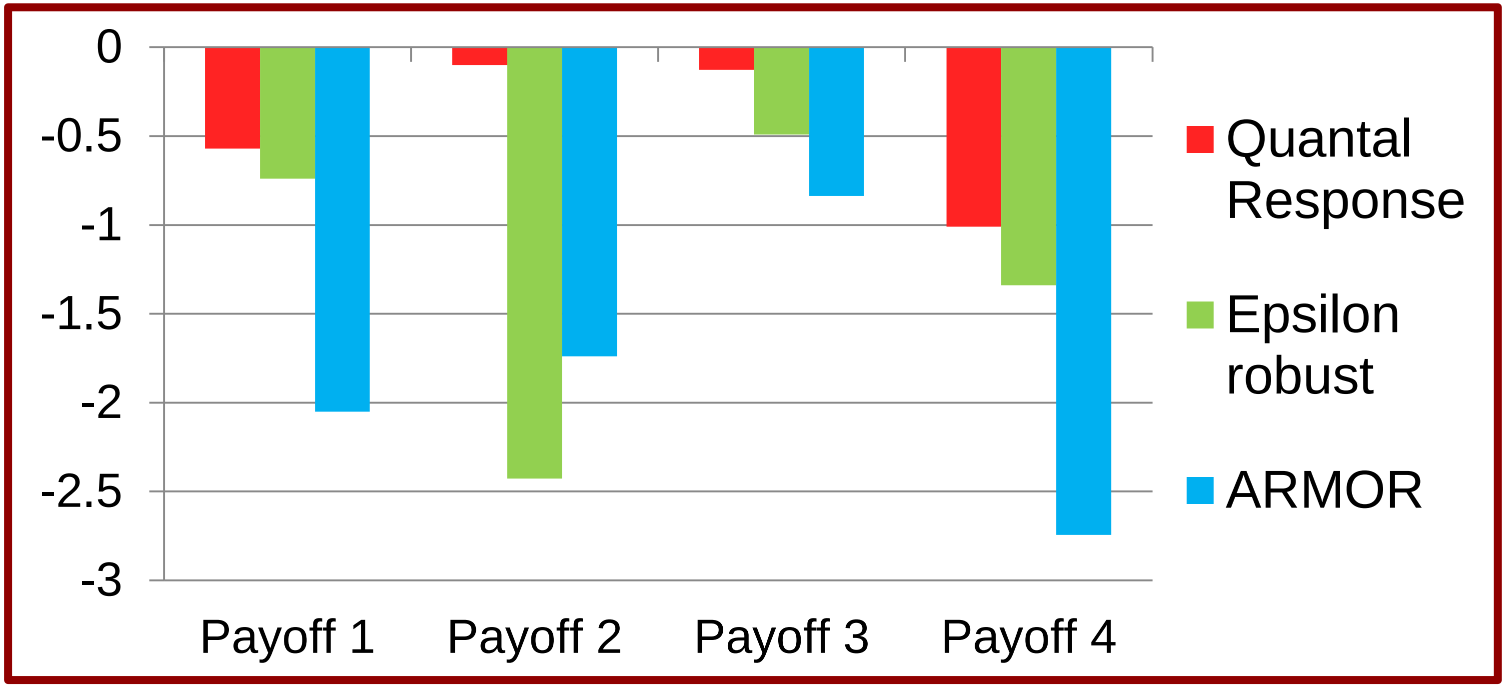

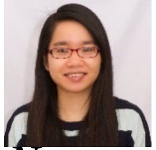

Nguyen

Subjective Utility Quantal Response(SUQR) [Nguyen 13]:

$$SEU^{adversary}(j) = w_1 \times Capture\,\text{Prob} + w_2 \times \text{Reward} + w_3 \times \text{Penalty}$$

Adversary's probability of **choosing target j**

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

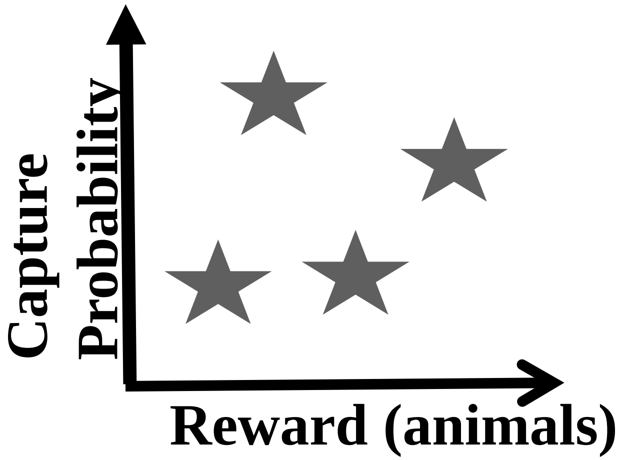

Poaching data from WWF

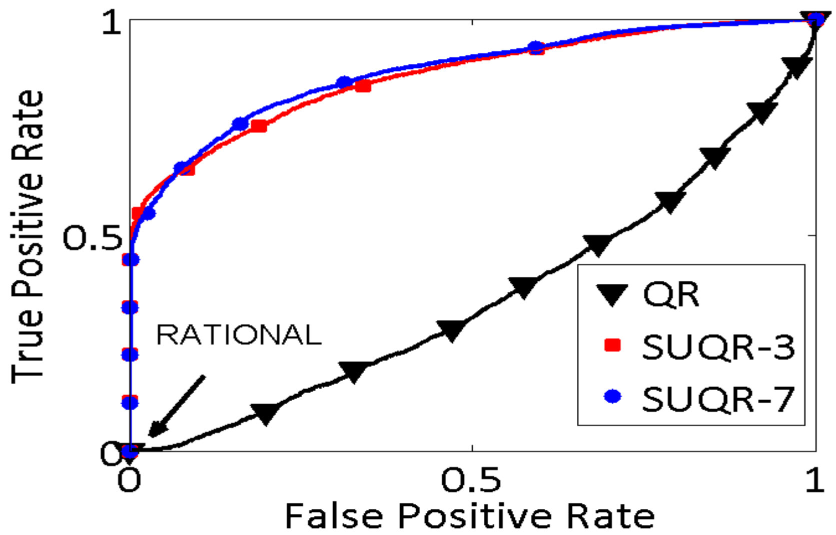

Green Security Games[2015]
Repeated Stackelberg Game


Kar

Bayesian SUQR:
Heterogeneous Poachers

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

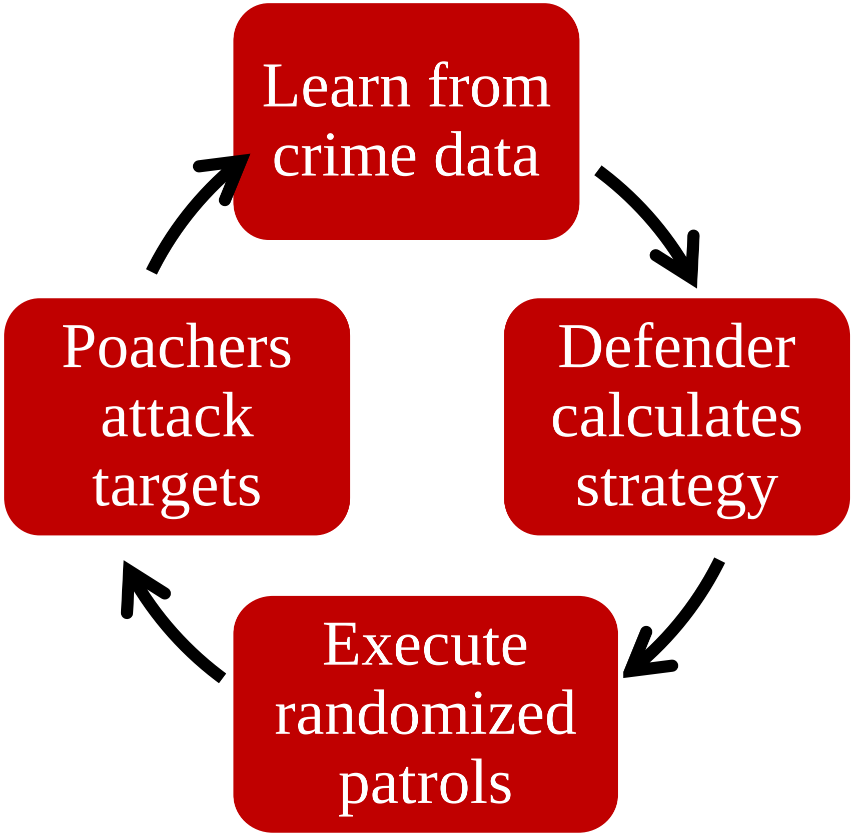

Logitudinal study on AMT:
35 weeks, 40 human subjects
10,000 emails!

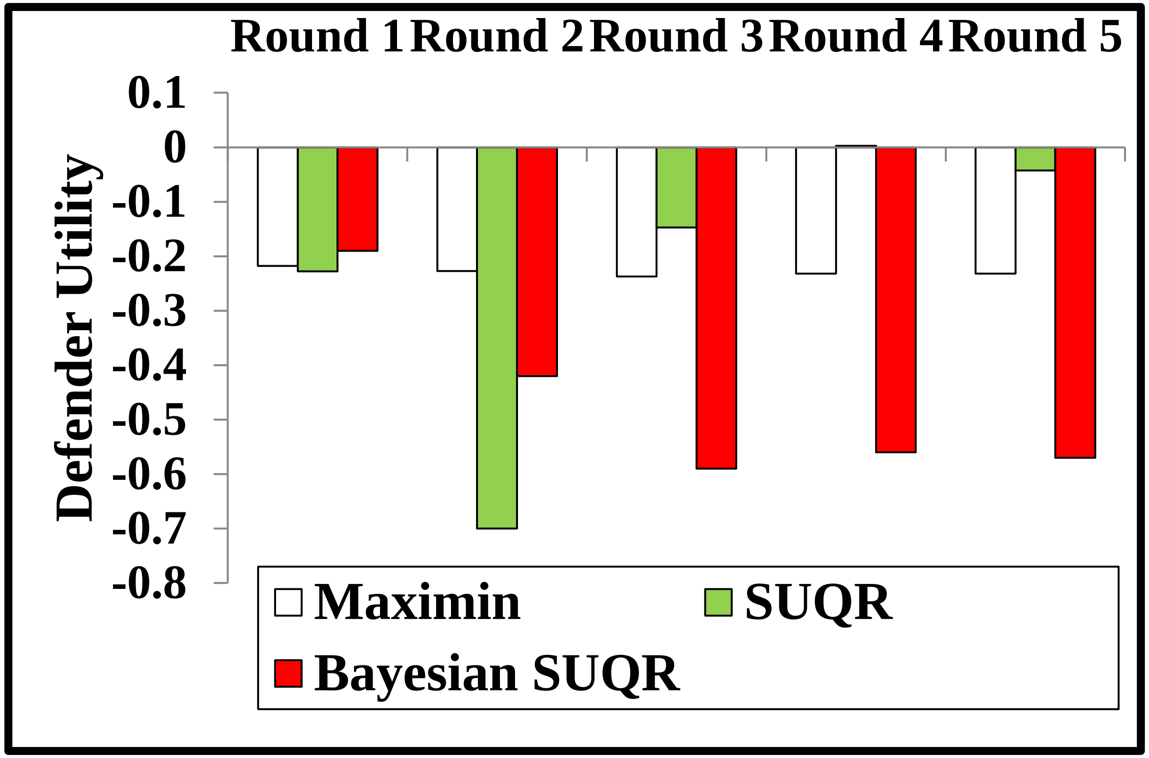

Lesson 3: SHARP and Repeated Stackelberg Games
Incorporate Past Success/Failure in SUQR

Bayesian SUQR: Heterogeneous Poachers

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

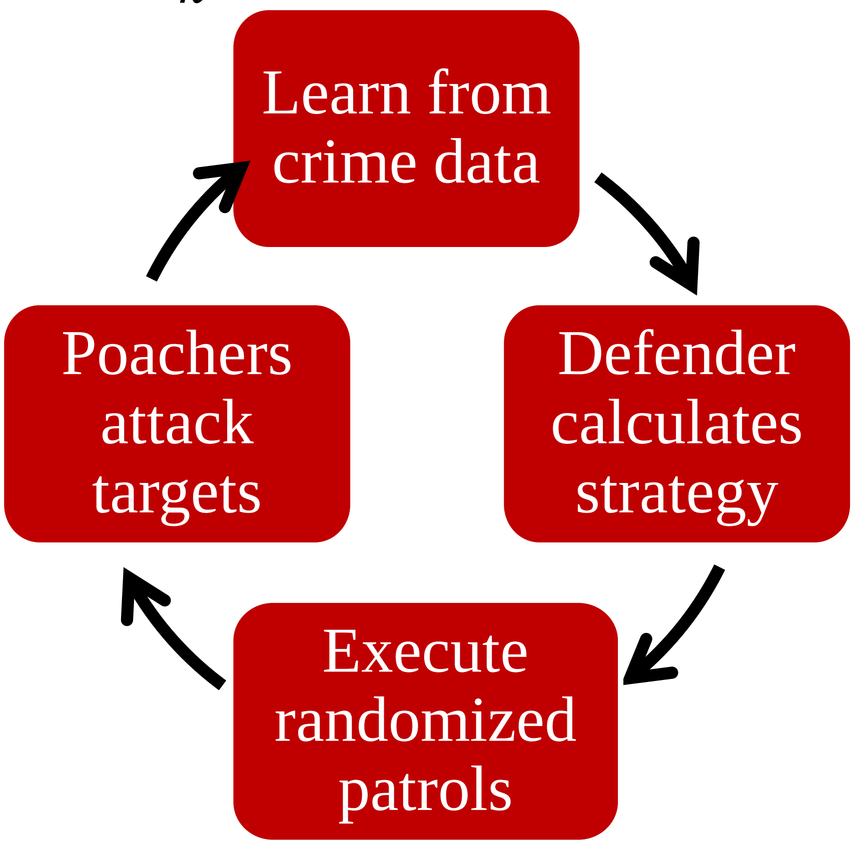

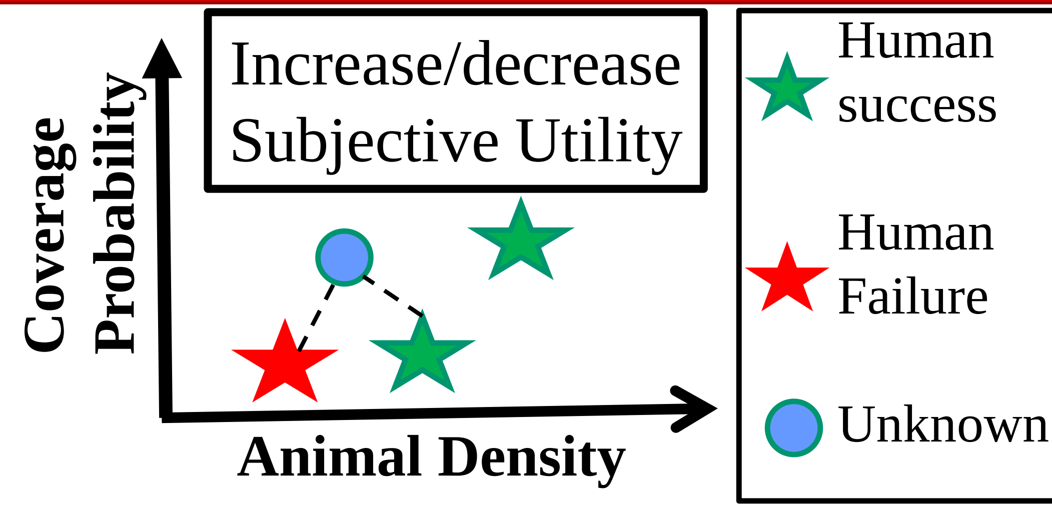

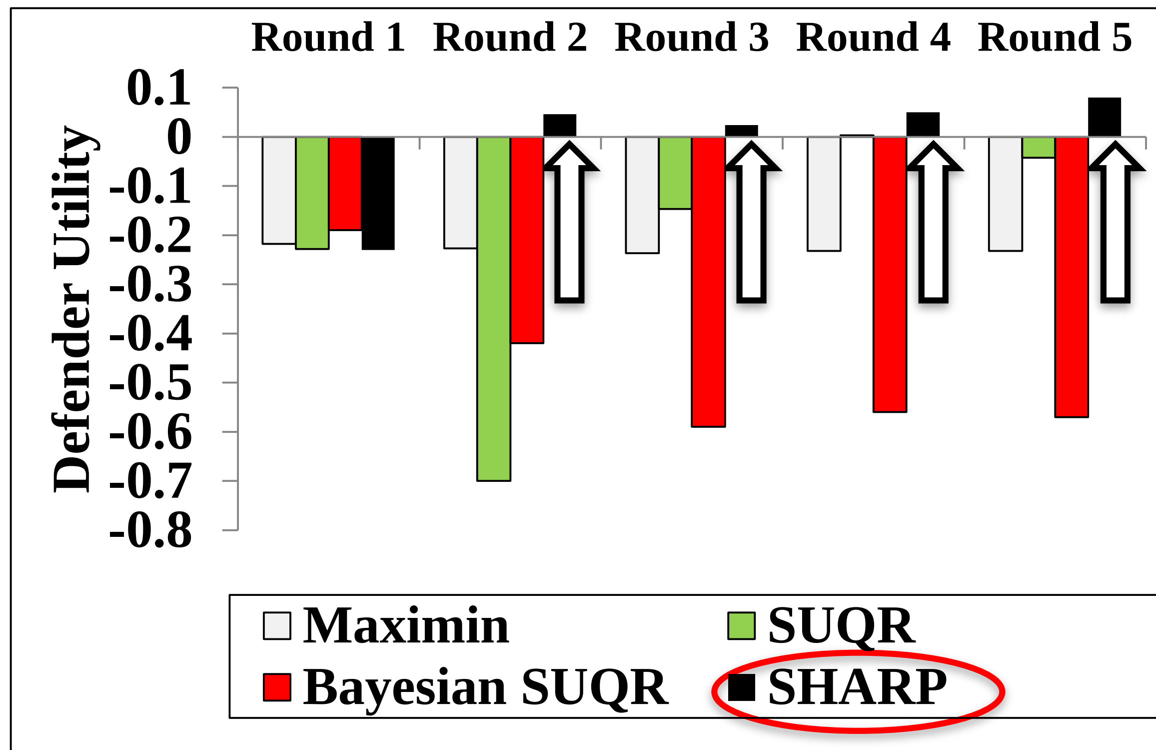

PAWS: Protection Assistant for Wildlife Security Trials in Uganda and Malaysia: [2014]

- **Important lesson: Geography!**

Uganda

Andrew Lemieux

Malaysia

Panthera

Malaysia

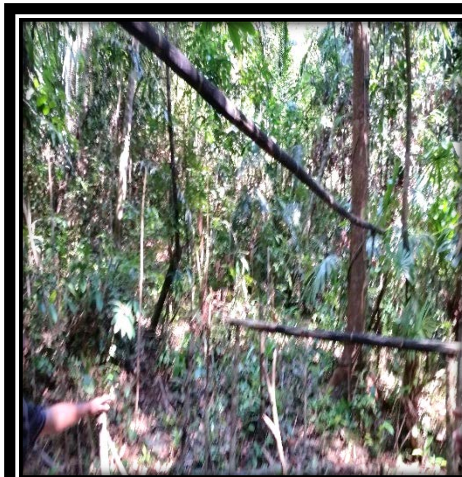

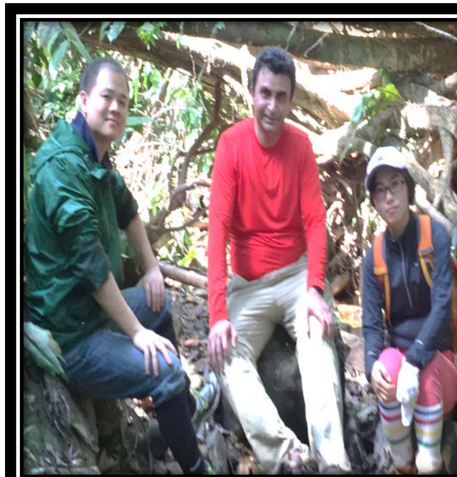

PAWS Next Steps 2015:
Southeast Asia (with Panthera and WWF)

PAWS Version 2: Features

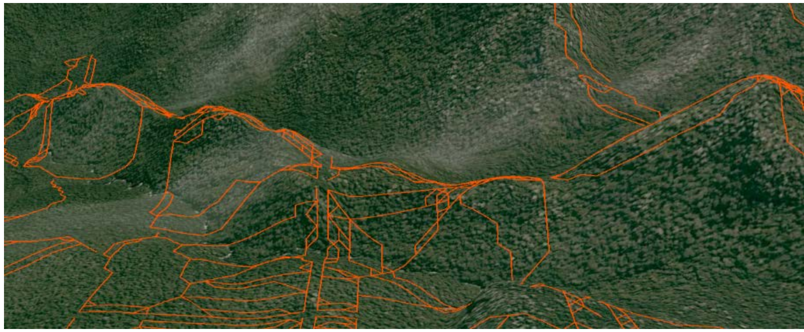

- Street map
  - *Ridgelines, rivers/streams*

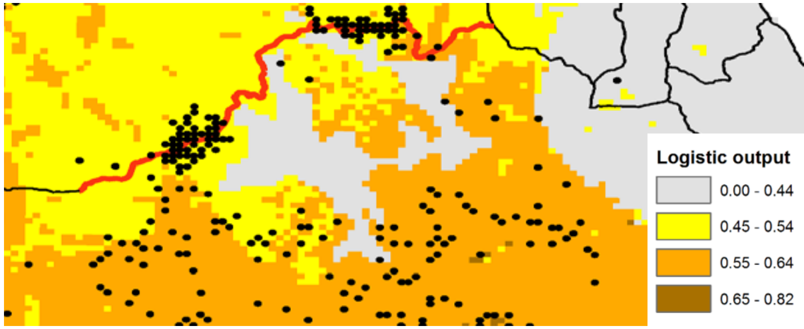

- Species Distribution Models (SDMs)
  - *From data points to distribution map*

Indonesia

Malaysia

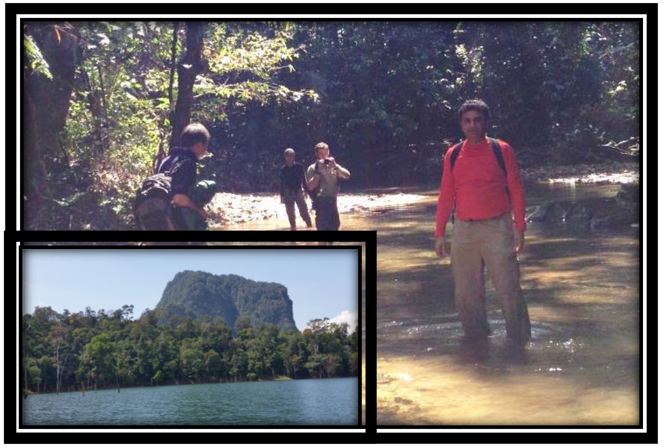

Opportunistic Crime Security Game[2015]
Urban Crime

Zhang

Sinha

- Crime prediction: uses past defender (police) allocation data

University of Southern California

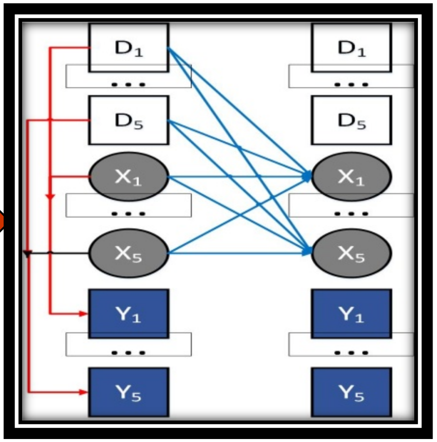

Results of learning

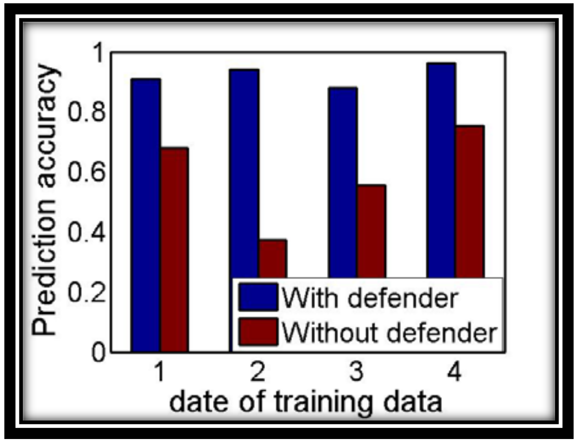

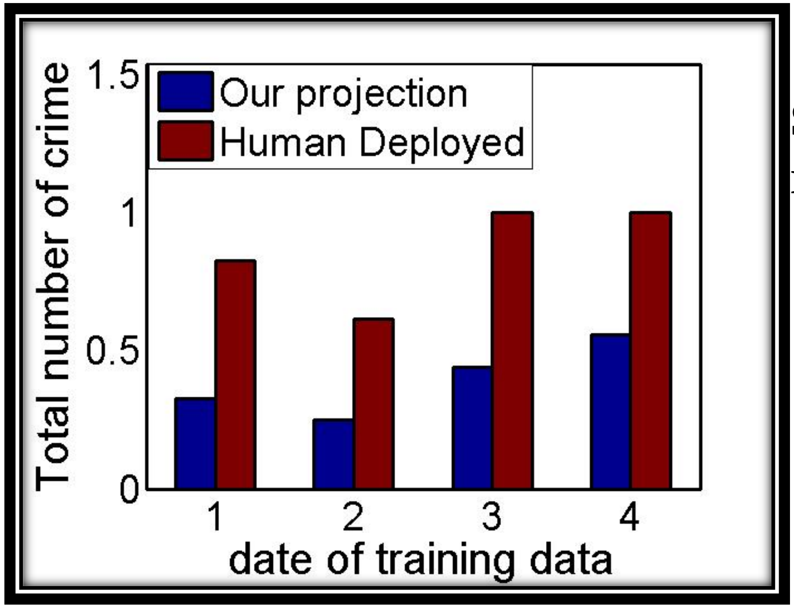

Simulation results

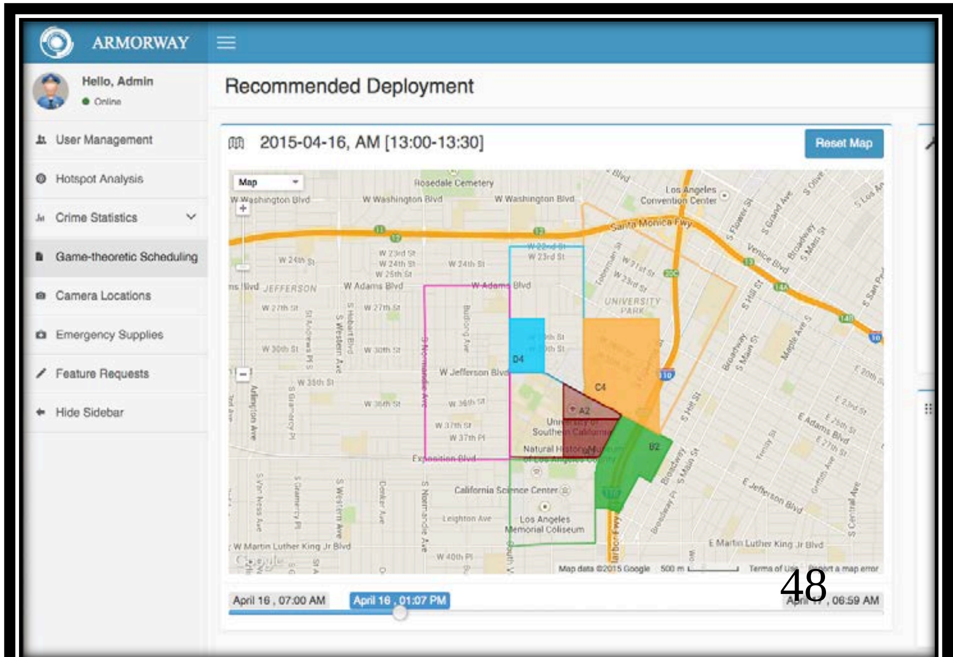

Security Resource Optimization:
Evaluating *Deployed Security Systems* Not Easy

- 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

Field Evaluation of Schedule Quality:

Improved Patrol Unpredictability & Coverage

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

Patrols Before PROTECT: Boston

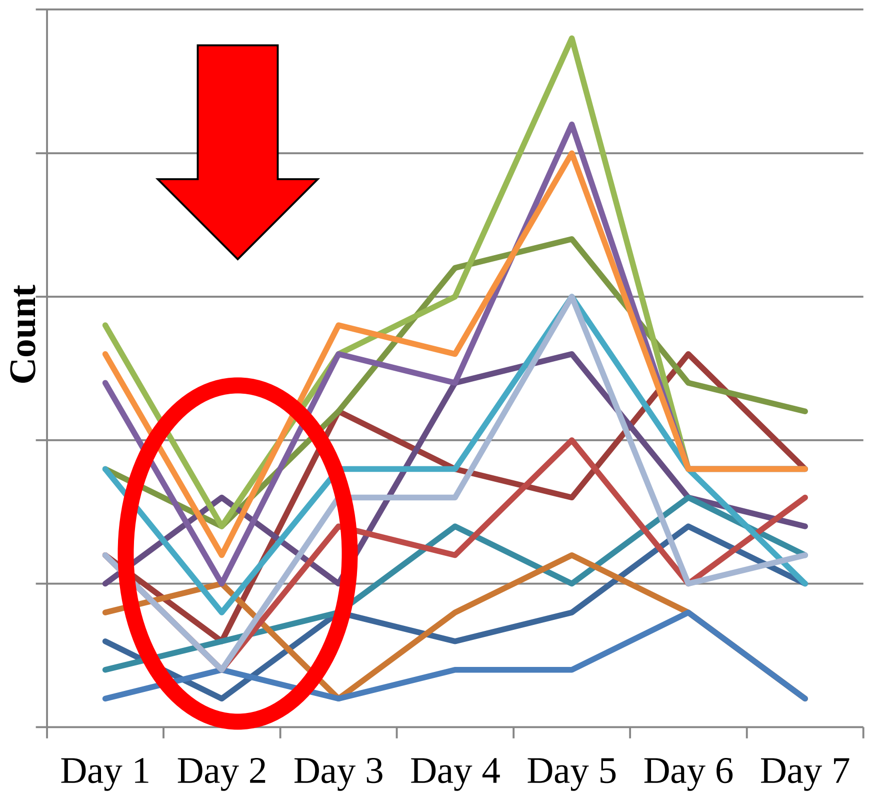

Patrols After PROTECT: Boston

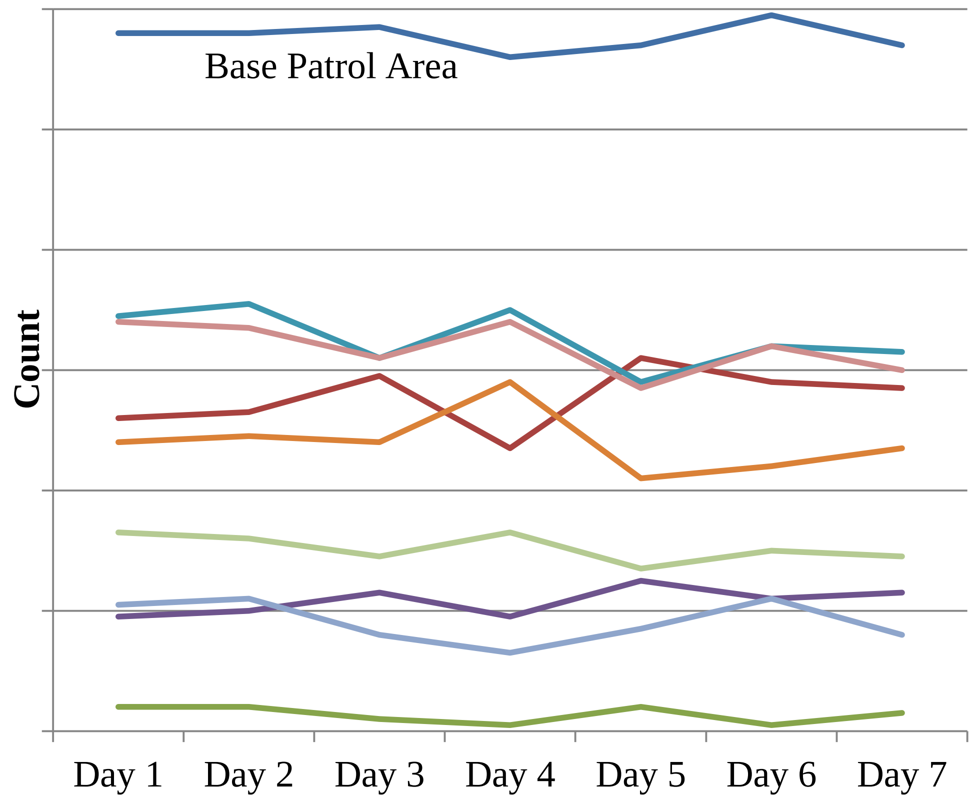

Field Evaluation of Schedule Quality:

Improved Patrol Unpredictability & Coverage for Less Effort

IRIS for FAMS: Outperformed expert human over six months

Report:GAO-09-903T

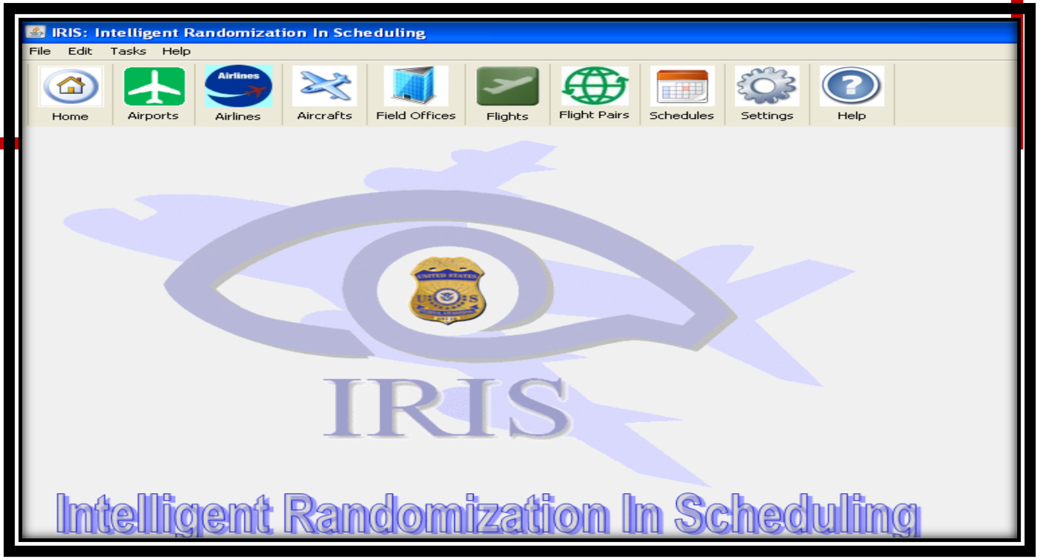

ARMOR at LAX: Savings of up to an hour a day in scheduling

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:

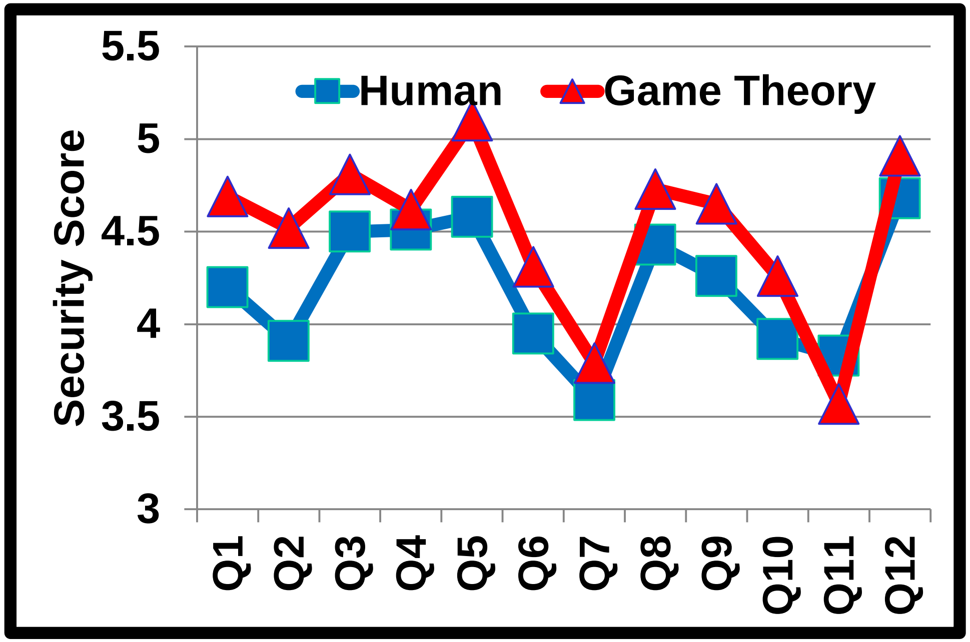

Field Test Against Adversaries: Mock Attackers Example from PROTECT

- "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*

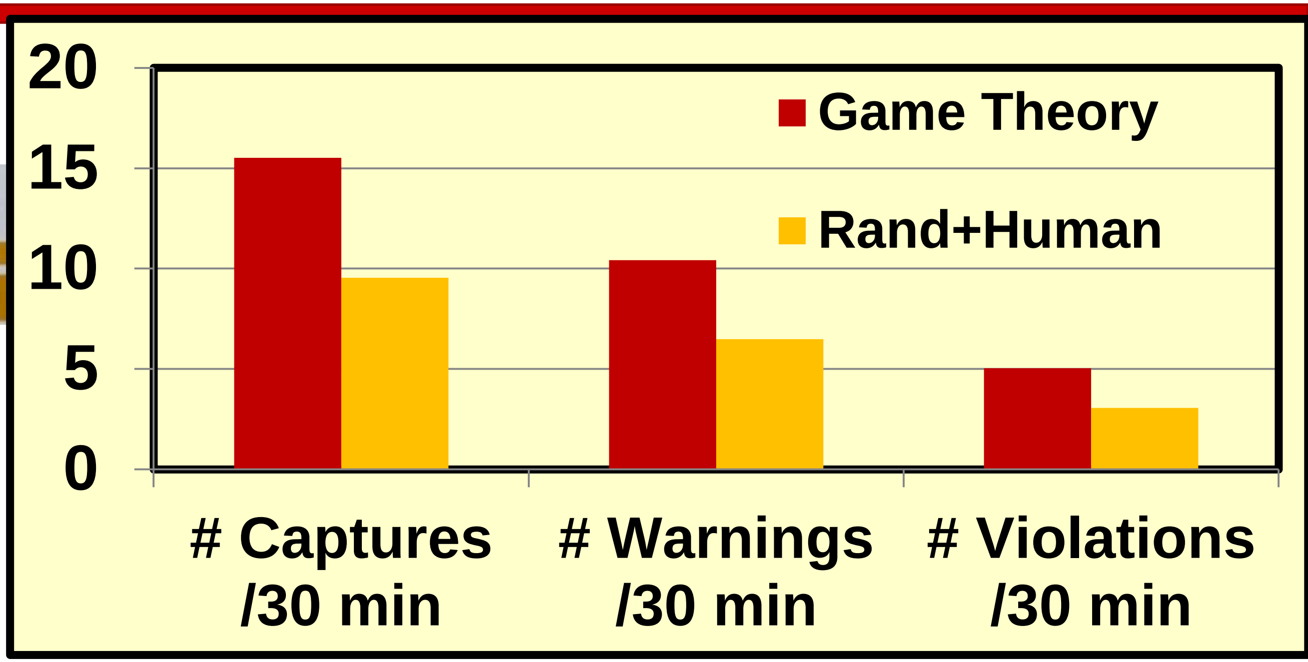

Not controlled

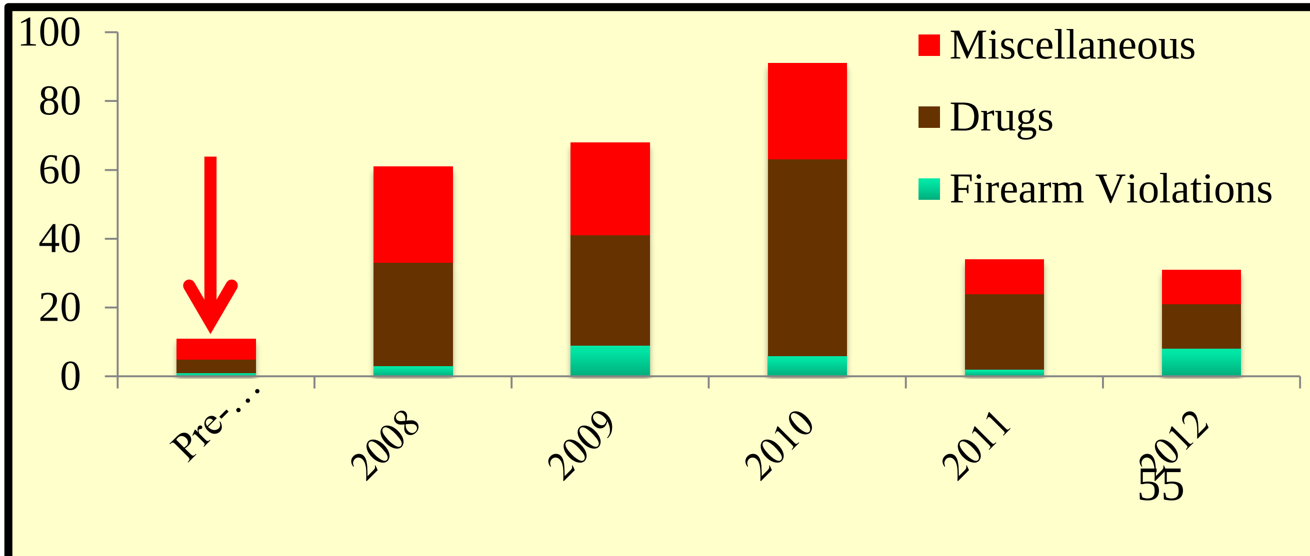

Expert Evaluation
Example from ARMOR, IRIS & PROTECT

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)

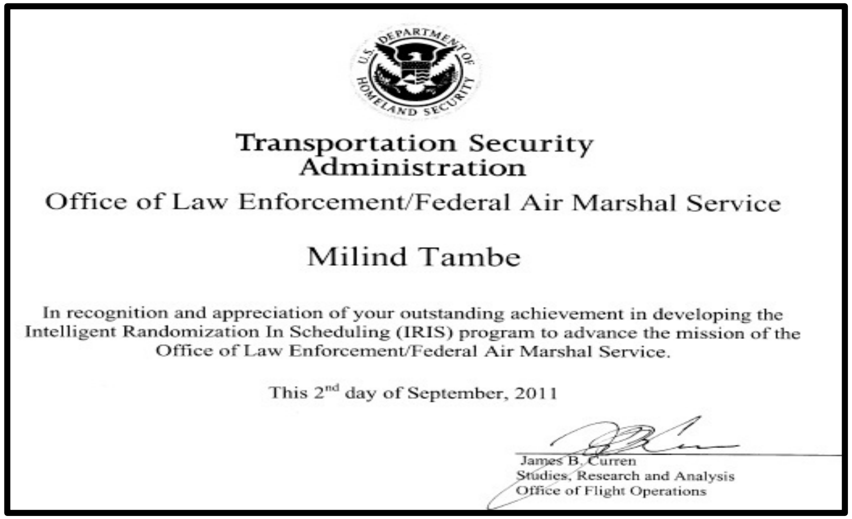

Transportation Security Administration

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)

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
- Human behavior models: Quantal response

"Green security games": Global, interdisciplinary challenges

- Intersection of criminology, computation, conservation

Global Efforts on Security Games: Yet Just the Beginning…

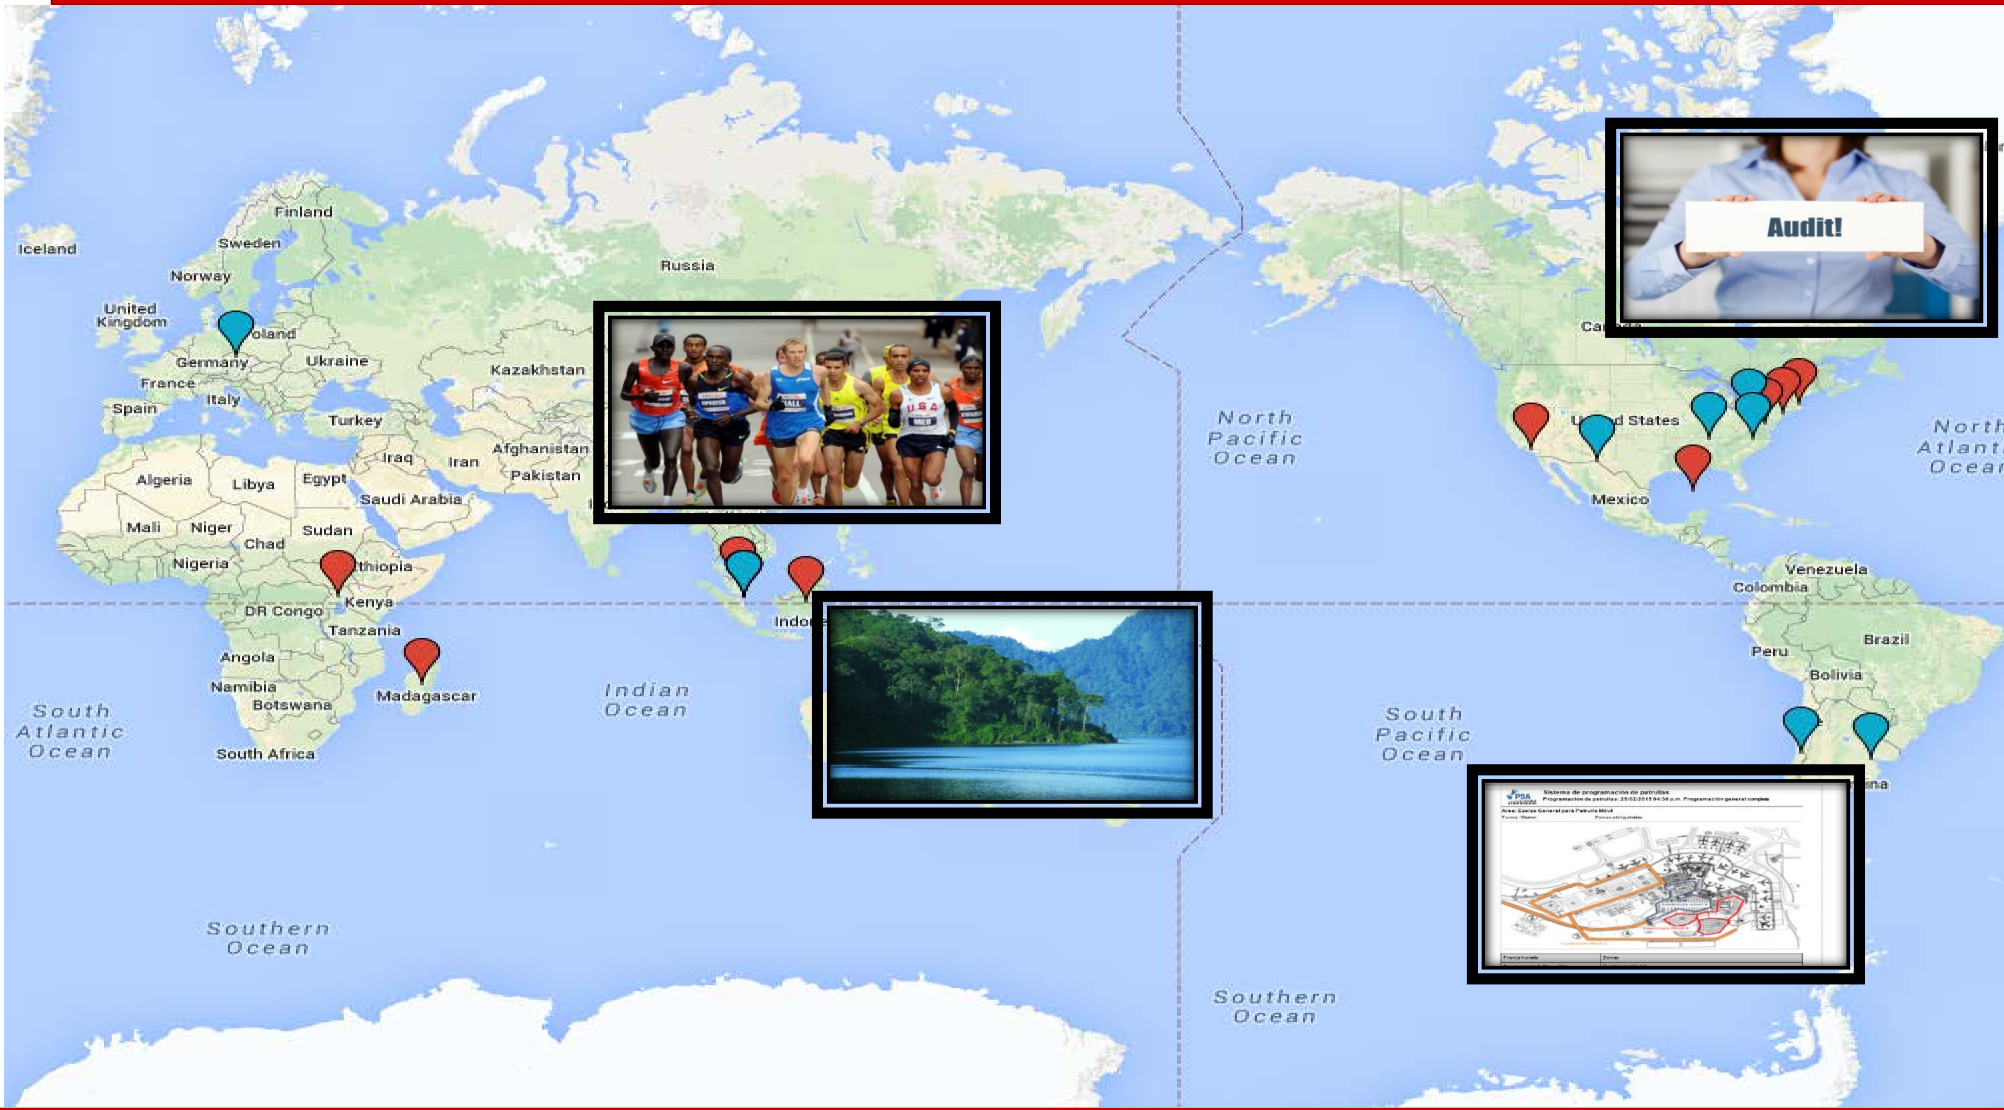

Thank you to sponsors:

CREATE HOMELAND SECURITY CENTER

Transportation Security Administration

Army Research Office

THANK YOU

tambe@usc.edu

http://teamcore.usc.edu/security

Lab Evaluation via Simulations: Example from IRIS (FAMS)

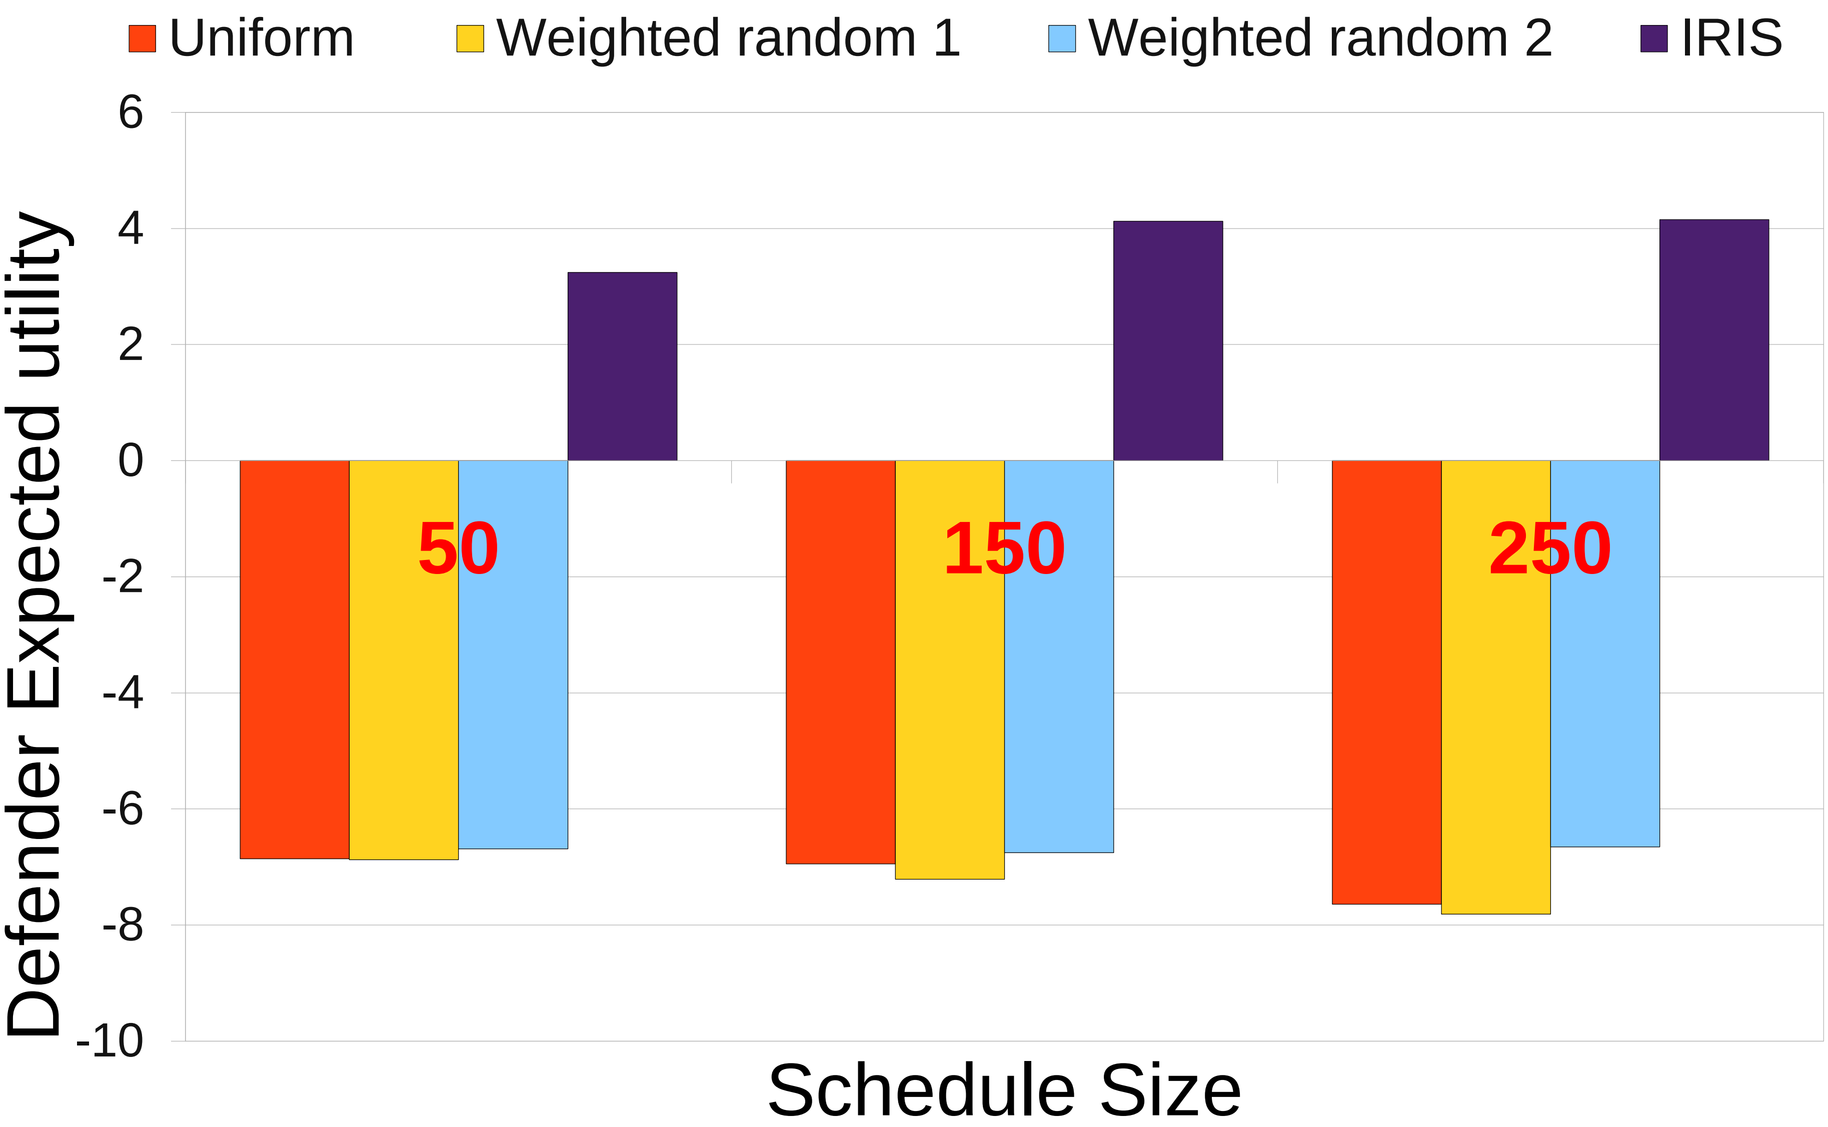

Uncertainty in Adversary Decision[2009]
Human subjects: Anchoring, ε-Optimality*

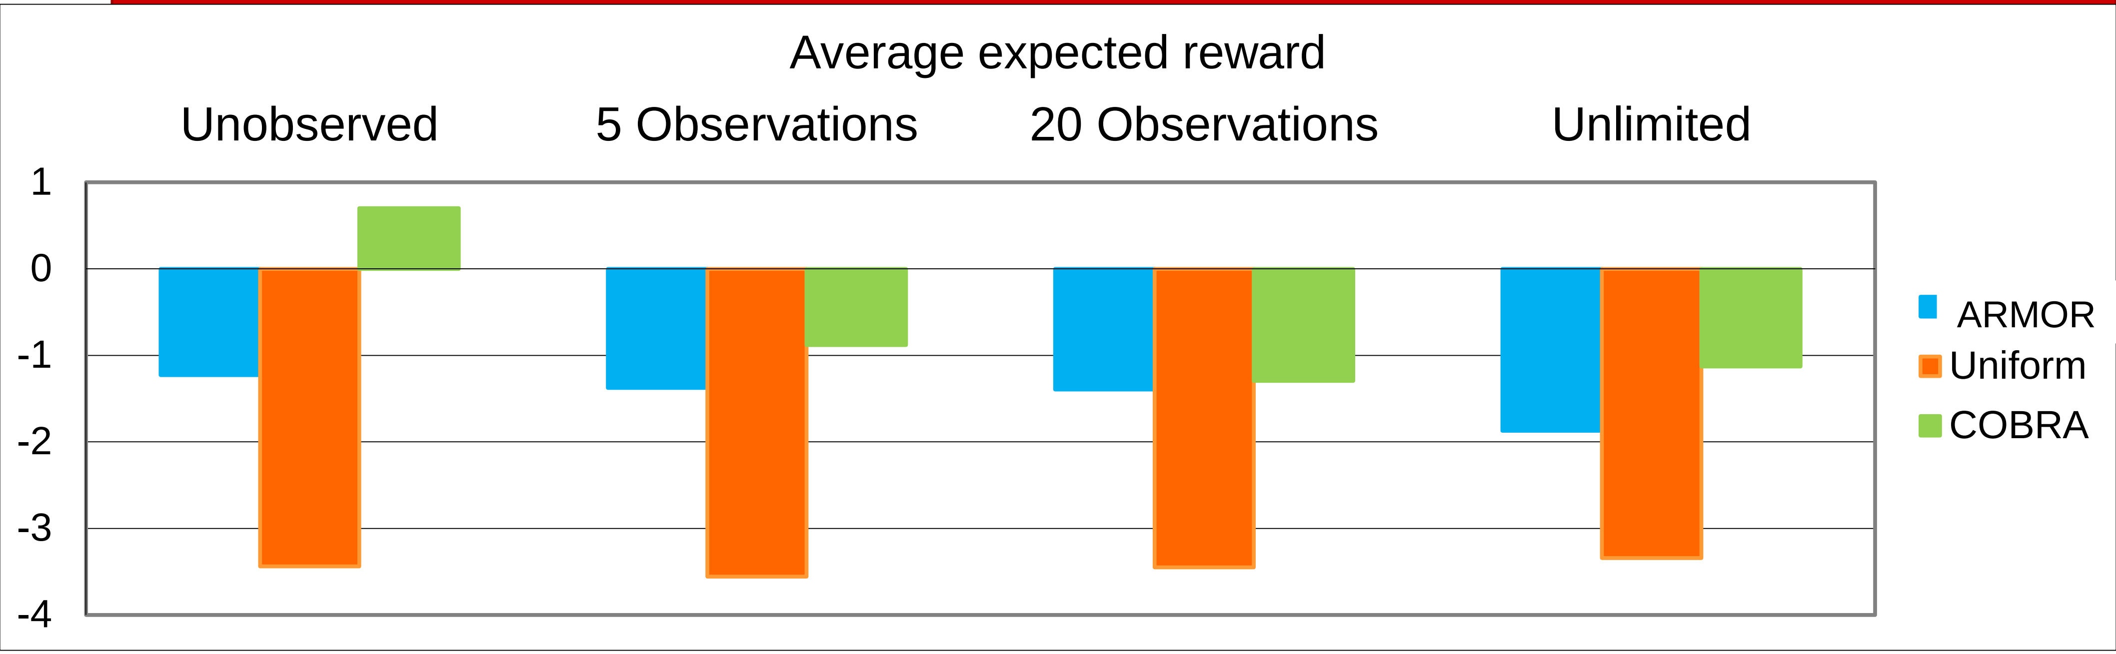

- **ARMOR:** Outperforms uniform random
- **COBRA:** $\max_{x,q} \sum_{i \in X} \sum_{j \in Q} R_{ij} x_i q_j$

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

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

Social networks:
e.g., counter-insurgency

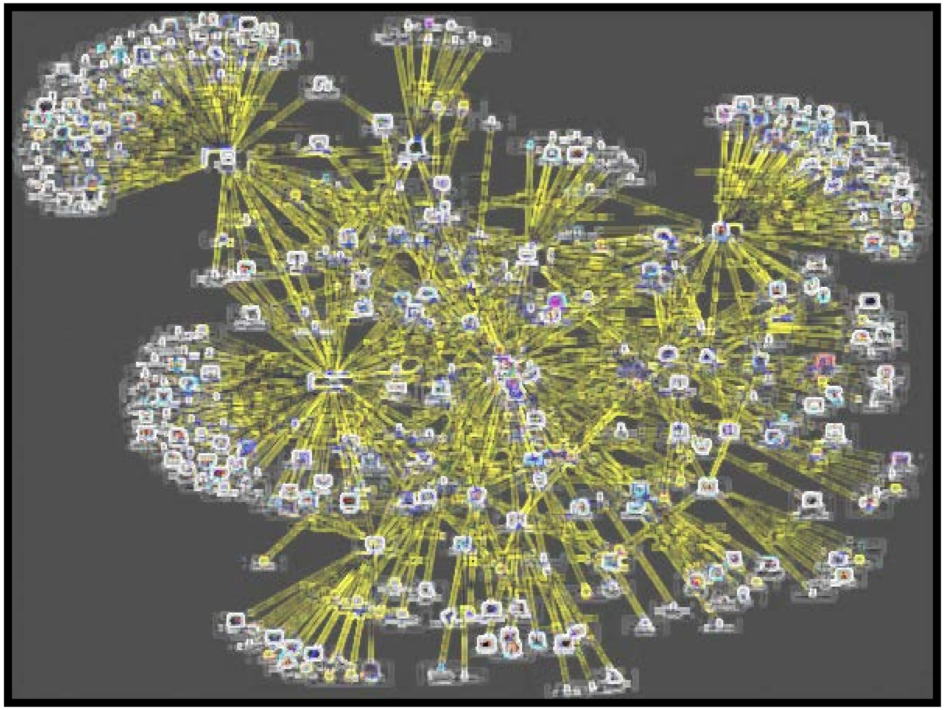

Cyber networks:

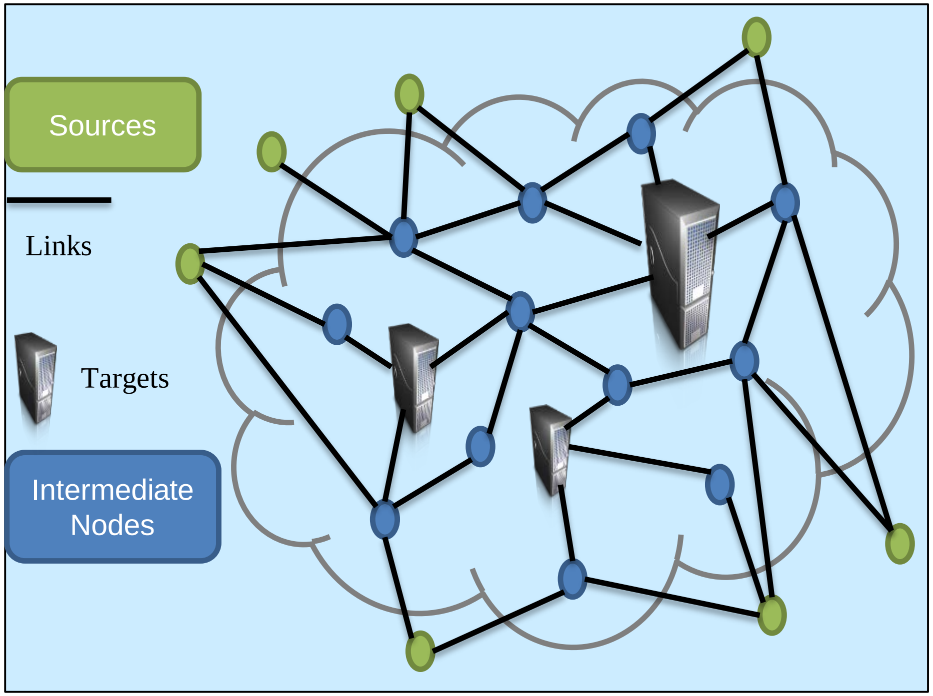

Just the Beginning of "Security Games"….

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

Thank you:

CREATE HOMELAND SECURITY CENTER
National Center for Risk and Economic Analysis of Terrorism Events

Army Research Office

Transportation Security Administration

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

- ***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

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

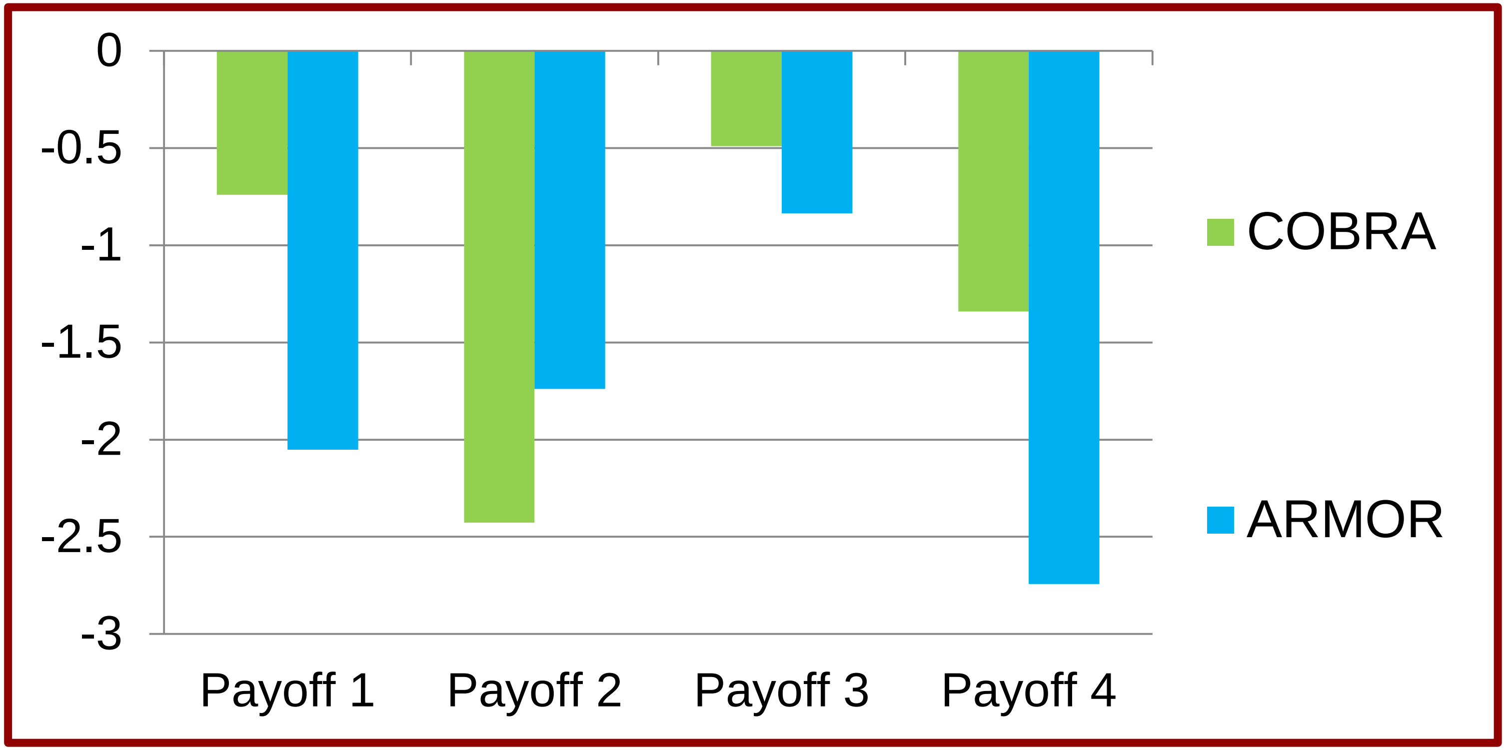

Global Presence of Security Games Efforts

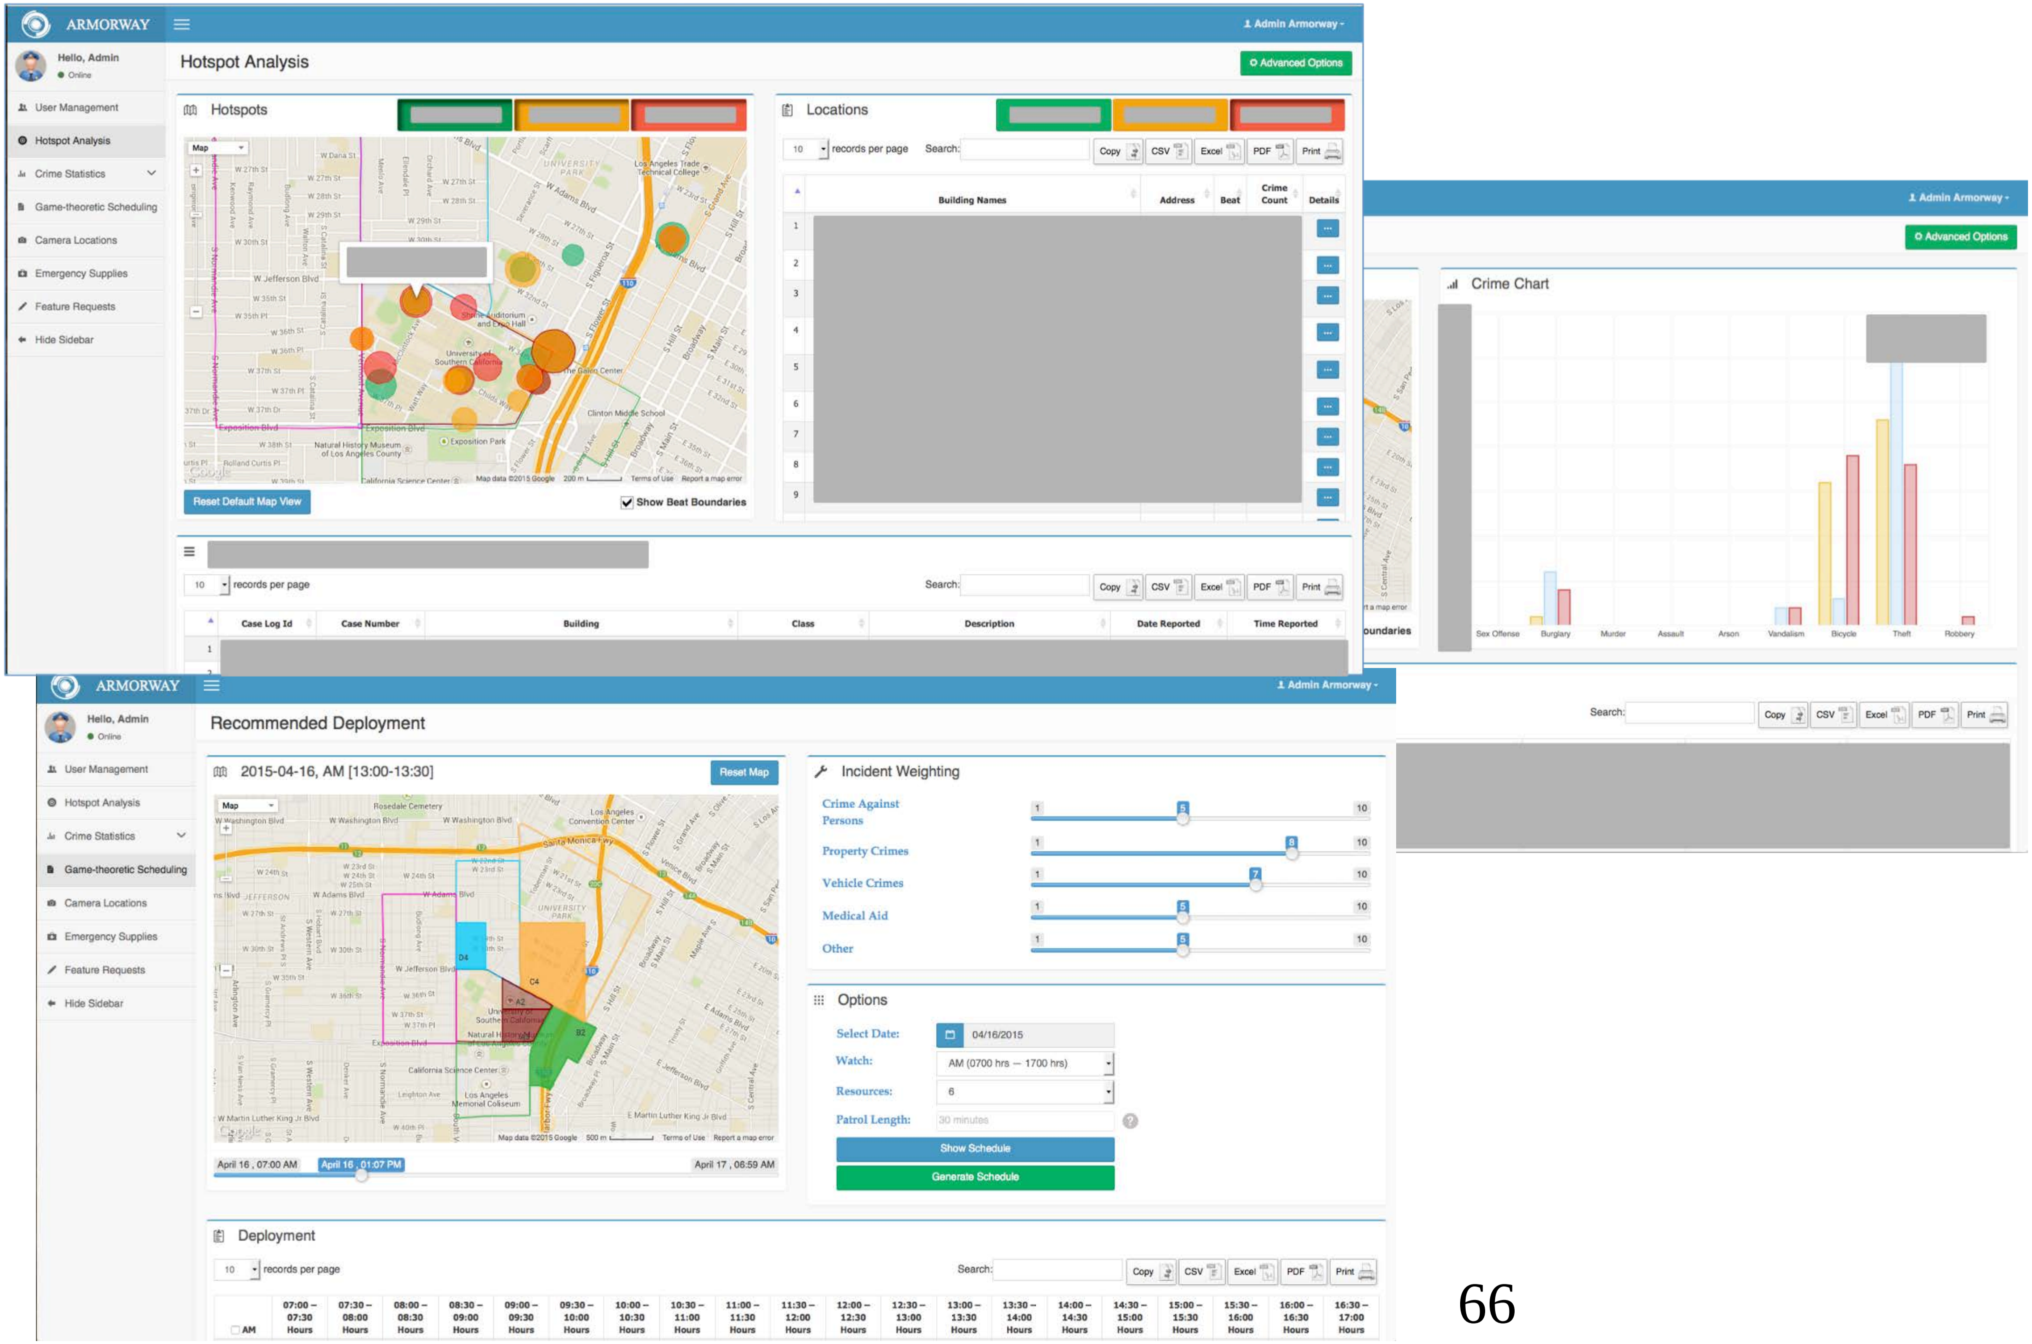

ARMORWAY
Admin Armorway
Hello, Admin
Online
User Management
Hotspot Analysis
Crime Statistics
Game-theoretic Scheduling
Camera Locations
Emergency Supplies
Feature Requests
Hide Sidebar
Hotspot Analysis
Advanced Options
Hotspots
Locations
records per page
Search:
Copy
CSV
Excel
PDF
Print
Building Names
Address
Beat
Crime Count
Details
Reset Default Map View
Show Beat Boundaries
Case Log Id
Case Number
Building
Class
Description
Date Reported
Time Reported
Crime Chart
Sex Offense
Burglary
Murder
Assault
Arson
Vandalism
Bicycle
Theft
Robbery
Recommended Deployment
2015-04-16, AM [13:00-13:30]
Reset Map
April 16 , 07:00 AM
April 16 , 01:07 PM
April 17 , 06:59 AM
Incident Weighting
Crime Against Persons
Property Crimes
Vehicle Crimes
Medical Aid
Other
Options
Select Date:
04/16/2015
Watch:
AM (0700 hrs — 1700 hrs)
Resources:
6
Patrol Length:
30 minutes
Show Schedule
Generate Schedule
Deployment
AM
07:00 – 07:30 Hours
07:30 – 08:00 Hours
08:00 – 08:30 Hours
08:30 – 09:00 Hours
09:00 – 09:30 Hours
09:30 – 10:00 Hours
10:00 – 10:30 Hours
10:30 – 11:00 Hours
11:00 – 11:30 Hours
11:30 – 12:00 Hours
12:00 – 12:30 Hours
12:30 – 13:00 Hours
13:00 – 13:30 Hours
13:30 – 14:00 Hours
14:00 – 14:30 Hours
14:30 – 15:00 Hours
15:00 – 15:30 Hours
15:30 – 16:00 Hours
16:00 – 16:30 Hours
16:30 – 17:00 Hours

Road, Social Networks[2013]
Results of Scale-up

Road networks:

Social networks:

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

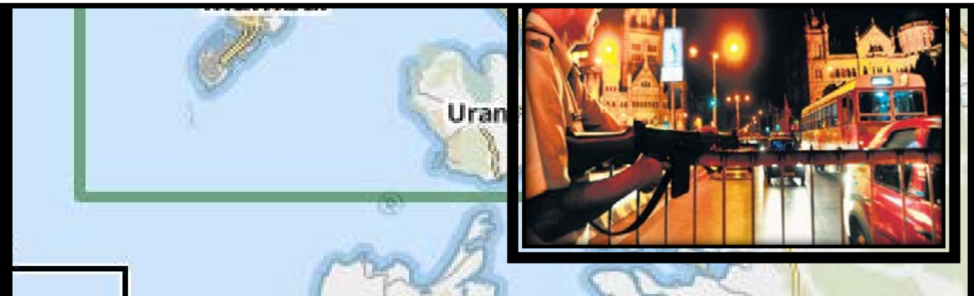

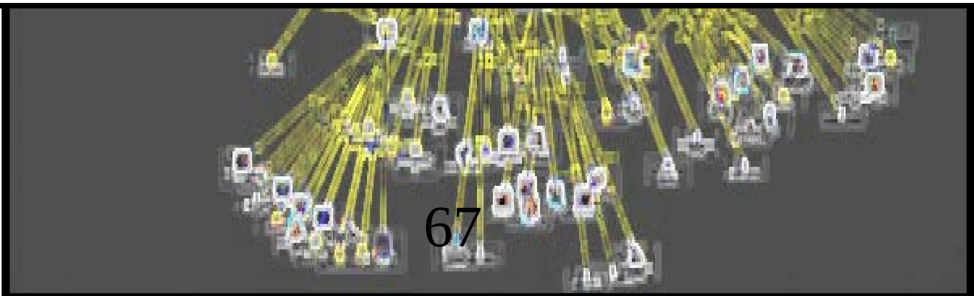

PAWS Version 2 [2015]

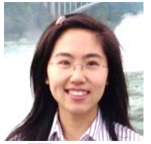

Fang

- Street map
  - *Ridgelines*
  - *Rivers and streams*
  - *Existing trails/roads*

- Species Distribution Models (SDMs)
  - *From data points to distribution map*

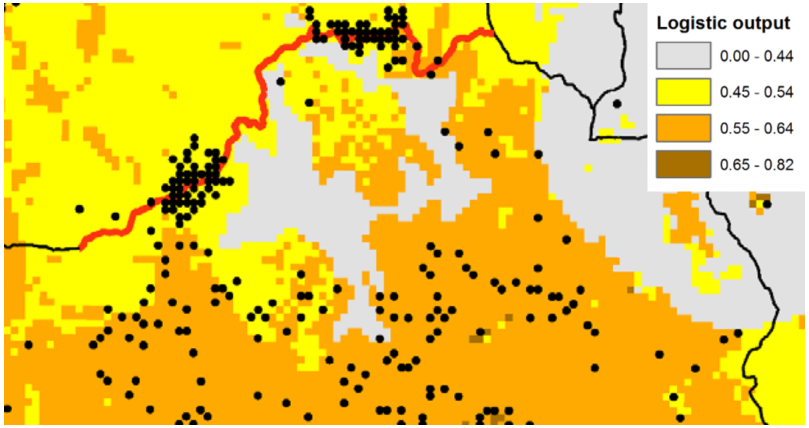

- Dynamic poacher behavior

Jan	Feb	Mar	Apr	May

Adaptive Subjective Utility of a target profile β_i for round R:

$$ASU_{\beta_i}^{R} = (1 - d * A_{\beta_i}^{R})\omega_1 f(x_{\beta_i}) + (1 + d * A_{\beta_i}^{R})\omega_2 \phi_{\beta_i} + (1 + d * A_{\beta_i}^{R})\omega_3 P_{\beta_i}^{a} + (1 - d * A_{\beta_i}^{R})\omega_4 D_{\beta_i}$$

Attractiveness of a target profile β_i for round R:

$$A_{\beta_i}^{R} = \frac{\sum_{r=1}^{R} V_{\beta_i}^{r}}{R}$$

Vulnerability of a target profile β_i for round r:

$$V_{\beta_i}^{r} = \frac{success_{\beta_i}^{r} - failure_{\beta_i}^{r}}{success_{\beta_i}^{r} + failure_{\beta_i}^{r}}$$

Computational Game Theory in the Field








PAWS: Protection Assistant for Wildlife Security Trials in Uganda and Malaysia: [2014]

- **Important lesson: Geography!**


Uganda


Andrew Lemieux


Malaysia


Panthera