USC CREATE HOMELAND SECURITY CENTER

A Game Theoretic Approach for Allocation of Limited Security Resources

Milind Tambe, University of Southern California
Fernando Ordonez, University of Southern California
tambe@usc.edu
fordon@usc.edu

1. Executive Summary

Security at major locations of economic or political importance and transportation infrastructure is a key concern around the world, particularly given the threat of terrorism. The protection of important locations includes tasks such as monitoring all entrances or checking all inbound traffic or patrolling trains and buses. However, limited security resources prevent comprehensive security coverage at all times, which allows adversaries to detect and exploit patterns in selective patrolling or monitoring, e.g., they can plan an attack by avoiding their knowledge of existing patrols. Randomizing schedules for patrolling, checking, or monitoring is thus an important tool in the police arsenal to avoid the vulnerability that comes with predictability.

In developing an automated program for randomization, we must address address three key challenges. First, we must provide distinct weights to different actions based on their complex costs and benefits. For example, if an attack on one part of an infrastructure will cause economic damage while an attack on another could potentially cost human lives, we must weigh the two options differently—giving higher weight (probability) to guarding the latter. Second, we must address the uncertainty in information that security forces have about the adversary. For example, while there may be a certain probability that a hard-core, high-capability terrorist group may be planning an attack on infrastructure, there may be a higher chance for a local gang, with lower capability and other motivations, that may be planning an attack. Third, we must take into account adversary's reaction to our randomized strategy.

To address these challenges in randomization, we rely on game theory to provide us appropriate randomization for protecting important infrastructure while providing some security guarantees. In particular, we have provided significant advances in game theoretic algorithms, crucial to deploying

"This research was supported by the United States Department of Homeland Security through the Center for Risk and Economic Analysis of Terrorism Events (CREATE) under grant number 2007-ST-061-000001. However, any opinions, findings, and conclusions or recommendations in this document are those of the authors and do not necessarily reflect views of the United States Department of Homeland Security."

these algorithms in critically important applications of counter-terrorism and homeland security. These algorithms have not only led to research advances, but also to actual deployments:

- Since August 2007 these algorithms are in use for the defense of the Los Angeles International Airport (LAX), embedded in a program called ARMOR (Assistant for Randomized Monitoring Over Routes).
- They are also currently (as of October 2009) deployed by the Federal Air Marshals for protection of international flights in a system called IRIS (Intelligent Randomization in Scheduling) in a pilot program.
- Our algorithms are also currently being deployed by the Transportation Security Administration (TSA) – the GUARDS system --- for randomizing their activities to protect airports – with LAX and Pittsburgh airports serving as pilots for evaluation. GUARDS is currently in phase I, with phase II to be completed in Feb 2010, when the system is to be fully operational.

Our game theoretic approach provides "intelligent" randomization of security forces' actions with security guarantees, significantly increasing adversary cost and uncertainty and providing a powerful deterrence. Game theoretic approaches are superior to human (police) randomization. Director Butts of LAX police points out "…over time security forces follow predictable patterns"; indeed, psychological studies show that humans are particularly bad at randomization. Of course, game-theoretic algorithms are much more powerful than simply rolling dice, which can only provide uniform or "uninformed" randomness, and are unable to take into account relative target importance (e.g. at the airport, it is more important to protect more crowded terminals than other areas) and adversary reaction to police strategy. These algorithms are also more powerful than simple weighted randomness because such weighting is unable to take into account an adversary's reaction – this is the crucial game theoretic aspect that we provide via our algorithms.

This project has provided an opportunity to marry latest game theoretic algorithms and research with practical deployed applications, with significant interest from law enforcement community. Beyond LAX police, Federal Air Marshals Service and the Transportation Security Administration, significant interest has been expressed by the Port of Los Angeles Police, the Coast Guard, and others.

2. Research Accomplishments

2.1. Game-theoretic approach

We model the interaction between the security agent and criminal adversaries as a Bayesian Stackelberg game, where the security agent is the leader and decides on a patrolling strategy first with the adversary making a decision subsequently aware of the decision of the leader. We assume that the leader may be unsure about the adversary faced and represent this with a probability distribution over possible adversary types. While the optimal policy selection for a Bayesian Stackelberg game is known to be NP-hard, our solution approach based on an efficient Mixed Integer Linear Program (MILP) provides significant speed-ups over existing approaches in obtaining the optimal solution. The resulting policy randomizes the agent's possible strategies, while taking into account the probability distribution over adversary types.

In protecting a number of targets there are multiple security resources with different capabilities collaborating in providing security. For instance in the LAX airport, in addition to the checkpoints and canine patrols, there are traffic personnel, TSA screening, LAWA police foot patrols, etc. These different resources collaborate to provide security for the airport by conducting different actions and with different capabilities. To represent this more elaborate

security action space we have expanded the Stackelberg game model to consider the fact that any one security agent does not provide perfect coverage at a location and the addition of additional resources can improve security. This new Stackelberg game, where security resources can select the same action, increases dramatically in size because there is no compact representation, as rewards are affected by the actions of the remaining agents, and because the problem considers a large number of security resources to begin with.

2.2. Evaluating the performance of security versus human opponents

Existing algorithms for Stackelberg games efficiently find optimal solutions, but they critically assume that the follower plays optimally. Unfortunately, in real-world applications, agents face human followers who — because of their bounded rationality and limited observation of the leader strategy—may deviate from their expected optimal response. Not taking into account these likely deviations when dealing with human adversaries can cause an unacceptable degradation in the leader's reward, including in security applications where these algorithms have seen real-world deployment. To address this crucial problem, we introduce three new mixed-integer linear programs (MILPs) for Stackelberg games to consider human followers, incorporating: (i) novel anchoring theories on human perception of probability distributions and (ii) robustness approaches for MILPs to address human imprecision. We rely on empirical validation to evaluate the effectiveness of the proposed methods. To that end, we consider two settings based on real deployed security systems, and compare 6 different approaches (three new with three previous approaches), in 4 different observability conditions, involving 98 human subjects playing 1360 games in total. The final conclusion was that a model which incorporates both the ideas of robustness and anchoring achieves statistically significant better rewards and also maintains equivalent or faster solution speeds compared to existing approaches.

2.3. Fast Solution Algorithms with Many Defense Resources and Scheduling Constraints

The game-theoretic methods to determine optimal randomized security policies do not scale well in general. This problem is exacerbated when in addition the defender must coordinate different resources to deploy security assets, which may be subject to complex scheduling constraints.

Given a specific structure in the payoff matrices we can obtain a compact representation of the strategies and payoffs for Bayesian Stackelberg games, and introduce a new exact solution algorithm using this representation. This is possible whenever the payoff matrix is such that the reward obtained by the leader and the adversary of an attack on a target depends only on whether that target is protected or not. This value does not depend on which other targets are protected. This compact representation leads to exponential improvements in both representation size and solution time over the best known solution algorithms for generic Stackelberg games.

We also extend the original algorithm to incorporate different types of resources with additional scheduling constraints, while still offering comparable performance improvements over previous algorithms. In the presence of resource constraints it is not always obvious how to obtain a distribution over the entire set of actions from the optimal solution to the compact representation. Since the entire set of actions is exponential in size relative to the compact representation we introduced a column generation method that seeks to obtain the distribution over the set of actions that is closest to the solution found in the compact representation. This column generation method begins considering a subset of the possible security actions and generates another action if it strictly decreases the distance to the compact representation solution. It stops when no distance reduction can be achieved by adding a new action, possibly because the distance is 0. We use this column generation method to show empirically that the compact representation is exact for large-scale real security domains.

2.4. Fast Solution Algorithms for Games with Network Structure

We have also developed new algorithms for cases where the domain is naturally represented by a graph structure. For example, subway systems, a network of city streets, and many similar situations can be modeled using a graph where nodes are possible targets (e.g., subway stations or buildings) and an attacker moves through the graph to attack a target (e.g., by taking a train or traversing roads). Security checkpoints and other measures can be deployed to prevent attacks in these domains. We have demonstrated that these domains can be modeled using a game-theoretic approach similar to ARMOR, but that the previous algorithms were unable to compute solutions for large graphs. In response, we developed new polynomial-time algorithms that exploit the graph structure to compute solutions efficiently. These algorithms compute exact optimal strategies for the defender in certain restricted cases, and empirically give very good approximate solutions even when these conditions are not met. These methods have not yet been deployed in a real domain, but were tested on a realistic case study based on the Mumbai attacks.

2.5. Different Attacker Observation Capabilities

In most work on security games to date (including our own), Stackelberg games are used to capture the possibility that the attacker will use surveillance or other methods to learn about the defense strategy when planning an attack. However, it is somewhat unrealistic to assume that the attacker always knows the defender's policy exactly; in some cases the attacker may choose not to make observations before an attack, or may use limited observations. The question is how different assumptions about the attacker's observation capabilities should change the defender's strategy. We proved theoretically that the defender's optimal strategy in a Stackelberg game is still optimal even if the attacker moves without observing the defender's strategy for several classes of games, including the games used in the ARMOR system. In other cases it is possible that the defender's optimal strategy depends on whether or not the attacker observes the strategy. However, we show experimentally that these differences are very rare, so the defender can safely play a Stackelberg strategy in most realistic situations.

3. Applied Relevance

3.1. Risk-Based Resource Allocation

ARMOR's use at LAX for the past two years has been deemed a success by LAX police and hence it is useful to discuss it in some detail. The application of ARMOR at LAX posed the following challenges:

- When and where to place checkpoints on inbound roads
- When and where to allocate canine units to terminals

These challenges bring with them some other important security concerns:

- Security is observable making any patterns or predictability vulnerable
- There are many different types of adversaries a security force will face
- These different adversary types may have differing probabilities of occurrence

ARMOR has addressed and overcome these challenges, and has been successfully deployed since August 2007 at the Los Angeles International Airport (LAX) to randomize checkpoints on the roadways entering the airport and the canine patrol routes within the airport terminals. After a successfully completed a six month trial, it was decided to "permanently" hand over ARMOR to LAX police.

Canine patrols at LAX

Checkpoints at LAX

ARMOR: Six month evaluation

ARMOR's effectiveness has been measured not only in mathematical studies such as the one shown above, but in continued evaluations by outside agencies of LAX police. Certainly, ARMOR has provided intelligent randomization, significantly increasing adversary cost and uncertainty. This is in contrast with the previous randomization approach at LAX (which was based on reliance on humans --- which on closer examination may have been more predictable).

ARMOR has also resulted in the numbers of arrests made of people attempting to carry guns and narcotics into the airport. The arrest record for the month of January 2009 was exceptional:

- January 3rd Loaded 9/mm pistol discovered in a car
- January 9th 16-handguns, 4-rifles, 1-black-powder pistol, 1-assault rifle (some fully loaded); 1000 rounds of ammunition.
- January 10th Two unloaded shotguns (no arrest)
- January 12th Loaded 22/cal rifle discovered in a car
- January 17th Loaded 9/mm pistol discovered in a car
- January 22nd Unloaded 9/mm pistol in a car (no arrest)

ARMOR has received significant external recognition; some of the awards are listed here:

- Commendation from the city of Los Angeles, given by the LAX police department, 2009
- Department of Homeland Security University Programs "Certificate of Recognition", 2009
- USC Viterbi School of Engineering "Use-inspired research award", 2009
- Best paper, industry track, International Conference on Autonomous Agents and Multiagent Systems, 2009
- Finalist for best paper, industry track, International Conference on Autonomous Agents and Multiagent Systems, 2008

4. Collaborative Projects

- We have deployed the ARMOR-Checkpoints and ARMOR-K9 in collaboration with the Los Angeles World Airport (LAWA) police. This collaboration has involved extensive meetings

with LAWA police personnel to calibrate the data and include specific constraints of each problem to the DOBSS model. We have provided LAWA police with alternative schedules of vehicle checkpoints and K-9 patrols obtained from the ARMOR software and have received feedback on the ease of use and effectiveness of the recommendations. In February we handed over the ARMOR software to LAWA police.

- We have had a close collaboration with Federal Air Marshals (FAM). This has resulted in the delivery of the IRIS program to the FAMS.

- In addition, we have had a close collaboration with the Transportation Security Administration (TSA), and have delivered the GUARDS software to them.

We have participated in the establishment of a Security Center in collaboration between the University of Chile and the Chilean Interior Minister. Among other projects this collaboration is aimed at facilitating a deployment of the ARMOR approach to police patrols in the city of Santiago.

In addition we have had inquiries from the Coast Guard, Port of Los Angeles Police, LA County Sheriff Department, NY/NJ Port Authority about ARMOR and its possible use.

5. Research Products

Research Products (Please detail below)		#
5a	# of peer-reviewed journal reports published	10
5a	# of peer-reviewed journal reports accepted for publication	
5a	# of non-peer reviewed publications and reports	1
5a	# of scholarly journal citations of published reports	115
5b	# of scholarly presentations (conferences, workshops, seminars)	13
5b	# of outreach presentations (non-technical groups, general public)	
5c	# of products delivered to DHS, other Federal agencies, or State/Local	3
5c	# of patents filed	2
5c	# of patents issued	
5c	# of products in commercialization pipeline (products not yet to market)	
5c	# of products introduced to market	

5.1. Publications and Reports

CREATE PUBLICATIONS	**Research Area**	**Referred**	**Not Referred**	**PDF Available for DHS**
Tambe, Milind; Ordóñez, Fernando - University of Southern California				
1. Taylor, M., Kiekintveld, C., Western, C., .Tambe, M., "A Framework for Evaluating Deployed Security Systems: Is There a Chink in Your ARMOR?" to appear, *In Informatica*	RM	x		x

CREATE PUBLICATIONS	Research Area	Referred	Not Referred	PDF Available for DHS
2. Kiekintveld, C., Pita, J., Jain, M., Tsai, J., Tambe, M., Ordonez, F., "Computing Optimal Randomized Resource Allocations for Massive Security Games," in proceedings of the *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2009	RM	x		x
3. Pita, J., Jain, M., Tambe, M., Ordonez, F., Kraus, S.,Magori-Cohen, R., "Effective Solutions for Real-World Stackelberg Games: When Agents Must Deal with Human Uncertainties," in proceedings of the *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2009	RM	x		x
4. Schurr, N., Marecki, J., Tambe, M., "Improving Adjustable Autonomy Strategies for Real-world Domains," in proceedings of the *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2009	RM	x		x
5. Tsai, J., Rathi, S., Kiekintveld, C., Tambe, M., Ordonez, F., "IRIS: A Tool for Strategic Security Allocation in Transportation Networks," in proceedings of the *International Conference on Autonomous Agents and Multiagent Systems (AAMAS Industry Track)*, Best Paper Award, industry track, May 2009	RM	x		x
6. Marecki, J., Tambe, M., "Planning with Continuous Resources for Agent Teams," in proceedings of the *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2009	RM	x		

5.2. Presentations - Conferences

1. Tambe, M., Ordonez, F., "Software Assistants for Real Patrol Planning Presentation," at *INFORMS* annual meeting, San Diego, October 2009
2. Kiekintveld, C., Pita, J., Jain, M., Tsai, J., Tambe, M., Ordonez, F., "Computing Optimal Randomized Resource Allocations for Massive Security Games," presentation at the *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2009
3. Pita,J., Jain, M., Tambe, M., Ordonez, F., Kraus, S ., Magori-Cohen, R., "Effective Solutions for Real-World Stackelberg Games: When Agents must Deal with Human Uncertainties," presentation at the *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2009
4. Schurr, N., Marecki, J., Tambe, M., "Improving Adjustable Autonomy Strategies for Real-world Domains," presentation at *the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2009
5. Tsai, J., Rathi, S., Kiekintveld, C., Tambe, M., Ordonez, F., "IRIS: A tool for Strategic Security Allocation in Transportation Networks," presentation at the *International Conference on Autonomous Agents and Multiagent Systems (AAMAS Industry Track)*, May 2009

6. Marecki, J., Tambe, M., "Planning with Continuous Resources for Agent Teams," presented at the *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2009
7. Tambe, M., Ordonez, F., "Models and Algorithms for Stackelberg Games with Incomplete Information," presentation at *INFORMS* annual meeting, October 2008

5.3. Presentations – Outreach

1. Tambe, M., "Multiagent Systems: Lessons Learned from Putting Theory into Practice," *International Workshop on Agent Design: Adapting from Practice to Theory (ADAPT)*, held in conjunction with *AAMAS '09*, 2009
2. Tambe, M., "Optimizing Multiagent Systems," *International Workshop on Optimization in Multiagent Systems (OPTMAS)*, held in conjunction with *AAMAS '09*, 2009

5.4. Models, Databases, and Software Tools and Products

Date Delivered	Item	Agency Receiving Product	Agency POC	Commercialization Status (D-delivered, P-Pipeline, M-Market)
Aug 2007	ARMOR: Assistant for Randomized monitoring over routes. Software to randomize scheduling of canines and checkpoints.	LAWA police	Director James Butts	D-P
Summer 2009	IRIS: Intelligent Randomization in International Scheduling. Randomize allocation of air marshals to flights	Federal Air Marshals Service	James Curren	D-P
Fall 2009	GUARDS: Game-theoretic Unpredictable and Randomly Deployed Security. For randomization of TSA's playbook activities.	Transportation security administration	Erin Steigerwald	D

- United States Patent application "DOBBS (Decomposed Optimal Bayesian Stackelberg Solver) is an Optimal Algorithm for Solving Bayesian Stackelberg Games" (co-inventors: M. Tambe, P. Paruchuri, F. Ordonez, J. P.Pearce, J. Marecki, S. Kraus)

- United States Patent application "ASAP Agent Security Via Approximate Policies Algorithm Is An Approximate Solver for Bayesian Stackelberg Games" (co-inventors: M. Tambe, P. Paruchuri, F. Ordonez, J. P. Pearce, J. Marecki, S. Kraus)

6. Education and Outreach Products

Education and Outreach Initiatives (Please detail below)	#
# of students supported (funded by CREATE)	10
# of students involved (funded by CREATE + any other programs)	18
# of students graduated	1
# of contacts with DHS, other Federal agencies, or State/Local (committees)	7

All students/postdoctoral researchers involved:

Postdoctoral researchers:
- Chris Kiekintveld
- Matthew Taylor

PhD students:
- Manish Jain
- James Pita
- Zhengyu Yin
- Scott Alfeld
- Jason Tsai

MS students:
- Shyamsundar Rathi
- Harish Bellamane
- Prakhar Garg
- Bharat Patel

Undergraduate students:
- Craig Western
- Alyssa Gottlieb
- Dana Li
- Michael Scott
- Andrew Ogden
- Arjun Srinivasan
- Goerg Ristock

7. Additional Information for DHS Data Base

Did project involve human subjects? No

Identify **three keywords or keyword phrases** to describe this project for keyword searches:
- Keyword 1: Game theory
- Keyword 2: Randomization
- Keyword 3: Security resource allocation