Getting Started on a Real-World Challenge Problem in Computational Game Theory and Beyond[*]

Milind Tambe, Bo An
Computer Science Department
University of Southern California
Los Angeles, CA 90089
{tambe, boa}@usc.edu

Abstract

The goal of this paper is to introduce a real-world challenge problem for researchers in multiagent systems and beyond, where our collective efforts may have a significant impact on activities in the real-world. The challenge is in applying game theory for security: Our goal is not only to introduce the problem, but also to provide exemplars of initial successes of deployed systems in this challenge problem arena, some key open research challenges and pointers to getting started in this research.

Introduction

Security is a critical concern around the world that arises in protecting our ports, airports, transportation or other critical national infrastructure from adversaries, in protecting our wildlife and forests from poachers and smugglers, and in curtailing the illegal flow of weapons, drugs and money; and it arises in problems ranging from physical to cyber-physical systems. In all of these problems, we have limited security resources which prevent full security coverage at all times; instead, limited security resources must be deployed intelligently taking into account differences in priorities of targets requiring security coverage, the responses of the adversaries to the security posture and potential uncertainty over the types, capabilities, knowledge and priorities of adversaries faced.

Game theory is well-suited to adversarial reasoning for security resource allocation and scheduling problems. Casting the problem as a Bayesian Stackelberg game, we have developed new algorithms for efficiently solving such games to provide randomized patrolling or inspection strategies. These algorithms have led to some initial successes in this challenge problem arena, leading to advances over previous approaches in security scheduling and allocation, e.g., by addressing key weaknesses of predictability of human schedulers. These algorithms are now deployed in multiple applications: ARMOR has been deployed at the Los Angeles International Airport (LAX) since 2007 to randomizes checkpoints on the roadways entering the airport and canine patrol

routes within the airport terminals (Pita et al. 2008); IRIS, is a game-theoretic scheduler for randomized deployment of the US Federal Air Marshals (FAMS) requiring significant scale-up in underlying algorithms has been in use since 2009 (Tsai et al. 2009); PROTECT, which uses a new set of algorithms based on quantal-response is deployed in the port of Boston for randomizing US coast guard patrolling (An et al. 2011b); PROTECT is now headed to New York; and GUARDS is under evaluation for national deployment by the US Transporation Security Administration (TSA) (Pita et al. 2011). These initial successes point the way to major future applications in a wide range of security arenas; with major research challenges in scaling up our game-theoretic algorithms, to addressing human adversaries' bounded rationality and uncertainties in action execution and observation, as well as in preference elicitation and multiagent learning.

This paper will provide pointers to our algorithms, key research challenges and how to get started in this research. While initial research has made a start, a lot remains to be done; yet these are large-scale interdisciplinary research challenges that call upon multiagent researchers to work with researchers in other disciplines, be "on the ground" with domain experts, and examine real-world constraints and challenges that cannot be abstracted away. Together as an international community of multiagent researchers, we can accomplish more!

Deployed and Emerging Security Applications

The last several years have witnessed the successful application of multi-agent systems in allocating limited resources to protect critical infrastructures (Basilico, Gatti, and Amigoni 2009; Korzhyk, Conitzer, and Parr 2010; Jain et al. 2010b; Pita et al. 2011; An et al. 2011b). The framework of game-theory (more precisely, Stackelberg games) is well suited to formulate the strategic interaction in security domains in which there are usually two players: the security force (defender) commits to a security policy first and the attacker (e.g., terrorist, poacher and smuggler) conducts surveillance to learn the policy and then takes his best attacking action.[1] Stackelberg games have been widely used for modeling/reasoning complex security problems and a variety of

---

---

[1]Or the attacker may be sufficiently deterred and dissuaded from attacking the protected target.

algorithms have been proposed to efficiently compute the equilibrium strategy, i.e., defender's best way of utilizing her limited security resources (there is actually a special class of Stackelberg games that often gets used in these security domains, and this class is referred to as security games). In the rest of this section, we describe the application of the Stackelberg game framework in multiple significant security domains.

ARMOR for Los Angeles International Airport

Los Angeles International Airport (LAX) is the largest destination airport in the United States and serves 60-70 million passengers per year. The LAX police use diverse measures to protect the airport, which include vehicular checkpoints and police units patrolling with canines. The eight different terminals at LAX have very different characteristics, like physical size, passenger loads, foot traffic or international versus domestic flights. Furthermore, the numbers of available vehicle checkpoints and canine units are limited by resource constraints. Thus it is challenging to optimally allocate these resources to improve their effectiveness while avoiding patterns in the scheduled deployments.

The ARMOR system (Assistant for Randomized Monitoring over Routes) focuses on two of the security measures at LAX (checkpoints and canine patrols) and optimizes security resource allocation using Bayesian Stackelberg games. Take the vehicle checkpoints model as an example. Assume that there are n roads, the police's strategy is placing $m < n$ checkpoints on these roads where m is the maximum number of checkpoints. The adversary may potentially choose to attack through one of these roads. ARMOR models different types of attackers with different payoff functions, representing different capabilities and preferences for the attacker. ARMOR uses DOBSS (Decomposed Optimal Bayesian Stackelberg Solver) to compute the defender's optimal strategy (Paruchuri et al. 2008). ARMOR has been successfully deployed since August 2007 at LAX to randomize checkpoints on the roadways entering the airport and canine patrol routes within the airport terminals.

IRIS for US Federal Air Marshals Service

The US Federal Air Marshals Service (FAMS) allocates air marshals to flights originating in and departing from the United States to dissuade potential aggressors and prevent an attack should one occur. Flights are of different importance based on a variety of factors such as the numbers of passengers, the population of source/destination, international flights from different countries, and special events that can change the risks for particular flights at certain times. Security resource allocation in this domain is significantly more challenging than for ARMOR: a limited number of FAMS need to be scheduled to cover thousands of commercial flights each day. Furthermore, these FAMS must be scheduled on tours of flights that obey various constraints (e.g., the time required to board, fly, and disembark). Therefore, we face significant computational challenge while generating the optimal scheduling policy that meets these scheduling constraints.

(a) PROTECT is being used in Boston

(b) Extending PROTECT to NY

Figure 1: USCG boats patrolling the ports of Boston and NY

Against this background, the IRIS system (Intelligent Randomization In Scheduling) has been developed and has been deployed by FAMS since October 2009 to randomize schedules of air marshals on international flights. In IRIS, the targets are the set of n flights and the attacker could potentially choose to attack one of these flights. The FAMS can assign $m < n$ air marshals that may be assigned to protect these flights. Since the number of possible schedules exponentially increases with the number of flights and resources, DOBSS is no longer applicable to the FAMS domain. Instead, IRIS uses the much faster ASPEN algorithm (Jain et al. 2010a) to generate the schedule for thousands of commercial flights per day. IRIS also use an attribute-based preference elicitation system to determine reward values for the Stackelberg game model.

PROTECT for US Coast Guard

The US Coast Guard's (USCG) mission includes maritime security of the US coasts, ports, and inland waterways; a security domain that faces increased risks due to threats such as terrorism and drug trafficking. Given a particular port and the variety of critical infrastructure that an adversary may attack within the port, USCG conducts patrols to protect this infrastructure; however, while the adversary has the opportunity to observe patrol patterns, limited security resources imply that USCG patrols cannot be at every location 24/7. To assist the USCG in allocating its patrolling resources, the PROTECT (Port Resilience Operational / Tactical Enforcement to Combat Terrorism) model is being designed to enhance maritime security and has been in use at the port of Boston since April 2011 (Figure 1). Similar to previous applications ARMOR and IRIS, PROTECT uses an attacker-defender Stackelberg game framework, with USCG as the defender against terrorist adversaries that conduct surveillance before potentially launching an attack.

While PROTECT builds on previous work, it offers some key innovations. First, this system is a departure from the assumption of perfect adversary rationality noted in previous work, relying instead on a quantal response (QR) model (McKelvey and Palfrey 1995) of the adversary's behavior. Second, to improve PROTECT's efficiency, a compact representation of the defender's strategy space is used by exploiting equivalence and dominance. Finally, the evaluation of PROTECT for the first time provides real-world

data: (i) comparison of human-generated vs PROTECT security schedules, and (ii) results from an Adversarial Perspective Team's (human mock attackers) analysis. The PROTECT model is now being extended to the port of New York and it may potentially be extended to other ports in the US.

GUARDS for US Transportation Security Agency

The United States Transportation Security Administration (TSA) is tasked with protecting the nation's over 400 airports which services approximately 28,000 commercial flights and up to approximately 87,000 total flights per day. To protect this large transportation network, the TSA employs approximately 48,000 Transportation Security Officers, who are responsible for implementing security activities at each individual airport. To aid the TSA in scheduling resources to protect airports, a new application called GUARDS (Game-theoretic Unpredictable and Randomly Deployed Security) has been developed. While GUARDS also utilizes Stackelberg games as ARMOR and IRIS, GUARDS faces three key challenges (Pita et al. 2011): 1) reasoning about hundreds of heterogeneous security activities; 2) reasoning over diverse potential threats; and 3) developing a system designed for hundreds of end-users. To address those challenges, GUARDS created a new game-theoretic framework that allows for heterogeneous defender activities and compact modeling of a large number of threats and developed an efficient solution technique based on general-purpose Stackelberg game solvers. GUARDS is currently under evaluation and testing for scheduling practices at an undisclosed airport. If successful, the TSA intends to incorporate the system into their unpredictable scheduling practices nationwide.

TRUSTS for Urban Security in Transit Systems

In some urban transit systems, including the Los Angeles Metro Rail system, passengers are legally required to purchase tickets before entering but are not physically forced to do so (Figure 2). Instead, patrol units move about through the transit system, inspecting tickets of passengers, who face fines for fare evasion. This setting yields the problem of computing optimal patrol strategies, to deter fare evasion and hence maximize revenue. The TRUSTS system (Tactical Randomization for Urban Security in Transit Systems) models the patrolling problem as a leader-follower Stackelberg game (Jiang et al. 2012). Urban transit systems, however, present unique computational challenges since there are exponentially many possible patrol strategies, each subject to both the spatial and temporal constraints of travel within the transit network under consideration. To overcome this challenge, TRUSTS uses a compact representation which captures the spatial as well as temporal structure of the domain. The system will be evaluated using real-world ridership data from the Los Angeles Metro Rail system.

Future Applications

Beyond the deployed and emerging applications above are a number of different application areas. One of those is protecting forests (Johnson et al. 2012), where we must protect

(a) Los Angeles Metro   (b) Barrier-free entrance to transit system

Figure 2: TRUSTS for transit systems

a continuous forest area from extractors. Since the attacker's behavior (e.g., extracting important resources from the forest) could be effected by spatial considerations, it is critical for the defender to incorporate spatial considerations into her enforcement decisions (Albers 2010). Another potential application is police patrols for crime suppression which is a data-intensive domain (Ordonez et al. 2008). Thus it would be promising to use data mining tools on a database of past reported crime and events to identify the locations to be patrolled, the times at which the game changes, and the types of adversaries faced. The idea is to exploit temporal and spatial patterns of crime on the area to be patrolled to determine the priorities on how to use the limited security resources. Even with all of these applications, we have barely scratched the surface of possibilities in terms of potential applications for multiagent researchers for applying game theory for security.

Open Research Issues

While the deployed applications have advanced the state of the art, significant future research remains to be done. In the following, we highlight some key research challenges, including scalability, robustness, human adversary modeling and mixed-initiative optimization. The main point we want to make is that this research does not require access to classified information of any kind. Problems, solution approaches and datasets are well specified in the papers discussed below,

Scalability: The first research challenge is improving the scalability of our algorithms for solving Stackelberg (security) games. The strategy space of both the defender and the attacker in these games may exponentially increase with the number of security activities, attacks, and resources. As we scale up to larger domains, it is critical to develop newer algorithms that scale up significantly beyond the limits of the current state of the art of Bayesian Stackelberg solvers. Driven by the growing complexity of applications, a sequence of algorithms for solving security games have been developed including DOBSS (Paruchuri et al. 2008), ERASER (Jain et al. 2010b), ASPEN (Jain et al. 2010a). However, existing algorithms still cannot scale up to very large scale domains such as scheduling randomized checkpoints in cities. In such graph based security games, the strategy space of the defender grows exponentially with the

number of available resources and the strategy space of the attacker grows exponentially with the size of the road network considered. The latest technique to schedule such checkpoints is based on a "double oracle approach" which does not require the enumeration of the entire strategy space for either of the players (Jain et al. 2011). However, existing algorithms still cannot scale up to large scale domains such as scheduling randomized checkpoints in cities of the size of Mumbai (Figure 3).

Figure 3: The terrorist attacks of 2008 in Mumbai.

Robustness: The second challenge is improving solutions' robustness. Classical game theory solution concepts often make assumptions on the knowledge, rationality, and capability (e.g., perfect recall) of players. Unfortunately, those assumptions could be wrong in real-world scenarios. Therefore, while computing the defender's optimal strategy, algorithms should take into account various uncertainties faced in the domain, including payoff noise (Kiekintveld, Marecki, and Tambe 2011), execution/observation error (Yin et al. 2011), uncertain capability (An et al. 2011c). While there are algorithms for dealing with different types of uncertainties, there is no general algorithm/framework that can deal with different types of uncertainty simultaneously. Furthermore, existing work assumes that the attacker knows (or with a small noise) the defender's strategy and there is no formal framework to model the attacker's belief update process and how it makes tradeoffs in consideration of surveillance cost, which remains an open issue for in future research.

One required research direction with respect to robustness is addressing bounded rationality of human adversaries, which is a fundamental problem that can affect the performance of our game theoretic solutions. Recently, there has been some research on applying ideas (e.g.,

prospect theory (Kahneman and Tvesky 1979), and quantal response (McKelvey and Palfrey 1995)) from social science or behavioral game theory within security game algorithms (Yang et al. 2011; Pita et al. 2010). Previous work usually applies existing frameworks and sets the parameters of these frameworks by experimental tuning or learning. However, in real-world security domains, we may have very limited data, or may only have some limited information on the biases displayed by adversaries. It is thus still a challenging problem to build high fidelity human adversary models that can address human bounded rationality. Furthermore, since real-world human adversaries are sometimes distributed coalitions of socially, culturally and cognitively-biased agents, acting behind a veil of uncertainty, we may need significant interdisciplinary research to build in social, cultural and coalitional biases into our adversary models.

Mixed-Initiative Optimization: Another challenging research problem in security games is mixed-initiative optimization in which human users and software assistants collaborate to make security decisions (An et al. 2011a). There often exist different types of constraints in security applications. For instance, the defender always has resource constraints, e.g., the numbers of available vehicle checkpoints, canine units, or air marshals. In addition, human users may place constraints on the defender's actions to affect the output of the game when they are faced with exceptional circumstances and extra knowledge. For instance, in the ARMOR system there could be forced checkpoints (e.g., when the Governor is flying) and forbidden checkpoints. Existing applications simply compute the optimal solution to meet all the constraints (if possible). Unfortunately, these user defined constraints may lead to poor (or infeasible) solutions due to the users' bounded rationality and insufficient information about how constraints affect the solution quality. Significantly better solution quality can be obtained if some of these constraints can be relaxed. However, there may be infinitely many ways of relaxing constraints and the software assistant may not know which constraints can be relaxed and by how much, as well as the real-world consequences of relaxing some constraints.

Thus, it is promising to adopt a mixed-initiative approach in which human users and software assistants collaborate to make security decisions. However, designing an efficient mixed-initiative optimization approach is not trivial and there are five major challenges. First, the scale of security games and constraints prevent us from using an exhaustive search algorithm to explore all constraint sets. Second, the user's incomplete information regarding the consequences of relaxing constraints requires preference elicitation support. Third, the decision making of shifting control between the user and the software assistant is challenging. Fourth, it is difficult to evaluate the performance of a mixed-initiative approach. Finally, it is a challenging problem to design good user interfaces for the software assistant to explain how constraints affect the solution quality. What remains to be done for the mixed-initiative approach includes sensitivity analysis for understanding how different constraints affect the solution quality, inference/learning for discovering directions of relaxing constraints, search for

finding constraint sets to explore, preference elicitation for finding the human user's preference of different constraint sets, and interface design for explaining the game theoretic solver's performance.

In addition to the above research challenges, there are other on-going challenges such as preference elicitation for acquiring necessary domain knowledge in order to build game models and evaluation of the game theoretic applications (Taylor et al. 2010).

Resources for Starting This Research

Security is recognized as a world-wide grand challenge and game theory is an increasingly important paradigm for reasoning about complex security resource allocation. While the deployed game theoretic applications have provided a promising start, very significant amount of research remains to be done. These are large-scale interdisciplinary research challenges that call upon multiagent researchers to work with researchers in other disciplines, be "on the ground" with domain experts, and examine real-world constraints and challenges that cannot be abstracted away.

There are a number of resources (mostly online) for starting this research. The research papers related to game theory for security have been extensively published at AAMAS conference [2] and the reader can also find some papers from AAAI [3] and IJCAI [4]. Additional resources:

- Key papers describing important algorithms and the deployed systems can also be found from a recently published book –*Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned* (Tambe 2011).

- The details of those deployed systems can also be found at http://teamcore.usc.edu/projects/security/.

- From http://teamcore.usc.edu/projects/security/, the reader can also find a tutorial at UAI'2011 – Game Theory for Security: Lessons learned from deployed applications.

While we have focused on research conducted by our Teamcore group, there are a few other research groups that have started addressing challenges in security games (Basilico, Gatti, and Amigoni 2009; Korzhyk, Conitzer, and Parr 2010; Dickerson et al. 2010; Korzhyk, Conitzer, and Parr 2011b; 2011a; Letchford and Vorobeychik 2011).

References

Albers, H. 2010. Spatial modeling of extraction and enforcement in developing country protected areas. *Resource and Energy Economics* 32(2):165–179.

An, B.; Jain, M.; Tambe, M.; and Kiekintveld, C. 2011a. Mixed-initiative optimization in security games: A preliminary report. In *Proc. of the AAAI Spring Symposium on Help Me Help You: Bridging the Gaps in Human-Agent Collaboration*, 8–11.

An, B.; Pita, J.; Shieh, E.; Tambe, M.; Kiekintveld, C.; and Marecki, J. 2011b. Guards and protect: Next generation applications of security games. *SIGECOM* 10:31–34.

An, B.; Tambe, M.; Ordonez, F.; Shieh, E.; and Kiekintveld, C. 2011c. Refinement of strong stackelberg equilibria in security games. In *Proc. of the 25th Conference on Artificial Intelligence*, 587–593.

Basilico, N.; Gatti, N.; and Amigoni, F. 2009. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *Proc. of The 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 500–503.

Dickerson, J. P.; Simari, G. I.; Subrahmanian, V. S.; and Kraus, S. 2010. A graph-theoretic approach to protect static and moving targets from adversaries. In *Proc. of The 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 299–306.

Jain, M.; Kardes, E.; Kiekintveld, C.; Ordonez, F.; and Tambe, M. 2010a. Security games with arbitrary schedules: A branch and price approach. In *Proc. of The 24th AAAI Conference on Artificial Intelligence*, 792–797.

Jain, M.; Tsai, J.; Pita, J.; Kiekintveld, C.; Rathi, S.; Tambe, M.; and Ordonez, F. 2010b. Software assistants for randomized patrol planning for the lax airport police and the federal air marshal service. *Interfaces* 40:267–290.

Jain, M.; Korzhyk, D.; Vanek, O.; Pechoucek, M.; Conitzer, V.; and Tambe, M. 2011. A double oracle algorithm for zero-sum security games on graphs. In *Proc. of The 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Jiang, A. X.; Yin, Z.; Kietkintveld, C.; Leyton-Brown, K.; Sandholm, T.; and Tambe, M. 2012. Towards optimal patrol strategies for urban security in transit systems. In *Proc. of the AAAI Spring Symposium on Game Theory for Security, Sustainability and Health*.

Johnson, M.; Fang, F.; Yang, R.; Tambe, M.; and Albers, H. 2012. Patrolling to maximize pristine forest area. In *Proc. of the AAAI Spring Symposium on Game Theory for Security, Sustainability and Health*.

Kahneman, D., and Tvesky, A. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47(2):263–291.

Kiekintveld, C.; Marecki, J.; and Tambe, M. 2011. Approximation methods for infinite bayesian stackelberg games: modeling distributional uncertainty. In *Proc. of The 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Korzhyk, D.; Conitzer, V.; and Parr, R. 2010. Complexity of computing optimal stackelberg strategies in security resource allocation games. In *Proc. of The 24th AAAI Conference on Artificial Intelligence*, 805–810.

Korzhyk, D.; Conitzer, V.; and Parr, R. 2011a. Security games with multiple attacker resources. In *Proc. of The International Joint Conference on Artificial Intelligence (IJCAI)*.

---

[2] www.aamas-conference.org.

[3] www.aaai.org/.

[4] ijcai.org/.

Korzhyk, D.; Conitzer, V.; and Parr, R. 2011b. Solving stackelberg games with uncertain observability. In *Proc. of The 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Letchford, J., and Vorobeychik, Y. 2011. Computing randomized security strategies in networked domains. In *Proc. of The AAAI Workshop on Applied Adversarial Reasoning and Risk Modeling (AARM)*.

McKelvey, R. D., and Palfrey, T. R. 1995. Quantal response equilibria for normal form games. *Games and Economic Behavior* 10(1):6–38.

Ordonez, F.; Tambe, M.; Jara, J. F.; Jain, M.; Kiekintveld, C.; and Tsai, J. 2008. *Handbook on Operations Research for Homeland Security*. chapter Deployed security games for patrol planning.

Paruchuri, P.; Pearce, J. P.; Marecki, J.; Tambe, M.; Ordonez, F.; and Kraus, S. 2008. Playing games with security: An efficient exact algorithm for Bayesian Stackelberg games. In *Proc. of The 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 895–902.

Pita, J.; Jain, M.; Western, C.; Portway, C.; Tambe, M.; Ordonez, F.; Kraus, S.; and Parachuri, P. 2008. Deployed ARMOR protection: The application of a game-theoretic model for security at the Los Angeles International Airport. In *Proc. of The 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 125–132.

Pita, J.; Jain, M.; Tambe, M.; Ordóñez, F.; and Kraus, S. 2010. Robust solutions to stackelberg games: Addressing bounded rationality and limited observations in human cognition. *Artificial Intelligence* 174(15):1142–1171.

Pita, J.; Tambe, M.; Kiekintveld, C.; Cullen, S.; and Steigerwald, E. 2011. Guards - game theoretic security allocation on a national scale. In *Proc. of The 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Tambe, M. 2011. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press.

Taylor, M. E.; Kiekintveld, C.; Western, C.; and Tambe, M. 2010. A framework for evaluating deployed security systems: Is there a chink in your armor? *Informatica* 34:129–139.

Tsai, J.; Rathi, S.; Kiekintveld, C.; Ordonez, F.; and Tambe, M. 2009. IRIS: a tool for strategic security allocation in transportation networks. In *Proc. of The 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 37–44.

Yang, R.; Kiekintveld, C.; Ordonez, F.; Tambe, M.; and John, R. 2011. Improving resource allocation strategy against human adversaries in security games. In *IJCAI*.

Yin, Z.; Jain, M.; Tambe, M.; and Ordonez, F. 2011. Risk-averse strategies for security games with execution and observational uncertainty. In *Proc. of The 25th AAAI Conference on Artificial Intelligence (AAAI)*, 758–763.