GUARDS and PROTECT: Next Generation Applications of Security Games

BO AN, JAMES PITA, ERIC SHIEH, MILIND TAMBE
University of Southern California, CA, USA
and
CHRIS KIEKINTVELD
University of Texas, El Paso, TX, USA
and
JANUSZ MARECKI
IBM T.J. Watson Research, NY, USA

We provide an overview of two recent applications of security games. We describe new features and challenges introduced in the new applications.

Categories and Subject Descriptors: I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems

General Terms: Security, Theory, Design, Performance

Additional Key Words and Phrases: Security, Game Theory, Agents, Resource Allocation

1. INTRODUCTION

The last five years have witnessed the successful application of game theory in reasoning about complex security problems [Basilico et al. 2009; Korzhyk et al. 2010; Dickerson et al. 2010; Jakob et al. 2010; Paruchuri et al. 2008; Pita et al. 2009; Pita et al. 2010; Kiekintveld et al. 2009; Jain et al. 2010]. Stackelberg games have been widely used to model patrolling or monitoring problems in security. In a Stackelberg security game, the defender commits to a strategy and the adversary makes its decision with knowledge of the leader's commitment. Two systems applying Stackelberg game models to assist with randomized resource allocation decisions are currently in use by the Los Angeles International Airport (LAX) [Pita et al. 2008] and the Federal Air Marshals Service (FAMS) [Tsai et al. 2009].

Two new applications called GUARDS (Game-theoretic Unpredictable and Randomly Deployed Security) [Pita et al. 2011] and PROTECT (Port Resilience Operational / Tactical Enforcement to Combat Terrorism) are under development for the Transportation Security Administration (TSA) and the United States Coast Guard respectively. Both are based on Stackelberg games. In contrast with previous applications at LAX and FAMS, which focused on one-off tailored applications and one security activity (e.g., canine patrol, checkpoints, or covering flights) per application, both GUARDS and PROTECT face new challenging issues due to the potential large scale deployment. This includes reasoning about hundreds of hetero-

Authors' addresses: {boa, jpita, eshieh, tambe}@usc.edu, cdkiekintveld@utep.edu, marecki@us.ibm.com

geneous security activities, reasoning over diverse potential threats, and developing a system designed for hundreds of end-users. In this article we will highlight several of the main issues that have arisen. We begin with an overview of the new applications and then discuss these issues in turn.

2. AN OVERVIEW OF GUARDS AND PROTECT

The fundamental novelty in the GUARDS system [Pita et al. 2011], compared to previous applications of such game-theoretic approaches, is the potential national scale deployment at over 400 airports. GUARDS is used to randomize a wide variety of TSA security activities focused on infrastructure protection at airports, and is being designed as a general system for use at any airport. GUARDS has been delivered to the TSA and 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.

The PROTECT (Port Resilience Operational / Tactical Enforcement to Combat Terrorism) model for the United States Coast Guard is being designed to enhance maritime security of coasts, ports, and inland waterways, a mission that faces increased risks given threats such as terrorism and drug trafficking. The PROTECT model casts the patrolling problem as a Bayesian Stackelberg game. We take a game theoretic approach in evaluating the scenario of adversaries (i.e terrorists) versus the defenders (Coast Guard) to generate weighted randomized patrols for the Coast Guard. We plan to demonstrate the PROTECT model in the Port of Boston in the spring/summer of 2011. The PROTECT model also has the potential to be deployed at multiple ports in the United States.

3. RESEARCH CHALLENGES

GUARDS and PROTECT introduce many new features and challenges beyond the previous applications at LAX and FAMS [Pita et al. 2009], mainly due to the potential large scale deployment. One immediate research challenge is improving the scalability of our algorithms for solving security games. Agents' strategy space may exponentially increase with the number of security activities, attacks, and resources. Existing algorithms [Jain et al. 2010; Paruchuri et al. 2008] find optimal randomized security schedules to allocate limited security resources to protect targets. 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 [Jain et al. 2011; Jain et al. 2011].

In both the TSA and the coast guard domain, the defender has to reason over heterogeneous security activities for each potential target and an adversary can execute heterogeneous attacks on a target. In addition, the defender may allocate more than one resource to cover a target. To address this challenge it is necessary to create a more expressive model than outlined in security games [Pita et al. 2011]. The more expressive model is able to reason over the numerous areas, security activities, and threats [Pita et al. 2011]. In fact, previous solution techniques [Jain et al. 2010; Kiekintveld et al. 2009] for traditional security games are no longer directly applicable.

Solving security games with uncertainty is also an important but challenging

problem. The first type of uncertainty is the defender's uncertainty regarding the payoff values of the attacker, which can be modeled as Bayesian Stackelberg games. While existing solution algorithms can handle discrete distributions over possible payoff values, it is challenging to represent and reason about many natural forms of uncertainty over the inputs [Kiekintveld et al. 2011]. Another type of uncertainty comes from the attacker's observation of the defender's strategy. A general Stackelberg game assumes perfect observation of the attacker, which may not be true in practice [Pita et al. 2010; Korzhyk et al. 2011]. In many situations, the attacker may act without observation of the defender's strategy, essentially converting the game into a simultaneous-move game model [Yin et al. 2010]. There are also some other sources of uncertainty regarding the attacker's decision making, e.g., uncertainty in the attacker's decision procedure due to its bounded rationality [Pita et al. 2010].

One of the most difficult issues we faced from the perspective of a potential national deployment was in acquiring the appropriate knowledge for the security challenge being considered. Due to the possibility of hundreds of end-users, it is not practical to sit down with each location and tailor the system to their individual needs. This presents a challenge in acquiring the necessary domain knowledge for such a large network of airports/ports to appropriately model their security challenge. Another interesting research issue is mixed-initiative interactions in which human users and software assistants collaborate to make security decisions [An et al. 2011]. An efficient human interaction process may potentially lead to models with higher overall solution quality.

In addition to the above challenges and lessons, there are on-going challenges of evaluation of this research [Taylor et al. 2010]. Overall, it is increasingly clear that there are a mounting number of interesting research challenges in the security games arena, and while the deployed applications have provided a promising start, very significant amount of research remains to be done.

ACKNOWLEDGMENTS

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.

REFERENCES

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

Basilico, N., Gatti, N., and Amigoni, F. 2009. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *Proceedings 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 *Proceedings 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. 2010. Security games with arbitrary schedules: A branch and price approach. In *Proceedings of The 24th AAAI Conference on Artificial Intelligence*. 792–797.

Jain, M., Kiekintveld, C., and Tambe, M. 2011. Quality-bounded solutions for finite bayesian stackelberg games: Scaling up. In *Proceedings of The 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

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 *Proceedings of The 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Jakob, M., Vanek, O., Urban, S., Benda, P., and Pechoucek, M. 2010. Employing agents to improve the security of international maritime transport. In *Proceedings of AAMAS 2010 Workshop on Agents In Traffic and Transportation*.

Kiekintveld, C., Jain, M., Tsai, J., Pita, J., Tambe, M., and Ordonez, F. 2009. Computing optimal randomized resource allocations for massive security games. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 689–696.

Kiekintveld, C., Marecki, J., and Tambe, M. 2011. Approximation methods for infinite bayesian stackelberg games: Modeling distributional uncertainty. In *Proceedings 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 *Proceedings of The 24th AAAI Conference on Artificial Intelligence*. 805–810.

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

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 *Proceedings of The 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 895–902.

Pita, J., Bellamane, H., Jain, M., Kiekintveld, C., Tsai, J., Ordez, F., and Tambe, M. 2009. Security applications: Lessons of real-world deployment. *ACM SIGecom Exchanges 8,* 2.

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., 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 *Proceedings of The 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 125–132.

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

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 *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 37–44.

Yin, Z., Korzhyk, D., Kiekintveld, C., Conitzer, V., , and Tambe, M. 2010. Stackelberg vs. nash in security games: Interchangeability, equivalence, and uniqueness. In *Proceedings of The 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 1139–1146.