

Game Theory for Security: An Important Challenge for Multiagent Systems

Bo An and Milind Tambe

Computer Science Department
University of Southern California
Los Angeles, CA 90089, USA
{boa, tambe}@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.

Keywords: Game Theory, Security, Multiagent Systems.

1 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 [1]; 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 [2]; 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 [3, 4]; PROTECT has been deployed in the port of Boston since April 2011 and is now in use at the port of New York; GUARDS is under evaluation for national deployment by the US Transportation Security Administration (TSA) [5], and TRUSTS is being tested by the Los Angeles Sheriffs Department (LASD) in the LA Metro system to schedule randomized patrols for fare inspection [6]. 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!

2 Stackelberg Games Background

A generic Stackelberg game has two players, a *leader*, and a *follower*. These players need not represent individuals, but could also be groups that cooperate to execute a joint strategy, such as a police force or a terrorist organization. Each player has a set of possible *pure strategies*, or the actions that they can execute. A *mixed strategy* allows a player to play a probability distribution over pure strategies. Payoffs for each player are defined over all possible pure-strategy outcomes for both the players. The payoff functions are extended to mixed strategies by taking the expectation over pure-strategy outcomes. The follower can observe the leader’s strategy, and then act in a way to optimize its own payoffs. Thus, the attacker’s strategy in a Stackelberg game is a best response to the leader’s strategy.

The most common solution concept in game theory is a *Nash equilibrium*, which is a profile of strategies for each player in which no player can gain by unilaterally changing to another strategy [7]. Strong Stackelberg equilibrium is a refinement of Nash equilibrium; it is a form of equilibrium where the leader commits to a strategy first, and the follower provides a best response while breaking ties in favor of the leader.¹ This Strong Stackelberg equilibrium is the solution concept adopted in security applications [7–10].

The Bayesian extension to the Stackelberg game allows for multiple types of players, with each type associated with its own payoff values [11, 12, 10]. For real-world security domains, we assume that there is only one leader type (e.g., only one police force), although there are multiple follower types (e.g. multiple groups of adversaries are trying to infiltrate security). Each follower type is represented by a different payoff matrix. The leader does not know the follower’s type. The goal is to *find the optimal*

¹ The leader can always induce the follower to strictly break ties in favor of the leader by perturbing his strategy by an infinitesimal amount [8].

mixed strategy for the leader to commit to, given that each follower type will know the mixed strategy of the leader when choosing its own strategy.

3 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 [13–15, 5, 3]. 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.² Stackelberg games have been widely used for modeling/reasoning complex security problems and a variety of 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.

3.1 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, police units patrolling the roads to the terminals, patrolling inside the terminals (with canines), and security screening and bag checks for passengers. The application of game-theoretic approach is focused on two of these measures: (1) placing vehicle checkpoints on inbound roads that service the LAX terminals, including both location and timing (2) scheduling patrols for bomb-sniffing canine units at the different LAX terminals. The eight different terminals at LAX have very different characteristics, like physical size, passenger loads, foot traffic or international versus domestic flights. These factors contribute to the differing risk assessments of these eight terminals. 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 [10]. ARMOR has been successfully deployed since August 2007 at LAX.

² Or the attacker may be sufficiently deterred and dissuaded from attacking the protected target.

3.2 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). Simply finding schedules for the marshals that meet all of these constraints is a computational challenge. Our task is made more difficult by the need to find a randomized policy that meets these scheduling constraints, while also accounting for the different values of each flight.

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 [16] 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.

3.3 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 and now is also in use at the port of New York (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.

The goal of PROTECT is to use game theory to assist the USCG in maximizing its effectiveness in the Ports, Waterways, and Coastal Security (PWCS) Mission. PWCS patrols are focused on protecting critical infrastructure; without the resources to provide one hundred percent on scene presence at any, let alone all of the critical infrastructure, optimization of security resource is critical. Towards that end, unpredictability creates



(a) PROTECT is being used in Boston



(b) Extending PROTECT to NY

Fig. 1. USCG boats patrolling the ports of Boston and NY

situations of uncertainty for an enemy and can be enough to deem a target less appealing. The PROTECT system, focused on the PWCS patrols, addresses how the USCG should optimally patrol critical infrastructure in a port to maximize protection, knowing that the adversary may conduct surveillance and then launch an attack. While randomizing patrol patterns is key, PROTECT also addresses the fact that the targets are of unequal value, understanding that the adversary will adapt to whatever patrol patterns USCG conducts. The output of PROTECT is a schedule of patrols which includes when the patrols are to begin, what critical infrastructure to visit for each patrol, and what activities to perform at each critical infrastructure.

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 [17] 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.

3.4 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. While many people are aware of common security activities, such as individual passenger screening, this is just one of many security layers TSA personnel implement to help prevent potential threats [18, 19]. These layers can involve hundreds of heterogeneous security activities executed by limited TSA personnel leading to a complex resource allocation challenge. While activities like passenger screening are performed for every passenger, the TSA cannot possibly run every security activity all the time. Thus, while the resources required for passenger screening are always allocated by the TSA, it must also decide how to appropriately allocate its remaining security officers among the layers of security to protect against a number of potential threats, while facing challenges such as surveillance and an adaptive adversary as mentioned before.

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 [5]: 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.

3.5 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, security personnel are dynamically deployed throughout the transit system, randomly inspecting passenger tickets. This proof-of-payment fare collection method is typically chosen as a more cost-effective alternative to direct fare collection, i.e., when the revenue lost to fare evasion is believed to be less than what it would cost to directly preclude it.

Take the Los Angeles Metro as an example. With approximately 300,000 riders daily, this revenue loss can be significant; the annual cost has been estimated at \$5.6 million [20]. The Los Angeles Sheriffs Department (LASD) deploys uniformed patrols on board trains and at stations for fare-checking (and for other purposes such as crime prevention), in order to discourage fare evasion. With limited resources to devote to patrols, it is impossible to cover all locations at all times. The LASD thus requires some mechanism for choosing times and locations for inspections. Any predictable patterns in such a patrol schedule are likely to be observed and exploited by potential fare-evaders. The LASD's current approach relies on humans for scheduling the patrols. However, human schedulers are poor at generating unpredictable schedules; furthermore such scheduling for LASD is a tremendous cognitive burden on the human schedulers who must take into account all of the scheduling complexities (e.g., train timings, switching time between trains, and schedule lengths).



Fig. 2. TRUSTS for transit systems

The TRUSTS system (Tactical Randomization for Urban Security in Transit Systems) models the patrolling problem as a leader-follower Stackelberg game [21]. The leader (LASD) precommits to a mixed patrol strategy (a probability distribution over all pure strategies), and riders observe this mixed strategy before deciding whether to buy the ticket or not. Both ticket sales and fines issued for fare evasion translate into revenue to the government. Therefore the optimization objective for the leader is to maximize total revenue (total ticket sales plus penalties). 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 LASD is currently testing TRUSTS in the LA Metro system by deploying patrols according to the generated schedules and measuring the revenue recovered.

3.6 Future Applications

Beyond the deployed and emerging applications above are a number of different application areas. One of those is protecting forests [22], where we must protect a continuous forest area from extractors by patrols through the forest that seek to deter such extraction activity. With limited resources for performing such patrols, a patrol strategy will seek to distribute the patrols throughout the forest, in space and time, in order to minimize the resulting amount of extraction that occurs or maximize the degree of forest protection. This problem can be formulated as a Stackelberg game and the focus is computing optimal allocations of patrol density [22].

Another potential application is police patrols for crime suppression which is a data-intensive domain [23]. 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.

The Stackelberg game framework can also be applied to adversarial domains that exhibit ‘contagious’ actions for each player. For example, word-of-mouth advertising / viral marketing has been widely studied by marketers trying to understand why one product or video goes ‘viral’ while others go unnoticed [24]. Counterinsurgency is the contest for the support of the local leaders in an armed conflict and can include a variety of operations such as providing security and giving medical supplies. Just as in word-of-mouth advertising and peacekeeping operations, these efforts carry a social effect beyond the action taken that can cause advantageous ripples through the neighboring population. Moreover, multiple intelligent parties attempt to leverage the same social network to spread their message, necessitating an adversary-aware approach to strategy generation. Game-theoretic approaches can be used to generate resource allocations strategies for such large-scale, real world networks. The interaction can be modeled as a graph with one player attempting to spread influence while the other player attempts to stop the probabilistic propagation of that influence by spreading their own influence. This ‘blocking’ problem models situations faced by governments/peacekeepers combatting the spread of terrorist radicalism and armed conflict with daily/weekly/monthly visits with local leaders to provide support and discuss grievances [25].

Game-theoretic methods are also appropriate for modeling resource allocation in cybersecurity [26] such as packet selection and inspection for detecting potential threats in large computer networks [27]. The problem of attacks on computer systems and corporate computer networks gets more pressing each year as the sophistication of the attacks increases together with the cost of their prevention. A number of intrusion detection and monitoring systems is being developed, e.g., deep packet inspection method that periodically selects a subset of packets in a computer network for analysis. However, there is a cost associated with the deep packet inspection, as it leads to significant delays in the throughput of the network. Thus, the monitoring system works under a constraint of limited selection of a fraction of all packets which can be inspected. The attacking/protecting problem can be formulated as a game between two players: the attacker (or the intruder), and the defender (the detection system) [27]. The intruder wants to gain control over (or to disable) a valuable computer in the network by scanning the network, hacking into a more vulnerable system, and/or gaining access to further devices on the computer network. The actions of the attacker can therefore be seen as sending malicious packets from a controlled computer (termed source) to a single or multiple vulnerable computers (termed targets). The objective of the defender is to prevent the intruder from succeeding by selecting the packets for inspection, identifying the attacker, and subsequently thwarting the attack. However, packet inspections cause unwanted latency and hence the defender has to decide where and how to inspect network traffic in order to maximize the probability of a successful malicious packet detection. The computational challenge is efficiently computing the optimal defending strategies [27].

4 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 [10], ERASER [15], ASPEN [16]. 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 [28]. 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).



Fig. 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 [29], execution/observation error [30], uncertain capability [31]. 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 [32], and quantal response [17]) from social science or behavioral game theory within security game algorithms [33, 34]. 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 [35]. 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.

Multi-objective Optimization: In existing applications such as ARMOR, IRIS and PROTECT, the defender is trying to maximize a single objective. However, there are domains where the defender has to consider multiple objectives simultaneously. For example, the Los Angeles Sheriff's Department (LASD) needs to protect the city's metro system from ticketless travelers, common criminals, and terrorists. From the perspective of LASD, each one of these attacker types provides a unique threat (lost revenue, property theft, and loss of life). Given this diverse set of threats, selecting a security strategy is a significant challenge as no single strategy can minimize the threat for all attacker types. Thus, tradeoffs must be made and protecting more against one threat may increase the vulnerability to another threat. However, it is not clear how LASD should weigh these threats when determining the security strategy to use. One could attempt to establish methods for converting the different threats into a single metric. However, this process can become convoluted when attempting to compare abstract notions such as safety and security with concrete concepts such as ticket revenue.

Multi-objective security games (MOSG) have been proposed to address the challenges of domains with multiple incomparable objectives [36]. In an MOSG, the threats posed by the attacker types are treated as different objective functions which are not aggregated, thus eliminating the need for a probability distribution over attacker types. Unlike Bayesian security games which have a single optimal solution, MOSGs have a set of Pareto optimal (non-dominated) solutions which is referred to as the Pareto frontier. By presenting the Pareto frontier to the end user, they are able to better understand the structure of their problem as well as the tradeoffs between different security strategies. As a result, end users are able to make a more informed decision on which strategy to enact. Existing approaches so far assume that each attacker type has a single objective and there is no uncertainty regarding each attacker type's payoffs. It is challenging to develop algorithms for solving multi-objective security games with multiple attacker objectives and uncertain attacker payoffs.

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 [37].

5 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 conference³ and the reader can also find some papers from AAAI⁴ and IJ-CAI⁵. 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* [38].
- The details of those deployed systems as well as related publications 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 [13, 14, 39–42].

Acknowledgments. This research is supported by MURI grant W911NF-11-1-0332 and ONR grant N00014-08-1-0733.

References

1. Pita, J., Jain, M., Western, C., Portway, C., Tambe, M., Ordonez, F., Kraus, S., Parachuri, P.: 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, pp. 125–132 (2008)
2. Tsai, J., Rathi, S., Kiekintveld, C., Ordonez, F., Tambe, M.: IRIS: a tool for strategic security allocation in transportation networks. In: Proc. of the 8th International Conference on Autonomous Agents and Multiagent Systems, AAMAS, pp. 37–44 (2009)
3. An, B., Pita, J., Shieh, E., Tambe, M., Kiekintveld, C., Marecki, J.: Guards and protect: Next generation applications of security games. SIGECOM 10, 31–34 (2011)
4. Shieh, E., An, B., Yang, R., Tambe, M., Baldwin, C., DiRenzo, J., Maule, B., Meyer, G.: PROTECT: A deployed game theoretic system to protect the ports of the united states. In: Proc. of the 11th International Conference on Autonomous Agents and Multiagent Systems, AAMAS (2012)
5. Pita, J., Tambe, M., Kiekintveld, C., Cullen, S., Steigerwald, E.: Guards - game theoretic security allocation on a national scale. In: Proc. of the 10th International Conference on Autonomous Agents and Multiagent Systems, AAMAS (2011)

³ www.aamas-conference.org.

⁴ www.aaai.org/.

⁵ ijcai.org/.

6. Yin, Z., Jiang, A., Johnson, M., Tambe, M., Kiekintveld, C., Leyton-Brown, K., Sandholm, T., Sullivan, J.: Trusts: Scheduling randomized patrols for fare inspection in transit systems. In: Proc. of the 24th Conference on Innovative Applications of Artificial Intelligence, IAAI (2012)
7. Osborne, M.J., Rubinstein, A.: *A Course in Game Theory*. MIT Press (1994)
8. von Stengel, B., Zamir, S.: Leadership with commitment to mixed strategies. Technical Report LSE-CDAM-2004-01, CDAM Research Report (2004)
9. Conitzer, V., Sandholm, T.: Computing the Optimal Strategy to Commit to. In: Proc. of the ACM Conference on Electronic Commerce, ACM-EC, pp. 82–90 (2006)
10. Paruchuri, P., Pearce, J.P., Marecki, J., Tambe, M., Ordonez, F., Kraus, S.: 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, pp. 895–902 (2008)
11. Harsanyi, J., Selten, R.: A Generalized Nash Solution for Two-Person Bargaining Games with Incomplete Information 18, 80–106 (1972)
12. Paruchuri, P., Pearce, J.P., Tambe, M., Ordonez, F., Kraus, S.: An efficient heuristic approach for security against multiple adversaries. In: Proc. of the 6th International Conference on Autonomous Agents and Multiagent Systems, AAMAS, pp. 311–318 (2007)
13. Basilico, N., Gatti, N., Amigoni, F.: 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, pp. 500–503 (2009)
14. Korzhyk, D., Conitzer, V., Parr, R.: Complexity of computing optimal stackelberg strategies in security resource allocation games. In: Proc. of the 24th AAAI Conference on Artificial Intelligence, pp. 805–810 (2010)
15. Jain, M., Tsai, J., Pita, J., Kiekintveld, C., Rathi, S., Tambe, M., Ordonez, F.: Software assistants for randomized patrol planning for the lax airport police and the federal air marshal service. *Interfaces* 40, 267–290 (2010)
16. Jain, M., Kardes, E., Kiekintveld, C., Ordonez, F., Tambe, M.: Security games with arbitrary schedules: A branch and price approach. In: Proc. of the 24th AAAI Conference on Artificial Intelligence, pp. 792–797 (2010)
17. McKelvey, R.D., Palfrey, T.R.: Quantal response equilibria for normal form games. *Games and Economic Behavior* 10(1), 6–38 (1995)
18. TSA: *Layers of Security: What We Do* (2011)
19. TSA: *Transportation Security Administration – U.S. Department of Homeland Security* (2011)
20. Hamilton, B.A.: *Faregating analysis*. Report commissioned by the la metro (2007)
21. Jiang, A.X., Yin, Z., Kiekintveld, C., Leyton-Brown, K., Sandholm, T., Tambe, M.: 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 (2012)
22. Johnson, M., Fang, F., Yang, R., Tambe, M., Albers, H.: Patrolling to maximize pristine forest area. In: Proc. of the AAAI Spring Symposium on Game Theory for Security, Sustainability and Health (2012)
23. Ordonez, F., Tambe, M., Jara, J.F., Jain, M., Kiekintveld, C., Tsai, J.: Deployed security games for patrol planning. In: *Handbook on Operations Research for Homeland Security* (2008)
24. Trusov, M., Bucklin, R.E., Pauwels, K.: Effects of word-of-mouth versus traditional marketing: Findings from an internet social networking site. *Journal of Marketing* 73 (2009)
25. Howard, N.J.: Finding optimal strategies for influencing social networks in two player games. Master’s thesis, MIT, Sloan School of Management (2011)
26. Alpcan, T.: *Network Security: A Decision and Game-Theoretic Approach*. Cambridge University Press (2010)

27. Vanek, O., Yin, Z., Jain, M., Bosansky, B., Tambe, M., Pechoucek, M.: Game-theoretic resource allocation for malicious packet detection in computer networks. In: Proc. of the 11th International Conference on Autonomous Agents and Multiagent Systems, AAMAS (2012)
28. Jain, M., Korzhyk, D., Vanek, O., Pechoucek, M., Conitzer, V., Tambe, M.: 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 (2011)
29. Kiekintveld, C., Marecki, J., Tambe, M.: Approximation methods for infinite bayesian stackelberg games: modeling distributional uncertainty. In: Proc. of the 10th International Conference on Autonomous Agents and Multiagent Systems, AAMAS (2011)
30. Yin, Z., Jain, M., Tambe, M., Ordóñez, F.: Risk-averse strategies for security games with execution and observational uncertainty. In: Proc. of the 25th AAAI Conference on Artificial Intelligence, AAAI, pp. 758–763 (2011)
31. An, B., Tambe, M., Ordóñez, F., Shieh, E., Kiekintveld, C.: Refinement of strong stackelberg equilibria in security games. In: Proc. of the 25th Conference on Artificial Intelligence, pp. 587–593 (2011)
32. Kahneman, D., Tversky, A.: Prospect theory: An analysis of decision under risk. *Econometrica* 47(2), 263–291 (1979)
33. Yang, R., Kiekintveld, C., Ordóñez, F., Tambe, M., John, R.: Improving resource allocation strategy against human adversaries in security games. In: IJCAI (2011)
34. Pita, J., Jain, M., Tambe, M., Ordóñez, F., Kraus, S.: Robust solutions to stackelberg games: Addressing bounded rationality and limited observations in human cognition. *Artificial Intelligence* 174(15), 1142–1171 (2010)
35. An, B., Jain, M., Tambe, M., Kiekintveld, C.: 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, pp. 8–11 (2011)
36. Brown, M., An, B., Kiekintveld, C., Ordóñez, F., Tambe, M.: Multi-objective optimization for security games. In: Proc. of the 11th International Conference on Autonomous Agents and Multiagent Systems, AAMAS (2012)
37. Taylor, M.E., Kiekintveld, C., Western, C., Tambe, M.: A framework for evaluating deployed security systems: Is there a chink in your armor? *Informatica* 34, 129–139 (2010)
38. Tambe, M.: *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press (2011)
39. Dickerson, J.P., Simari, G.I., Subrahmanian, V.S., Kraus, S.: 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, pp. 299–306 (2010)
40. Korzhyk, D., Conitzer, V., Parr, R.: Solving stackelberg games with uncertain observability. In: Proc. of the 10th International Conference on Autonomous Agents and Multiagent Systems, AAMAS (2011)
41. Korzhyk, D., Conitzer, V., Parr, R.: Security games with multiple attacker resources. In: Proc. of the International Joint Conference on Artificial Intelligence, IJCAI (2011)
42. Letchford, J., Vorobeychik, Y.: Computing randomized security strategies in networked domains. In: Proc. of the AAAI Workshop on Applied Adversarial Reasoning and Risk Modeling, AARM (2011)