Security Applications: Lessons of Real-World Deployment

JAMES PITA, HARISH BELLAMANE, MANISH JAIN, CHRIS KIEKINTVELD, JASON TSAI, FERNANDO ORDÓÑEZ, MILIND TAMBE
University of Southern California, CA, USA

1. INTRODUCTION

Game theory has played an important role in security decisions. Recent work using Stackelberg games [Fudenberg and Tirole 1991] to model security domains has been particularly influential [Basilico et al. 2009; Kiekintveld et al. 2009; Paruchuri et al. 2008; Pita et al. 2008; Pita et al. 2009]. In a Stackelberg game, a leader (in this case the defender) acts first and commits to a randomized security policy. The follower (attacker) optimizes its reward *considering the strategy chosen by the leader*. These games are well-suited to representing the problem security forces face in allocating limited resources, such as officers, canine units, and checkpoints. In particular, the fact that the attacker is able to observe the policy reflects the way real terrorist organizations plan attacks using extensive surveillance and long planning cycles.

Stackelberg game models are not just theoretical models; they are at the heart of deployed decision-support software now in use by the Los Angeles World Airport (LAWA) police and the United States Federal Air Marshals Service (FAMS). A new application is under development for the Transportation Security Administration (TSA), also using game-theoretic analysis. Moving from theoretical analysis to applying game theory in real applications posed many new challenges, and there remain many open questions to be solved in this exciting area of work. In this article we will highlight several of the main issues that have come up, including (i) developing efficient algorithms to solve large-scale Stackelberg Security Games, (ii) evaluating deployed security systems, (iii) knowledge acquisition from security experts to specify the game models, and (iv) handling mixed-initiative interactions. We begin with an overview of the deployed systems and then discuss these issues in turn.

2. DEPLOYED SYSTEMS OVERVIEW: ARMOR, IRIS, AND GUARDS

The ARMOR (Assistant for Randomized Monitoring Over Routes) system [Pita et al. 2008] has been in use by the LAWA police at Los Angeles International Airport (LAX) since August 2007. This system was designed to assist LAWA police in assigning vehicle checkpoints to inbound roads, and canine units to different airport terminals. ARMOR uses a Bayesian Stackelberg framework to optimally allocate limited resources based on security information provided by LAWA experts. The system uses a mixed-initiative software interface to allow police to adjust the solution returned by the game solver based on any specific constraints or intelligence

Authors' addresses: {jpita, bellaman, manishja, kiekintv, fordon, tambe, jasontts}@usc.edu

for a particular day. Since its deployment, the police have reported an increase in the number of arrests at the airport for offenses such as drug violations and concealed firearms.

The IRIS (Intelligent Randomization In Scheduling) system [Tsai et al. 2009] was designed to address a similar resource allocation problem for the Federal Air Marshals Service (FAMS). FAMS has a limited number of air marshals that may be assigned to protect commercial airline flights. We formulate this as a Stackelberg game, similar to the ARMOR formulation in principle. However, this is a massive scheduling problem involving thousands of personnel, tens of thousands of flights, and complex constraints. New solution methods were necessary to find optimal allocation strategies for this domain in a reasonable amount of time. The input necessary for this problem is also significantly more complicated, which required new user interfaces and model elicitation techniques. After an extensive internal review, IRIS is being used in a pilot deployment since October 2009.

The GUARDS (Game-theoretic Unpredictable and Randomly Deployed Security) system is currently under development for the Transportation Security Administration (TSA). GUARDS will be used to randomize a wide variety of TSA security activities at airports, and is being designed as a general system for use at any airport. The randomization strategy will be based on game-theoretic analysis, as in ARMOR and IRIS. An initial version of the system is currently undergoing testing at the Pittsburgh International (PIT) and Los Angeles International (LAX) airports.

3. LESSONS LEARNED AND OPEN PROBLEMS

Our experience working on these real-world security applications has raised issues, and revealed areas where new research is needed to improve and expand the use of game-theoretic reasoning in practical situations. The first important lesson is that real domains require fast algorithms that can scale to very large and complex problems. Computing a Nash equilibrium is hard (PPAD-complete) in the general case [Daskalakis et al. 2006; Chen and Deng 2006]. Results are somewhat more encouraging for Stackelberg games, which can be solved in polynomial time for some cases [Conitzer and Sandholm 2006]. However, Bayesian Stackelberg games, multiple defense resources, and other complications can lead to large strategy spaces and difficult computation problems [Conitzer and Sandholm 2006; Kiekintveld et al. 2009]. The main trajectory of research in developing algorithms for ARMOR, IRIS, and GUARDS has been to find and exploit structure in the games to avoid combinatorial explosions in representation size and computational requirements [Jain et al. 2008; Paruchuri et al. 2008; Kiekintveld et al. 2009]. While the algorithms developed so far are very effective for known real-world problems, there remains much to be done to develop algorithms that exploit different kinds of structure and can be applied to more general classes of games.

Evaluating deployed security systems such as ARMOR and IRIS is also an important but challenging problem. It is often not possible to run ideal controlled experiments in a deployed setting, and laboratory experiments are somewhat unsatisfying as they abstract away from some features of the problem. We have developed a framework for multi-faceted evaluation and identified some of the key

challenges in evaluating deployed security systems [Taylor et al. 2010], including (i) there are security concerns with making evaluations of security publicly available, leading to difficulties obtaining data and evaluations from security forces (ii) there are ethical difficulties with experimenting in a deployed setting, especially if this could lead to not providing the best possible security at all times, and (iii) there are many external variables that cannot be controlled in any real-world evaluation, such as economic conditions, geopolitics, the number of travelers, and the number of planned attacks. To mitigate these difficulties, we gather evidence from as many sources as possible, including qualitative expert evaluations, statistical data from deployed systems, laboratory experiments with human subjects, and extensive evidence from simulations [Pita et al. 2009; Taylor et al. 2010].

Another key problem in developing an application using game theory is to elicit the knowledge necessary to define a game model. Somehow the data and knowledge that experts have about possible attack scenarios, the available defense resources, and the possible outcomes of different scenarios must be used to generate an explicit Stackelberg game model with well-defined actions and payoffs. In large domains like FAMS, this may require specifying thousands or hundreds of thousands of distinct parameters. One of the important developments was a decomposed preference elicitation scheme that allowed us to automate much of this process and reduce the number of parameters that needed to be specified to a more manageable number [Tsai et al. 2009]. However, the process of developing a game model is still very labor intensive, and further improvements in this process would be of great value.

Finally we examine the lessons learned from providing mixed-initiative interactions, where the final decisions are the result of a multi-stage interaction between the user and the software decision aid. Security settings are often fluid and dynamic, and it may be impractical to capture all of the necessary special cases in a general game model. If the system does not allow the security forces the flexibility to adapt to changing circumstances on the ground, it is not likely to be adopted for everyday use. Therefore, it is necessary to build in mechanisms that allow users to modify the inputs and even in some cases to alter the final suggested schedule output by the game solver. However, altering schedules can affect the proposed outcome of a Stackelberg game and addressing these disturbances is an interesting problem. ARMOR has a mixed-initiative interface that allows such modifications, though in practice this capability is rarely used. Including it was important for building confidence in the system and achieving organization acceptance. Understanding theoretically the impact of such mixed-initiative interactions remains an important open problem.

There is a growing interest in security systems such as ARMOR, IRIS, and GUARDS in new application domains with additional security agencies. Developing systems that meet the demands of these new applications will require applying the existing techniques, but also improving upon these methods and meeting the new challenges that are sure to arise.

REFERENCES

Basilico, N., Gatti, N., and Amigoni, F. 2009. Leader-follower strategies for robotic patrolling in environments with arbitrary topology. In *AAMAS*.

Chen, X. and Deng, X. 2006. Settling the complexity of two-player Nash equilibrium. In *FOCS*. 261–272.

Conitzer, V. and Sandholm, T. 2006. Computing the optimal strategy to commit to. In *ACM EC-06*. 82–90.

Daskalakis, C., Goldberg, P., and Papadimitriou, C. H. 2006. The complexity of computing a Nash equilibrium. In *STOC*. 71–78.

Fudenberg, D. and Tirole, J. 1991. *Game Theory*. MIT Press.

Jain, M., Pita, J., Tambe, M., Ordóñez, F., Paruchuri, P., and Kraus, S. 2008. Bayesian Stackelberg games and their application for security at Los Angeles International Airport. *SIGEcom Exchanges 7,* 1.

Kiekintveld, C., Jain, M., Tsai, J., Pita, J., Tambe, M., and Ordóñez, F. 2009. Computing optimal randomized resource allocations for massive security games. In *AAMAS*.

Paruchuri, P., Marecki, J., Pearce, J., Tambe, M., Ordóñez, F., and Kraus, S. 2008. Playing games for security: An efficient exact algorithm for solving bayesian stackelberg games. In *AAMAS*.

Pita, J., Jain, M., Marecki, J., Ordóñez, F., Portway, C., Tambe, M., Western, C., Paruchuri, P., and Kraus, S. 2008. Deployed ARMOR protection: The application of a game theoretic model for security at the los angeles international airport. In *AAMAS*.

Pita, J., Jain, M., Ordóñez, F., Tambe, M., Kraus, S., and Magori-Cohen, R. 2009. Effective solutions for real-world stackelberg games: When agents must deal with human uncertainties. In *AAMAS*.

Taylor, M., Kiekintveld, C., Western, C., and Tambe, M. 2010. A framework for evaluating deployed security systems: Is there a chink in your ARMOR? In *Informatica*. Vol. 29.

Tsai, J., Rathi, S., Kiekintveld, C., Ordóñez, F., and Tambe, M. 2009. IRIS - a tool for strategic security allocation in transportation networks. In *AAMAS*.