Securing Networks Using Game Theory: Algorithms and Applications

(Extended Abstract)

Manish Jain
University of Southern California
Los Angeles, CA 90089
manish.jain@usc.edu

ABSTRACT

Extensive transportation networks have become the economic backbone of the modern age. Thus, securing these networks against the increasing threat of terrorism is of vital importance. However, protecting critical infrastructure using limited security resources against intelligent adversaries in the presence of the uncertainty and complexities of the real-world is a major challenge. While game-theoretic approaches have been proposed for security domains, traditional methods cannot scale to realistic problem sizes (up to billions of action combinations), even in the absence of uncertainty.

My thesis proposes new models and algorithms that have not only advanced the state of the art in game-theory, but have actually been successfully deployed in the real-world. For instance, IRIS has been in use by the Federal Air Marshal Service for scheduling officers on some international flights since October 2009. My thesis contributes to a very new area that uses insights from large-scale optimization for game-theoretic problems. It represents a successful transition from game-theoretic advancements to real-world applications that are already in use, and it has opened exciting new avenues to greatly expand the reach of game theory.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence

General Terms

Algorithms, Optimization, Experimentation

Keywords

Game Theory, Bayesian Stackelberg Games, Security

1. INTRODUCTION

Protecting critical infrastructure and targets such as airlines and airports, historical landmarks, and power generation facilities is a challenging task for police and security agencies worldwide. The growing threat of international terrorism has exacerbated this challenge in recent years. This work studies the problem of protecting transportation networks for airplanes, trains, and buses which carry millions of people per day to their destinations, making them a prime target for terrorists. For example, in 2001, the 9/11 attack via commercial airliners resulted in $27.2 billion of direct short term costsas well as a loss of 2,974 lives. The 2008 terrorist attacks in Mumbai resulted in 195 lives lost and nearly 300 wounded.

Measures for protecting potential target areas include monitoring entrances or inbound roads, checking inbound traffic and patrols aboard transportation vehicles. Stackelberg games have been used to model the security resource allocation problem [8], however, the scale of the problem in networked domains makes it challenging for existing techniques to be applied. For example, the Federal Air Marshals Service (FAMS) schedule armed officers onboard passenger aircrafts. The enormity of the challenge faced by the FAMS can be revealed by a small example: an instance with 100 flights and 10 officers would have more than a billion possible assignments; in reality, there are an estimated 3,000–4,000 officers and about 30,000 flights. Another example domain is protecting urban road networks. In response to the attacks in 2008, the Mumbai police have started to schedule a limited number of inspection checkpoints on the road network throughout the city. They have to consider millions of combinations of checkpoints along with billions of paths that the attackers could choose. Additionally, uncertainty in the real-world further increases complexity. For example, the police may be facing either a well-funded hard-lined terrorist or criminals from local gangs. These two groups may have entirely different preferences, and the police may not know what type of attacker they would be facing on any given day. Similar problems are faced in other real-world domains as well.

The objective of a Stackelberg solution algorithm is to compute the allocation of limited security resources to security measures that maximize the expected utility of the defender under the presence of domain dependent scheduling constraints when facing an adaptive intelligent attacker. A significant limitation of existing solution methods [1, 8] is that they handle multiple security resources by enumerating all possible combinations of resource assignments. This grows combinatorially in the number of resources and the size of the network, which makes it computationally infeasible to solve real-world problems, since there may be billions of combinations in the real-world. Moreover, existing algorithms do not scale-up in the presence of uncertainty. My work provides newer models and algorithms specifically designed to handle security challenges faced in large networked domains.

2. CONTRIBUTIONS

Many security domains involve allocating multiple resources to cover many potential targets. Such problems are compactly repre-

Cite as: Securing Networks Using Game Theory: Algorithms and Applications (Extended Abstract), Manish Jain, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Tumer, Yolum, Sonenberg and Stone (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. 1345-1346.

sented using *security games* [5], where only payoffs for successful and unsuccessful outcomes for both the defender and the attacker are required. I have developed new models and algorithms to compute optimal defender strategies for these games. In particular, my contributions are as follows: (i) use insights from large-scale optimization to solve massive security games; (ii) identify and exploit domain structure; and (iii) provide a new framework for Bayesian games that is applicable to all Stackelberg solvers.

Use large-scale optimization techniques: Real world problems, like the FAMS and urban road networks, present billions of action choices (pure strategies) to both the defender and the attacker. Such large problem instances cannot even be represented in modern computers, let alone solved using naïve techniques. I have developed algorithms, ASPEN [2] and RUGGED [3], that use strategy generation to provide scale-ups in domains with massive pure strategy spaces. The algorithms start by considering a minimal set of pure strategies for both the players (defender and attacker). 'Useful' strategies are then generated and added to the set, until the optimal solution is obtained. ASPEN uses branch and price, which is a combination of branch and bound and column generation. It is applicable in domains with massive number of defender actions and few (polynomially many) attacker actions, like the FAMS domain where the defender can have billions of possible flight tours but the attacker can only attack the fixed set of flights. Branch and price is not an "out of the box" approach, and ASPEN provides a novel master-slave decomposition to facilitate strategy generation. Additionally, conventional linear relaxation techniques perform poorly in this domain, and ASPEN uses novel branch and bound heuristics that improve its performance by orders of magnitude [2]. Similarly, RUGGED is designed for domains which have a massive number of actions for both players, like in urban road network security, and provides novel best-response formulations that enable strategy generation for both the defender and the attacker.

Exploiting domain structure: The algorithms are designed to exploit the structure of the underlying network. This also enables them to handle specific scheduling constraints presented by the domain. For example, the FAMS need to assign flight tours to every air marshal, where each tour should satisfy the logistical and spatio-temporal domain constraints. This problem of finding the optimal defender strategy in the presence of such scheduling constraints is NP-hard [6]. ASPEN uses a novel decomposition of the problem instance into a master problem and a network flow subproblem, which allows it to efficiently consider all the scheduling constraints while generating new strategies. ASPEN is indeed the first known method for efficiently solving real-world-sized security games with arbitrary schedules, and forms the core of IRIS, the scheduling assistant in use by the FAMS since October 2009. Similarly, RUGGED also uses a network flow formulation to efficiently compute best response paths of the attacker.

Handling uncertainty via Bayesian games: The different preferences of different attacker types are modeled through Bayesian Stackelberg games. Computing the optimal leader strategy in Bayesian Stackelberg game is NP-hard [1], and polynomial time algorithms cannot achieve approximation ratios better than $O(types)$ [7]. I have developed a new technique for solving large Bayesian Stackelberg games that decomposes the entire game into many hierarchically organized *restricted* games, which are used to improve the performance of branch and bound search. The solutions obtained for the restricted games at the 'child' nodes are used to provide: (i) pruning rules, (ii) tighter bounds, and (iii) efficient branching heuristics to solve the bigger game at the 'parent' node faster. Such hierarchical techniques have seen little application towards obtaining optimal solutions in Bayesian games, while Stackelberg settings have not seen any application of such hierarchical decomposition. Additionally, these algorithms are naturally designed for obtaining quality bounded approximations, and provide a further order of magnitude scale-up without any significant loss in quality.

Real-world Results: Game-theoretic approaches for security scheduling have been successfully deployed in the real world, with applications like ARMOR and IRIS in use by the Los Angeles airport police and the FAMS since August 2007 and October 2009 respectively [4]. IRIS uses the ASPEN algorithm for scheduling air marshals on board few international flights; FAMS is indeed working towards increasing the scope of IRIS towards domestic and other sectors. Furthermore, game-theoretic software assistants for other agencies like the Coast Guard and Border Patrol are under development as well.

3. FUTURE WORK

Thus far, my contributions have been in developing models and algorithms for massive security games for transportation networks. In the future, I would like to develop scalable algorithms for more complex security domains: specifically, for domains with multiple levels of security and multiple attackers. Additionally, current models assume that (i) the actions of the defender are executed perfectly, (ii) the attacker observes the defender strategy perfectly, and (iii) the attacker acts rationally. This may not be case in the real-world due to human errors or other unforeseen circumstances. Given that Stackelberg games have already seen real-world deployments in security domains, the requirement of developing robust solution techniques is urgent. Robust strategy generation in Stackelberg games is largely unexplored, and I plan to develop a new framework that can relax the aforementioned assumptions and model uncertainties of the real-world. Finally, I would like to generalize all the insights from this work and build towards a unified scalable robust solution technique.

4. REFERENCES

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

[2] M. Jain, E. Kardes, C. Kiekintveld, F. Ordonez, and M. Tambe. Security games with arbitrary schedules: A branch and price approach. In *AAAI*, 2010.

[3] M. Jain, D. Korzhyk, O. Vanek, V. Conitzer, M. Pechoucek, and M. Tambe. A double oracle algorithm for zero-sum security games on graphs. In *AAMAS*, 2011.

[4] M. Jain, J. Tsai, J. Pita, C. Kiekintveld, S. Rathi, M. Tambe, and F. Ordonez. Software Assistants for Randomized Patrol Planning for the LAX Airport Police and the Federal Air Marshals Service. *Interfaces*, 40:267–290, 2010.

[5] C. Kiekintveld, M. Jain, J. Tsai, J. Pita, M. Tambe, and F. Ordonez. Computing optimal randomized resource allocations for massive security games. In *AAMAS*, pages 689–696, 2009.

[6] D. Korzhyk, V. Conitzer, and R. Parr. Complexity of computing optimal stackelberg strategies in security resource allocation games. In *AAAI*, pages 805–810, 2010.

[7] J. Letchford, V. Conitzer, and K. Munagala. Learning and approximating the optimal strategy to commit to. In *Second International Symposium on Algorithmic Game Theory (SAGT)*, pages 250–262, 2009.

[8] P. Paruchuri, J. P. Pearce, J. Marecki, M. Tambe, F. Ordonez, and S. Kraus. Playing games with security: An efficient exact algorithm for Bayesian Stackelberg games. In *AAMAS-08*, pages 895–902, 2008.