Dynamic Allocation of Security Resources for Protecting Public Events

(Extended Abstract)

Yue Yin[1,2], Bo An[3], Manish Jain[4]
[1]University of Chinese Academy of Sciences, Beijing 100049, China
[2]The Key Laboratory of Intelligent Information Processing, ICT, CAS, Beijing, 100190, China
[3]School of Computer Engineering, Nanyang Technological University, Singapore 639798
[4]Department of Computer Science, Virginia Tech, Blacksburg, VA 24061
[1]melody1235813@gmail.com, [3]boan@ntu.edu.sg,[4]jmanish@cs.vt.edu

ABSTRACT

Large scale public events are attractive targets for terrorist attacks. It is of great significance to intelligently allocate limited security resources to protect such events. In most public events, the impact of an attack at different targets changes over time. For instance, in marathon, the impact of an attack at different locations changes over time as the participants move along the race course. In addition, the police can relocate security resources among potential attacked targets at any time during the event and an attacker may act at any time, thus the strategy spaces of both agents are continuous. Furthermore, a certain kind of public events, e.g., the Olympic Games, is usually held infrequently. Thus the attacker does not get an opportunity to conduct surveillance and respond to a distribution of defender strategies. In this paper, we aim to address the security resource allocation problem in public events domain with time-critical payoff, continuous strategy spaces, and low frequency.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Intelligent agents, Multiagent systems

General Terms

Security, Algorithms

Keywords

Game theory, Security, Optimization

1. INTRODUCTION

Protecting large public events is extremely important, since such events usually provide easy access to a large number of targets for the adversary. In addition, an attack on any target can cause terrible damage. Recently, two bombs exploded near the ending point of the Boston Marathon on April 15, 2013, killing 3 people and injuring an estimated 264 others (See Figure 1).

Figure 1: Boston marathon bombings

However, the security resource allocation problem in such domains is challenging. First, the importance of targets changes over time. For example, the value of targets along a marathon track changes over time with the changing number of participants and spectators at any specific area over the course of the race. Since the number of security resources like policemen or patrolling cars is limited, such resources need to be intelligently allocated to protect the event. Second, the attacker may attack at any time and the defender can relocate resources among targets at any time, thus the strategy space of each agent is continuous and infinite. Furthermore, due to the rarity of such events, the attacker does not get an opportunity to conduct surveillance and respond to a distribution of defender strategies as in [5, 6]. Therefore, a pure defender strategy sampled from the optimal mixed strategy does not necessarily outperform the one-shot optimal pure strategy in terms of ex-post payoff. In this work, we aim to address the security resource allocation problems in such domains with time-critical payoff, continuous strategy spaces and low frequency.

In recent years, game theory has gained attention in security resource allocation problems by researchers [1, 2, 8]. Some systems based on game theoretic approaches have been successfully deployed in the real world, e.g., ARMOR [7] for setting checkpoints on the roadways entering the Los Angeles International Airport; PROTECT [3] for scheduling patrols of the United States Coast Guard. However, a key assumption underlying the technique in these systems is that the payoffs of targets are static over time. While some researchers addressed time-critical domains [4], they arbitrarily discretize the defender strategy space and as such, their solution is not optimal when the continuous defender strategy space is considered. In addition, they compute the mixed strategies for the defender, which is substantially dif-

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

ferent from the solution concept in our domain, where we compute the optimal pure strategy for the defender.

In this paper, we design a game model to minimize the worst-case loss of the defender. In our model, both the defender and the attacker have continuous and infinite strategy spaces. The payoff of an attack for both agents depends on the timing of the attack as well as the number of resources assigned to the target of the attack. In our model, the defender can relocate resources at any time during the proceeding of the public events, while the attacker can attack any target at any time. To address the rarity feature of our problem, we consider pure strategies of both agents and adopt the maximin strategy as the solution concept.

2. SECURITY GAMES IN PUBLIC EVENTS

In a public event, there are n targets (e.g., segments in the marathon scenario) represented by $\mathcal{T} = \{1, \ldots, n\}$ and the defender has m identical security resources (each resource can be a police patrol team in the marathon domain). The value of a target i varies over time t, which can be represented by a continuous function $v_i(t)$. We assume that $v_i(t)$ is piecewise linear and the value functions of targets are common knowledge for the defender and the attacker. We also assume that a public event starts at time 0 and ends at time $t_e > 0$. The defender executes an assignment of resources at time 0 when the event begins. As the event proceeds, the defender may move resources from some targets to other targets. Formally, let $Q^0 = \langle q_i^0 \rangle$ represent the initial assignment of resources where q_i^0 is the number of resources assigned to target i when the event begins. Denote all resource transfers during $[0, t_e]$ as $C = \langle C_k \rangle$ where C_k represents the k^{th} transfer. $C_k = \langle c_{ij}^k : i, j \in \mathcal{T} \rangle$ where c_{ij}^k represents the number of resources transferred from target i to target j in the k^{th} transfer. Let τ_k denote the time when the k^{th} transfer begins. Thus a pure defender strategy is fully represented by a tuple $S = (Q^0, C)$. Let $\mathcal{S}$ be the defender strategy space.

Let $Q^t(S) = \langle q_i^t(S) : i \in \mathcal{T} \rangle$ denote the resource assignment at t, with $q_i^t(S)$ representing the number of resources assigned to target i at time t in S. We represent the time required to transfer resources from target i to target j as d_{ij}. Given the set of transferring time of all target pairs, $D = \langle d_{ij} : i, j \in \mathcal{T} \rangle$, and a defender strategy $S = (Q^0, C)$, $q_i^t(S)$ can be computed as follows.

$$q_i^t(S) = q_i^0 + \sum_{C_k \in C, \tau_k \leq t - d_{ji}, j \in \mathcal{T}} c_{ji}^k - \sum_{C_k \in C, \tau_k \leq t, j \in \mathcal{T}} c_{ij}^k \quad (1)$$

The attacker's pure strategy is represented as (i, t), representing that the attacker attacks target i at time $t, t \in [0, t_e]$. Let $p(r)$ be the probability of a successful attack if the target of the attack is protected by r resources. We set $p(r) = \frac{1}{e^{\lambda r}} (\lambda > 0)$, satisfying $p(r) \in [0, 1]$. λ is a parameter measuring the marginal utility of adding one more security resource. For the attacker, the payoff of attacking target i at time t when the defender plays the strategy S is $U^a(i, t, S) = p(q_i^t(S))v_i(t)$. We assume a zero-sum game to reduce the complexity of the model. Thus the defender's payoff is opposite to the attacker's payoff, i.e., $U^d(i, t, S) = -U^a(i, t, S)$.

We model the problem as a one-shot game due to the rarity of public events and adopt the maximin strategy as solution concept. Namely, the defender chooses a strategy maximizing the worst case defender utility, which indicates that the attacker maximizes his utility under the zero-sum game assumption. We focus on computing optimal pure defender strategy in this work. Let the attacker's response function be $f(S) = \{f_{tg}(S) : S \to i, f_{tm}(S) : S \to t\}$ where $f_{tg}(S)$ is the target attacked and $f_{tm}(S)$ is the time of attack. A pair of strategies $(S, f(S))$ form a maximin equilibrium if they satisfy the following:

$$U^a(f_{tg}(S), f_{tm}(S), S) \geq U^a(i, t, S), \forall i \in \mathcal{T}, t \in [0, t_e],$$
$$U^d(f_{tg}(S), f_{tm}(S), S) \geq U^d(f_{tg}(S'), f_{tm}(S'), S'), \forall S' \in \mathcal{S}.$$

3. CHALLENGES

Though we consider pure strategy equilibrium, the continuous strategy spaces of both agents still make the problem computationally challenging. Specifically, both the number of transfers made by the defender in an event and the timing of each transfer are unknown, which makes the typical mixed integer linear programming approach to solve security games (as in [7]) infeasible in our model. In addition, the time needed to transfer resources between different pairs of targets can be different. When a resource is in transfer, it is not used to protect any target. These facts exacerbate the difficulties of computing the optimal defender strategies.

4. ACKNOWLEDGEMENTS

This work is partially supported under the Singapore Institute of Manufacturing Technology-Nanyang Technological University (SIMTech-NTU) Joint Laboratory and Collaborative research Programme on Complex Systems, and the Center for Computational Intelligence (C2I) at NTU. This work is also supported by Singapore MOE AcRF Tier 1 grant MOE RG33/13, and NSFC grant No. 61202212.

5. REFERENCES

[1] B. An, M. Brown, Y. Vorobeychik, and M. Tambe. Security games with surveillance cost and optimal timing of attack execution. In *AAMAS*, pages 223–230, 2013.

[2] B. An, D. Kempe, C. Kiekintveld, E. Shieh, S. Singh, M. Tambe, and Y. Vorobeychik. Security games with limited surveillance. In *AAAI*, pages 1241–1248, 2012.

[3] B. An, F. Ordónez, M. Tambe, E. Shieh, R. Yang, C. Baldwin, J. DiRenzo III, K. Moretti, B. Maule, and G. Meyer. A deployed quantal response-based patrol planning system for the US Coast Guard. *Interfaces*, 43(5):400–420, 2013.

[4] F. Fang, A. X. Jiang, and M. Tambe. Optimal patrol strategy for protecting moving targets with multiple mobile resources. In *AAMAS*, pages 957–964, 2013.

[5] D. Korzhyk, V. Conitzer, and R. Parr. Solving Stackelberg Games with uncertain observability. In *AAMAS*, pages 1013–1020, 2011.

[6] J. Letchford and Y. Vorobeychik. Optimal interdiction of attack plans. In *AAMAS*, pages 199–206, 2013.

[7] J. Pita, M. Jain, J. Marecki, F. Ordóñez, C. Portway, M. Tambe, C. Western, P. Paruchuri, and S. Kraus. Deployed ARMOR protection: The application of a game theoretic model for security at the Los Angeles International Airport. In *AAMAS*, pages 125–132, 2008.

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