Two-phase Security Games

Extended Abstract

Andrzej Nagórko
University of Warsaw, IDEAS NCBR
Warsaw, Poland
amn@mimuw.edu.pl

Paweł Ciosmak
IDEAS NCBR
Warsaw, Poland
pawel.ciosmak@ideas-ncbr.pl

Tomasz Michalak
University of Warsaw, IDEAS NCBR
Warsaw, Poland
tpm@mimuw.edu.pl

ABSTRACT

A standard model of a security game assumes a one-off assault during which the attacker cannot update their strategy even if new actionable insights are gained in the process. In this paper, we propose a version of a security game that takes into account a possibility of a two-phase attack. Specifically, in the first phase, the attacker makes a preliminary move to gain extra information about this particular instance of the game. Based on this information, the attacker chooses an optimal concluding move. We derive a compact-form mixed-integer linear program that computes an optimal strategy of the defender. Our simulation shows that this strategy mitigates serious losses incurred to the defender by a two-phase attack while still protecting well against less sophisticated attackers.

KEYWORDS

Security games; Two-phase attack; Mixed-Integer formulation

ACM Reference Format:
Andrzej Nagórko, Paweł Ciosmak, and Tomasz Michalak. 2023. Two-phase Security Games: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 INTRODUCTION

In a classic economic model of a Stackelberg game [7], the leader chooses his strategy first, and while doing this, he is observed by the followers, who can adjust their response accordingly. In the last two decades, this model has received significant attention in the context of security applications, where a defender (the leader in the Stackelberg game) distributes limited security resources to guard a set of targets against an attacker (the follower in the Stackelberg game). For instance, Stackelberg games were applied in such domains as infrastructure security (ARMOR [4], IRIS [6], PROTECT [5]), green security (PAWS [8], MIDAS [2]), opportunistic crimes (TRUSTS [9]), as well as cybersecurity [10]. In all these contexts, Stackelberg games are often called *security games*.

The attack in security games is typically modeled as a one-off assault during which the attacker has no chance to update their strategy even if new valuable information is gained in the process. This, however, does not cover certain tactics that can be applied by ever more agile covert organizations. In particular, given the improvements in border control technologies that result in significant quantities of cocaine being seized in Latin America and Europe, drug cartels have to look for more innovative smuggling methods and routes. Unfortunately, according to a report by the European Monitoring Center for Drugs and Drug Addiction [1, p. 4]: "*These groups are innovative and skilled in switching and modifying both trafficking routes and* modi operandi *to circumvent law enforcement activities. They are quick to identify and exploit new opportunities for cocaine trafficking (...) shift transit routes and storage points to capitalize on the presence of ineffective border controls.*" To look for such new routes and access points, in the first phase of an operation, drug cartels can send "small-time" couriers whose key goal is to gain information. In the second phase, given the extra insight, the decision is made on which routes should be chosen for transports of much larger quantities and value. This paper stems from an observation that most of the existing models are vulnerable to such two-phase attacks which may have a significant security repercussions.

Against this background, we propose a security game that takes into account a possibility of a two-phase attack. We show that a strategy computed with our model mitigates serious losses of the defender from a two-phase attack while still protecting well against less sophisticated attackers, comparing to the strategies computed using standard techniques [4]. Moreover, we present time complexity experimental analysis comparing various algorithms for computing defender's strategies in the two-phase security games.

2 MODEL DESCRIPTION

In the Baysian Stackelberg two-phase security game the defender picks his mixed strategy x first. Then, with the knowledge of x, the attacker of type t (encountered with probability p_t) picks his first-phase mixed strategy y. After both the defender and the attacker make their random moves independently according to x and y, the attacker learns his first-phase payoff c. With this information he picks his second-phase mixed strategy z. The outcome of the game for the defender is $r + r'$, where r denotes the first-phase defender's payoff and r' denotes the second-phase one. The outcome for the attacker is $c + c'$, where c' is the second phase payoff. Optimal defender's strategy is found by maximizing the expected payoff $\mathrm{E}(r+r')$ assuming that the attacker is maximizing his payoff $\mathrm{E}(c + c')$. This optimization problem can be formulated as a mixed quadratic linear programming problem (MQLP). This, in turn, can be reformulated as a mixed integer linear programming problem (MILP).

3 EXAMPLE: AIRPORT PROTECTION

In order to present efficiency of our model let us examine a scaled-down version of the two-phase security game in the airport setting. We assume that there are four airport terminals (S_1, S_2, S_3, and S_4) and two patrol units to protect them.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

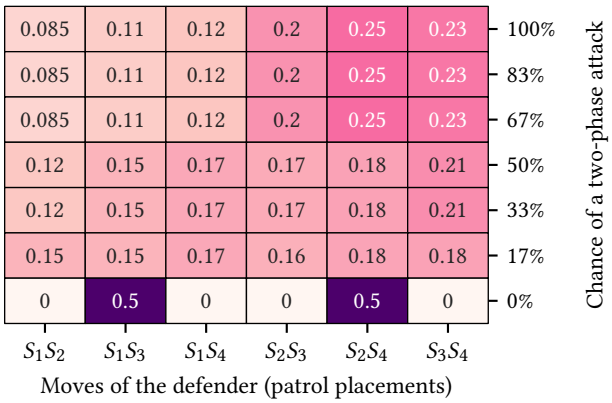

Figure 1: Each row presents an optimal mixed strategy of the defender against a group of attackers with a given expected chance of encountering a two-phase attack. As we can see in the last row, without presence of two-phase attackers the Stackelberg equilibrium heavily over-fits to the random noise in payoff matrices.

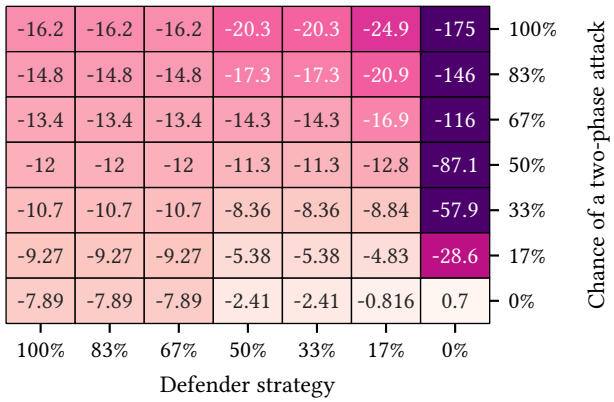

Figure 2: Expected defender payoff when playing a strategy from Figure 1 against a given chance of a two-phase attack. As we can see in the last column, the loss incurred by playing a strategy that ignores the possibility of a two-phase attack is an order of magnitude larger than over-cautious protection against such attacks.

Pure strategies of the defender are placements of the patrols to the terminals $I = \{S_1S_2, S_1S_3, S_1S_4, S_2S_3, S_2S_4, S_3S_4\}$. Two first types of the attacker include low- and high-profile human traffickers (type 1 and 2, respectively), who can choose one terminal as a target of back-off, i.e., $J_1 = J_2 = \{S_1, S_2, S_3, S_4, \emptyset\}$. For a high-profile attackers payoffs are 50, 100, 150 and 200 respectively and for a low-profile attacker the payoffs are five times smaller. The defender payoffs are opposite, with small random noise added uniformly from interval $[-5, 5]$. Third type has the resources and the capabilities of both the low-profile human trafficker and the high-profile one, and he is able to try two terminals in phases. Let $t \in \{0\%, 17\%, 33\%, 50\%, 67\%, 83\%, 100\%\}$ be a chance of encountering a two-phase attacker, $(1-t)\cdot 80\%$ be a probability of encountering a low-profile attacker and $(1-t)\cdot 20\%$ be a likelihood of encountering a high-profile attacker. For $t = 0\%$ this is the standard one-phase model, while $t = 100\%$ describes a pure two-phase attack.

In Figure 1 we show strategies of the defender computed using our model, where the rows are parameterised by the expected probability of a two-phase attack. According to the strategy in the lowest row, representing single-phase game, terminals S_1 and S_2 are never protected simultaneously. Such a situation is typical for Stackelberg equilibria in one-phase games and can be easily exploited by performing a two-phase attack.

Variation of the expected payoff of the defender with respect to the different compositions of attacker groups are show in Figure 2. Notice that the expected payoff of the defender **0.7** against a single-phase attack drops to **-175** when single-phase strategy is pitted against a two-phase attacker. However, with our security model, the expected payoff against coordinated attackers jumps from **-175** to **-16.2** (the defender is still at a disadvantage). Note that this comes at a cost: for the uncoordinated (one-phase) attack, when low- and high-profile attackers act independently, this strategy brings payoff **-7.89** to the defender (a drop from **0.7**).

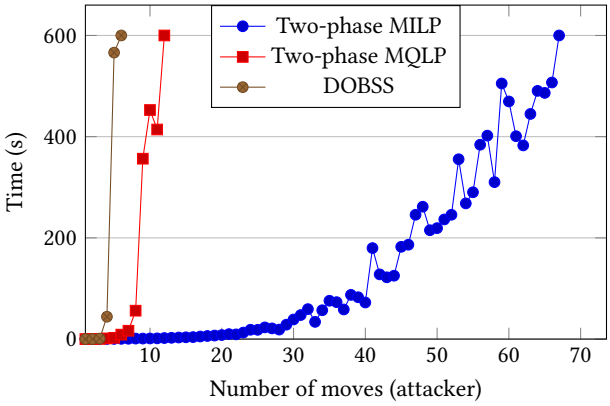

Figure 3: Running time (averaged) for random problems with 3 defender's moves against a number of attacker's moves marked on the x axis. Two-phase attack with 2 attacker types. Lower time is better, computation time limit 600 seconds. Averaged over 20 runs.

4 TIME COMPLEXITY COMPARISON

Finally, let us compare time complexity of MILP with MQLP and DOBSS [3] algorithms. The application of the last one is possible by the reduction of the Bayesian Stackelberg two-phase game to the standard Bayesian Stackelberg game using Harsanyi transformation. However, this reduction results in an exponential explosion of the problem size. In the Figure 3 we present comparison of the growth of running time across three algorithms mentioned above, with respect to the number of moves of the attacker. The computation was performed with SCIP solver on a single core of Intel Xeon 3.60GHz processor.

REFERENCES

[1] European Monitoring Center for Drugs and Drug Addiction. Perspectives on drugs. cocaine trafficking to Europe, 2016.

[2] William Haskell, Debarun Kar, Fei Fang, Milind Tambe, Sam Cheung, and Elizabeth Denicola. Robust protection of fisheries with compass. In *Twenty-Sixth IAAI Conference*, 2014.

[3] Praveen Paruchuri, Jonathan P Pearce, Janusz Marecki, Milind Tambe, Fernando Ordonez, and Sarit Kraus. Playing games for security: An efficient exact algorithm for solving Bayesian Stackelberg games. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 2*, pages 895–902, 2008.

[4] James Pita, Manish Jain, Fernando Ordónez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, and Sarit Kraus. Using game theory for Los Angeles airport security. *AI magazine*, 30(1):43–43, 2009.

[5] Eric Shieh, Bo An, Rong Yang, Milind Tambe, Craig Baldwin, Joseph DiRenzo, Ben Maule, and Garrett Meyer. Protect: A deployed game theoretic system to protect the ports of the United States. In *Proceedings of the 11th international conference on autonomous agents and multiagent systems-volume 1*, pages 13–20. Citeseer, 2012.

[6] Jason Tsai, Shyamsunder Rathi, Christopher Kiekintveld, Fernando Ordonez, and Milind Tambe. Iris-a tool for strategic security allocation in transportation networks. *AAMAS (Industry Track)*, pages 37–44, 2009.

[7] Heinrich Von Stackelberg. *Marktform und gleichgewicht*. J. Springer, 1934.

[8] Rong Yang, Benjamin J Ford, Milind Tambe, and Andrew Lemieux. Adaptive resource allocation for wildlife protection against illegal poachers. In *AAMAS*, pages 453–460, 2014.

[9] Zhengyu Yin, Albert Xin Jiang, Matthew P Johnson, Christopher Kiekintveld, Kevin Leyton-Brown, Tuomas Sandholm, Milind Tambe, and John P Sullivan. Trusts: Scheduling randomized patrols for fare inspection in transit systems. In *Twenty-Fourth IAAI Conference*, 2012.

[10] Yunxiao Zhang and Pasquale Malacaria. Bayesian Stackelberg games for cybersecurity decision support. *Decision Support Systems*, 148:113599, 2021.