Balance Between Scalability and Optimality in Network Security Games

Doctoral Consortium

Kai Wang
Harvard University
kaiwang@g.harvard.edu

ABSTRACT

Network security games (NSGs) are widely used in security related domain to model the interaction between the attacker and the defender. However, due to the complex graph structure of the entire network, finding a Nash equilibrium even when the attacker is fully rational is not well-studied yet. There is no efficient algorithms known with valid guarantees. We identify two major issues of NSGs: i) non-linearity ii) correlation between edges. NSGs with non-linear objective function are usually hard to optimize, while correlated edges might create exponentially many strategies and impact the scalability. In this paper, we analyze the distortion of linear and non-linear formulations of NSGs with fully rational attacker. We provide theoretical bounds on these different formulations, which can quantify the approximation ratio between linear and non-linear assumption. This result can help us understand how much loss will the linearization incur in exchange for the scalability.

KEYWORDS

Game theory for practical applications; non-cooperative games: theory & analysis

ACM Reference Format:
Kai Wang. 2020. Balance Between Scalability and Optimality in Network Security Games. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 3 pages.

1 INTRODUCTION

Many real-world security problems present the challenge of limited budget, including airport protection, wildlife conservation, and web security. Stackelberg security games (SSGs) are commonly used to model the interaction between the attacker and the defender. This game theory analysis can usually help us to plan ahead and strategically allocate the limited resource to protect our targets. Network security games (NSGs) are a generalized version of SSGs, which involves an underlying graph structure. The defender moves first, choosing a set of edges under budget constraint to allocate checkpoints, while the attacker moves second, observing the defender's mixed strategy and then choosing the optimal path to reach one of the target. Since the strategies of both the attacker and the defender are graph dependent, the underlying game theory structure becomes much more sophisticated.

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

The previous works of NSGs with fully rational attacker generally face into two main challenges: i) edge coverage dependency ii) non-linear objective function. While assuming the attacker to be perfectly rational, the equilibrium finding problem can be encoded as a bilevel optimization problem. RANGER [9] and ESVVT algorithm [8] were proposed to solve NSGs. However, both of them require the objective function to be linear and the coverage of each edge to be independent. Double oracle [4, 5] was then proposed to address coverage dependency. They argued that ignoring the dependency can significantly reduce the defender utility. However, this approach does not have a theoretical guarantee and still cannot deal with non-linear objective function, which implicitly rules out the cases where the checkpoint could be imperfect and we re-screen check the same person multiple times. Mc Carthy et al. introduced the imperfect checkpoints and allowed the attacker to be checked multiple times, allowing the objective function to be non-linear. Unfortunately, their approach does not have a theoretical guarantee.

Boundedly rational attacker is relatively less studied in the domain of NSGs. Yang et al. proposed to fit a quantal response model and the historical data by a neural network. Ford et al. proposed to fit a subjective utility quantal response model [7], which allows more flexibility in the adversarial behavior. These approaches generally ignore the underlying graph structure and just rely on the local feature. Wang et al. proposed to adopt graph convolutional networks to further utilize the graph structure. However, the common challenges of boundedly rational behavior is the lack of theoretical guarantee. The performance depends on the sufficiency of the data and the generalizability of the behavioral model.

In this paper, we focus on fully rational attacker. We provide theoretical analysis of the difference between non-linear formulation and linear formulation. This theoretical analysis helps understand how much loss will be incurred if we linearize the non-linear term in NSGs. Our approximation result shows that at most a constant ratio of the optimality will be lost due to linearization, where we can achieve scalability while maintaining a reasonable optimality.

2 DEFINITION

Given a graph $G = (V, E)$, we denote S to be the set of all possible sources and T to be the set of all possible targets. The defender is trying to allocate checkpoints on some edges in E to prevent the attacker from reaching to a target. The resources are limited so the defender can only allocate up to B checkpoints. The attacker, after observing the defender strategy, is trying to avoid being caught and maximize his own expected reward. Each target comes with payoffs $U_{c/u}^{a/d}(t)$ which respectively refer to the uncaught (u) and caught (c) payoffs of the attacker (a) and defender (d). In this paper (so far),

we only consider the zero-sum case. Thus, $U_c^a(t) = -U_c^d(t)$ and $U_u^a(t) = -U_u^d(t)$. We assume all the uncaught payoffs $U_u^a(t)\ \forall t \in T$ for the attacker are non-negative and all the caught penalties $U_c^a(t)$ are non-positive.

3 FULLY RATINOAL ATTACKER

When all the resources are homogeneous, we can use the marginal probability of covering edge $e \in E$ as the defender mixed strategy. We assume the independence across edge coverage x_e, which implies that there is probability x_e that the attacker will get caught when the attacker passes edge e, and this probability is independent of all other edges. The defender also suffers from a budget constraint B, where in expectation $\sum_{e\in E} x_e \le B$.

Therefore, given a defender strategy x, the attacker's payoff of using a path P toward target t can be given by:

$$\prod\nolimits_{e\in P}(1-x_e)U_u^a(t) + (1-\prod\nolimits_{e\in P}(1-x_e))U_c^a(t) \tag{1}$$

Since the attacker is fully rational, he can enumerate all the possible paths and find the best one to use. From the defender's perspective, the defender aims to minimize the attacker's expected payoff by optimally allocating the limited resource. For simplicity, we assume the NSGs to be zero-sum, where the attacker wants to maximize his utility while the defender wants to minimize the attacker's utility. Therefore, NSGs can be written as a minimax problem.

3.1 Non-linear Minimax Problem

The defender's optimization problem can be written as a minimax problem with exponentially many constraints:

$$\min \quad U^a \tag{2}$$

$$\text{s.t.} \quad U^a \ge \prod\nolimits_{e\in P}(1-x_e)U_u^a(t) + (1-\prod\nolimits_{e\in P}(1-x_e))U_c^a(t) \quad \forall t \in T, \forall \text{ path } P \text{ toward } t \tag{3}$$

$$\sum\nolimits_{e\in E} x_e \le B$$

Let us simply denote the minimax value of optimization problem 2 with budget B by $f(B)$. Notice that this formulation is non-linear and non-convex in x. In practice, we can still apply non-convex optimization solver to find a solution.

3.2 Linear Minimax Problems

If we replace the right hand side of inequality (3) in the optimization problem by only the first order term $\sum_{e\in P} x_e$, then the entire minimax optimization problem will underestimate the defender payoff, providing a lower bound formulation.

Lower bound: the new optimization problem can be written as:

$$\min \quad U^a \tag{4}$$

$$\text{s.t.} \quad U^a \ge (1-y_P)U_u^a(t) + y_P U_c^a(t)\ \forall P, t \in T$$

$$y_P = \min(\sum\nolimits_{e\in P} x_e, 1), \quad \sum\nolimits_{e\in E} x_e \le B, \quad 0 \le x_e \le 1$$

Let us denote the optimum of this optimization problem by $g(B)$.

3.3 Approximation Ratio

Given the lower bound $g(B)$, we can analyze the approximation ratio between the real optimum and the lower bound, which eventually gives us the following result:

Theorem 3.1. *The optimum of optimization problem 4 is* $(1-e^{-1})^2$ *approximate to the optimum of the optimization problem 2.*

In practice, many existing works [8, 9] chose to adopt the linear formulation (Equation 4). Other works using the non-linear formulation either require best response oracle [4, 5] or a constraint programming optimizer [6] to find the optimal solution. Theorem 3.1 provides an insight of the balance of scalability and optimality. In exchange of the scalability, we can adopt the linear formulation but need to sacrifice at most constant ratio of the optimality.

4 BOUNDEDLY RATIONAL ATTACKER

When the attacker is boundedly rational, we can use a behavior function $\mathbf{q}(\mathbf{x}, \xi)$ to represent the attacker's boundedly rational behavior, where $\mathbf{x}$ refers to the marginal coverage that the attacker perceives and ξ refers to the context or information revealed to both defender and the attacker, e.g., target values and graph structure. The defender utility under coverage $\mathbf{x}$, context ξ, and attacker behavior $\mathbf{q}$ can be given by:

$$\text{DefU}(\mathbf{x};\mathbf{q}) = \sum\nolimits_{P\in\mathcal{A}} \mathbf{q}_P(\mathbf{x},\xi)\left(U_u^d(P)\prod\nolimits_{e\in P}(1-\mathbf{x}_e) + U_c^d(P)(1-\prod\nolimits_{e\in P}(1-\mathbf{x}_e))\right) \tag{5}$$

However, unlike the fully rational case, the number of attacker pure strategies $\mathcal{A}$ could be exponential or infinite when cycles exist. It is impossible to exactly compute Equation 5.

Interestingly, we found that when the attacker follows Markovian behavior, the attacker's movement can be represented as an absorbing Markov chain, where the attacker only gets absorbed when he reaches a target or gets caught. More specifically, we say the attacker is Markovian if when the attacker is choosing the next edge, he makes decision based on the current location and static information only. In other words, the attacker is memoryless. Under the memoryless condition, we can efficiently compute the defender utility in Equation 5 as discussed in [2, 3, 10]. Then the optimization problem can be solved by any non-convex optimization solver.

It is interesting that how much we will lose when we approximate a memory-dependent behavior with a memoryless behavior. Similar to the fully rational case, we know that we will lose a constant ratio of optimality. We are also interested in how much we will lose by using a memoryless approximation.

5 CONCLUSION

In Section 3, we introduce two common formulation of network security games (Equation 2, 4). Although Equation 2 correctly encodes the non-linear defender utility, it is usually hard to solve, requiring non-convex optimization solver which might also lead to suboptimality. Instead, Equation 4 is linear and easy to compute. The result in Section 3.3 shows that we can approximate the non-linear formulation by the linear formulation with a constant ratio of optimality gap. In Section 4, we present an defender optimization problem when the attacker is boundedly rational. We point out that exponentially many attacker pure strategies commonly exist in network security games with boundedly rational attacker. This can be resolved by approximating the attacker behavior with a memoryless behavior, where the defender utility can be efficiently computed and optimized.

REFERENCES

[1] Benjamin Ford, Thanh Nguyen, Milind Tambe, Nicole Sintov, and Francesco Delle Fave. 2015. Beware the soothsayer: From attack prediction accuracy to predictive reliability in security games. In *GameSec-15*. 35–56.

[2] Alexander Gutfraind, Aric Hagberg, and Feng Pan. 2009. Optimal interdiction of unreactive Markovian evaders. In *CPAIOR-09*. Pittsburgh, 102–116.

[3] Alexander Gutfraind, Aric A Hagberg, David Izraelevitz, and Feng Pan. 2011. Interdiction of a Markovian evader. In *Proc. of INFORMS Computing Society*. Monterey, CA.

[4] Manish Jain, Vincent Conitzer, and Milind Tambe. 2013. Security scheduling for real-world networks. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems, 215–222.

[5] Manish Jain, Dmytro Korzhyk, Ondřej Vaněk, Vincent Conitzer, Michal Pěchouček, and Milind Tambe. 2011. A double oracle algorithm for zero-sum security games on graphs. In *AAMAS-11*. Taipei, 327–334.

[6] Sara Marie Mc Carthy, Milind Tambe, Christopher Kiekintveld, Meredith L Gore, and Alex Killion. 2016. Preventing illegal logging: Simultaneous optimization of resource teams and tactics for security. In *AAAI-16*. New York.

[7] Thanh Hong Nguyen, Rong Yang, Amos Azaria, Sarit Kraus, and Milind Tambe. 2013. Analyzing the Effectiveness of Adversary Modeling in Security Games. In *AAAI-13*. Bellevue, Washington.

[8] Jason Tsai, Zhengyu Yin, Jun-young Kwak, David Kempe, Christopher Kiekintveld, and Milind Tambe. 2010. Strategic Security Placement in Network Domains with Applications to Transit Security. *Quantitative Risk Analysis for Security Applications (QRASA)* (2010), 17.

[9] Jason Tsai, Zhengyu Yin, Jun-young Kwak, David Kempe, Christopher Kiekintveld, and Milind Tambe. 2010. Urban security: Game-theoretic resource allocation in networked domains. In *Twenty-Fourth AAAI Conference on Artificial Intelligence*.

[10] Kai Wang, Andrew Perrault, Aditya Mate, and Milind Tambe. 2020. Scalable Game-Focused Learning of Adversary Models: Data-to-Decisions in Network Security Games. In *AAMAS-20*. Auckland.

[11] Rong Yang, Fei Fang, Albert Xin Jiang, Karthik Rajagopal, Milind Tambe, and Rajiv Maheswaran. 2012. Designing better strategies against human adversaries in network security games. In *AAMAS-12*. Valencia.