General-Sum Cyber Deception Games under Partial Attacker Valuation Information

Extended Abstract

Omkar Thakoor, Milind Tambe, Phebe Vayanos
University of Southern California
Los Angeles, CA
{othakoor,tambe,vayanou}@usc.edu

Haifeng Xu
Harvard University
Cambridge, MA
hxu@seas.harvard.edu

Christopher Kiekintveld
University of Texas at El Paso
El Paso, TX
cdkiekintveld@utep.edu

ABSTRACT

The rapid increase in cybercrime, causing a reported annual economic loss of $600 billion [20], has prompted a critical need for effective cyber defense. Strategic criminals conduct network reconnaissance prior to executing attacks to avoid detection and establish situational awareness via scanning and fingerprinting tools. Cyber deception attempts to foil these reconnaissance efforts; by disguising network and system attributes, among several other techniques. Cyber Deception Games (CDG) is a game-theoretic model for optimizing strategic deception, and can apply to various deception methods. Recently introduced initial model for CDGs assumes zero-sum payoffs, implying directly conflicting attacker motives, and perfect defender knowledge on attacker preferences. These unrealistic assumptions are fundamental limitations of the initial zero-sum model, which we address by proposing a general-sum model that can also handle uncertainty in the defender's knowledge.

KEYWORDS

Game Theory; Cyber Deception; Cyber Security; Uncertainty

ACM Reference Format:
Omkar Thakoor, Milind Tambe, Phebe Vayanos, Haifeng Xu, and Christopher Kiekintveld. 2019. General-Sum Cyber Deception Games under Partial Attacker Valuation Information. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

1 INTRODUCTION

The ubiquity of Internet connectivity has spurred a vast increase in cybercrime. Recent major attacks include data breaches at Equifax [15], Yahoo [14], as well as government agencies like OPM [24]. Rather than attempting "brute force" exploits, which can lead to detection and arrest, adept attackers conduct reconnaissance as the first stage for an effective cyber attack [16, 22]. Scanning tools such as NMap [21], xProbe2 [3], or fingerprinting techniques like sinFP [4], are used to identify vulnerabilities to develop specific plans to infiltrate the network without the risk of detection.

A defensive measure for mitigating the reconnaissance abilities of attackers is using deception and concealment techniques to make it more difficult to gain an accurate understanding of the true network configuration. Additional uncertainty about the network can lead attackers to spend more time in reconnaissance efforts or alter their tactics for infiltration, boost the chances of detecting their activities, and consequently reduce the efficacy of the infiltration strategies attempted. Examples of cyber deception techniques include the use of honeypots or decoys [13], real systems using deceptive defenses [10], obfuscated responses to fingerprinting [5, 27], and software-defined networking to obfuscate network infrastructure [1]. Canary [31] is a deception-based tool in real-world deployment, while CyberVAN [8] is a test-bed for simulating various deception algorithms.

Effective strategizing is particularly vital due to costs and feasibility constraints that must be satisfied. Some ways of masking may be infeasible due to interference with functionality of the network for legitimate users. Deception via counter-fingerprinting techniques such as HoneyD, OSfuscate, IPMorph etc. typically costs performance degradation [27]. Typically, one must also consider costs of developing, deploying, and maintaining deceptive strategies which may include both computational resources and developer time.

Another key challenge is modeling the preferences and capabilities of the attacker. The attacker's motives can greatly vary — they may exactly conflict the defender's, or they could be orthogonal. E.g., the attacker could be economically motivated whereas the defender may prioritize protecting the information on national security. Often, the preferences may be strongly governed by the available exploits. Given such diversely motivated and equipped real-world adversaries, it is often impossible to know precise information about them, prompting the need to account for uncertainty in defender's knowledge when modeling the attacker.

A game-theoretic approach allows to capture the adverserial nature of the attackers. The Cyber Deception Game (CDG) [28] is a game theoretic model of deception via concealment of the real configuration of the network to mitigate attacker reconnaissance. Yet, it is limited to zero-sum games, and thus, also cannot handle situations where the defender does not accurately know attacker's preferences. We propose a general-sum model which allows the players' preferences to be different, and paves the way for the problem arising from the defender's uncertainty about the attacker's payoffs. Another model of this kind is the Two Stage Deception Game model [32] which considers reconnaissance in two different stages, however, this too relies on the perfect defender knowledge.

The AHEAD architecture for active defense [10] provides a realistic architecture for deploying the deceptive strategies in which real hosts attempt to disguise themselves actively. However, it does not consider the strategic question we address about optimizing

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

the use of these capabilities under practical constraints. Game theoretic approaches have been adopted to model other problems in cyber defense [2, 19, 29, 30], including several that consider using game theory to strategically deploy honeypots for cyber deception [11, 12, 25]. However, these models do not consider the possibility of concealing or disguising the configuration of the real network, as our model does. Previous works have also considered uncertainty about the adversary in security games, however, the results are not applicable due to the specific constraints and objectives of CDGs. These include work on modeling uncertainty about human attackers [26], Bayesian [18] and interval-based approaches [17] for modeling uncertainty in basic security games, and regret-based approaches similar to ours but for other types of security games [23] that do not apply to CDGs.

2 MODEL

Various components of the General-sum CDG model are as follows.

Network Configurations. We view the network as a set $\mathcal{K}$ of machines. Each machine has a *true configuration* (TC), which can be thought of as a tuple of several attributes such as [OS Linux, Webserver TomCat 8]. Thus, it is an abstract and exhaustive categorization of a machine from the security perspective. $\mathcal{I}$ denotes the set of TCs present in the network. The *true state of the network* (TSN) is defined by a vector $\boldsymbol{n} = (n_i)_{i \in \mathcal{I}}$ where each n_i is the no. of machines having TC i. Through deception techniques, the defender "masks" each machine with an *observed configuration* (OC); $\mathcal{J}$ denotes the set of all possible OCs.

Deception Strategies. The defender's deception strategy can be encoded with an integer matrix Φ, where Φ_{ij} denotes the number of machines with TC i, that are masked with OC j. The *observed state of the network* (OSN) is, unlike the TSN, a function of the deception strategy, given by $\boldsymbol{m}(\Phi) := (m_j(\Phi))_{j \in \mathcal{J}}$, where $m_j(\Phi) = \sum_i \Phi_{ij}$ is the no. of machines masked by OC j, under strategy Φ.

Strategy feasibility and costs. Achieving deception is often costly and not arbitrarily feasible. Hence, we have *feasibility* constraints, denoted by a (0,1)-matrix Π, where $\Pi_{ij} = 1$ iff TC i can be masked with OC j. Further, we assume that masking a TC i with an OC j has a net cost of c_{ij} incurred by the defender. The defender requires the total cost of masking to not exceed a limit B, called the *budget*. $\mathcal{F}$ denotes the set of strategies that are *feasible* and *affordable* — $\mathcal{F}$ can be described with linear constraints.

Defender and Attacker Valuations. If the attacker procures a machine with TC i, he gets a utility v_i^{a}, his *valuation* of TC i. Collectively, these are represented as a vector $\boldsymbol{v}^{\mathbf{a}}$. Analogously, we define valuations $\boldsymbol{v}^{\mathbf{d}}$ for the defender; a higher valuation v_i^{d} reflects a smaller loss when TC i is compromised.

Game Model. Then, we consider a Stackelberg game where the defender is the leader who knows TSN $\boldsymbol{n}$ and plays a deception strategy Φ. The attacker is the follower, who then chooses a machine to attack. Since only the OC distinguishes the machine from the attacker's perspective, he must choose an OC to attack as his best response, based on their expected utilities (described momentarily) and randomly attack a machine masked by this OC.

We assume that the defender can only play a pure strategy since it is usually not possible to change the network frequently, making the attacker's view of the network static. We assume the attacker perfectly knows this defender strategy Φ with which he can compute his best response. This assumption on the attacker's knowledge is carried from the earlier work on CDG [28], and justified via insider information leakage or other means of surveillance.

Thus, when the defender plays a strategy Φ, her expected utility when OC j is attacked (with $m_j(\Phi) > 0$), is given by

$$u^{\mathrm{d}}(\Phi, j) = \mathbb{E}[v_i^{\mathrm{d}} | \Phi, j] = \sum_{i \in \mathcal{I}_j} \mathbb{P}(i|\Phi, j) v_i^{\mathrm{d}} = \sum_{i \in \mathcal{I}} \frac{\Phi_{ij}}{m_j(\Phi)} v_i^{\mathrm{d}}.$$

The attacker's expected utility is similarly defined, and denoted as $u^{\mathrm{a}}(\Phi, j, \boldsymbol{v}^{\mathrm{a}})$ since it depends on the attacker valuations.

CDG Example: Consider a CDG with 6 machines, 4 TCs and 3 OCs. Let the TSN be $\boldsymbol{n} = (2, 2, 1, 1)$. Let the valuations be $\boldsymbol{v}^{\mathbf{d}} = (8, 2, 7, 11)$ and $\boldsymbol{v}^{\mathbf{a}} = (7, 2, 5, 11)$. Let $\mathcal{J}_1 = \{1\}$, $\mathcal{J}_2 = \{2\}$, $\mathcal{J}_3 = \{1, 3\}$ and $\mathcal{J}_4 = \{2, 3\}$. Let the costs be $c_{31} = 5$, and $c_{ij} = 1$ for all other feasible (i, j) pairs, and let the budget $B = 7$. Thus, machines with TC 1 and 2 have only 1 choice of OC to mask due to feasibility constraint. Masking TC 3 with OC 1 at cost 5 is too expensive, since masking the remaining machines costs at least 3. Thus, due to the budget constraint, TC 3 has OC 3 as the unique choice. Thus,

$$\mathcal{F} = \left\{ \Phi = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}, \ \Phi' = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{bmatrix} \right\}$$

If the defender plays Φ, attacker's best response is to attack OC 1, yielding expected utilities $u^{\mathrm{a}}(\Phi, 1, \boldsymbol{v}^{\mathbf{a}}) = 7$, and $u^{\mathrm{d}}(\Phi, 1) = 8$ for the attacker and the defender, respectively.

3 OPTIMIZATION PROBLEM

Previous works on Stackelberg games for security domains have typically adopted *Strong Stackelberg equilibrium* (SSE) as the solution concept, which requires mixed strategies to be feasible for guaranteed inducibility. Since CDG only allows pure strategies to the defender, adopting SSE is unjustified. Hence, we consider the *robust* assumption that the attacker breaks ties against the defender, i.e., minimizing her utility, which leads to a *Weak Stackelberg Equilibrium* (WSE) [7]. Consequently, the defender obtains a utility $u^{\min}(\Phi, \boldsymbol{v}^{\mathrm{a}})$ as given by the following optimization problem (OP):

$$\min_j \ u^{\mathrm{d}}(\Phi, j) \mid u^{\mathrm{a}}(\Phi, j, \boldsymbol{v}^{\mathrm{a}}) \geq u^{\mathrm{a}}(\Phi, j', \boldsymbol{v}^{\mathrm{a}}) \ \ \forall j' \in \mathcal{J}. \tag{1}$$

Hence, the defender needs to choose Φ to maximize $u^{\min}(\Phi, \boldsymbol{v}^{\mathbf{a}})$, making this a bi-level optimization problem. A WSE can be guaranteed to exist here, since the leader only plays from a finite set of pure strategies. It has been shown that in the zero-sum setting, when the constraints on feasibility and budget are absent, an optimal strategy is simply to mask all the machines by the same OC [28]. However, such a strategy can be shown to be suboptimal with counter-examples in the general-sum setting.

A key domain challenge is that the defender may not accurately know the attacker's valuations for different TCs. These situations can be modeled by extending the formulation above to incorporate notions of robustness, e.g., minimax regret (MMR) [6, 9].

4 ACKNOWLEDGEMENTS

This research was sponsored by the Army Research Office and accomplished under Grant Number W911NF-17-1-0370.

REFERENCES

[1] Stefan Achleitner, Thomas La Porta, Patrick McDaniel, Shridatt Sugrim, Srikanth V Krishnamurthy, and Ritu Chadha. 2016. Cyber deception: Virtual networks to defend insider reconnaissance. In *Proceedings of the 8th ACM CCS international workshop on managing insider security threats*. ACM, 57–68.

[2] Tansu Alpcan and Tamer Başar. 2010. *Network security: A decision and game-theoretic approach*. Cambridge University Press.

[3] O. Arkin and F. Yarochkin. 2003. *A fuzzy approach to remote active operating system fingerprinting*. http://www.syssecurity.com/archive/papers/Xprobe2.pdf

[4] Patrice Auffret. 2010. SinFP, unification of active and passive operating system fingerprinting. *Journal in Computer Virology* 6, 3 (01 Aug 2010), 197–205. https://doi.org/10.1007/s11416-008-0107-z

[5] David Barroso Berrueta. 2003. A practical approach for defeating Nmap OS-Fingerprinting. *Retrieved March* 12 (2003), 2009.

[6] Craig Boutilier, Relu Patrascu, Pascal Poupart, and Dale Schuurmans. 2006. Constraint-based Optimization and Utility Elicitation Using the Minimax Decision Criterion. *Artif. Intell.* 170, 8-9 (June 2006), 686–713. https://doi.org/10.1016/j.artint.2006.02.003

[7] M. Breton, A. Alj, and A. Haurie. 1988. Sequential Stackelberg equilibria in two-person games. *Journal of Optimization Theory and Applications* 59, 1 (01 Oct 1988), 71–97. https://doi.org/10.1007/BF00939867

[8] R. Chadha, T. Bowen, C. J. Chiang, Y. M. Gottlieb, A. Poylisher, A. Sapello, C. Serban, S. Sugrim, G. Walther, L. M. Marvel, E. A. Newcomb, and J. Santos. 2016. CyberVAN: A Cyber Security Virtual Assured Network testbed. In *MILCOM 2016 - 2016 IEEE Military Communications Conference*. https://doi.org/10.1109/MILCOM.2016.7795481

[9] D. P. de Farias and B. Van Roy. 2003. On constraint sampling in the linear programming approach to approximate linear programming. In *42nd IEEE International Conference on Decision and Control (IEEE Cat. No.03CH37475)*, Vol. 3. 2441–2446 Vol.3. https://doi.org/10.1109/CDC.2003.1272986

[10] Fabio De Gaspari, Sushil Jajodia, Luigi V Mancini, and Agostino Panico. 2016. AHEAD: A New Architecture for Active Defense. In *Proceedings of the 2016 ACM Workshop on Automated Decision Making for Active Cyber Defense*. ACM, 11–16.

[11] Karel Durkota, Viliam Lisỳ, Branislav Bošanskỳ, and Christopher Kiekintveld. 2015. Approximate solutions for attack graph games with imperfect information. In *GameSec*. Springer, 228–249.

[12] Karel Durkota, Viliam Lisỳ, Branislav Bosanskỳ, and Christopher Kiekintveld. 2015. Optimal Network Security Hardening Using Attack Graph Games.. In *IJCAI*. 526–532.

[13] KJ Ferguson-Walter, DS LaFon, and TB Shade. 2017. Friend or Faux: Deception for Cyber Defense. *Journal of Information Warfare* 16, 2 (2017), 28–III.

[14] Vindu Goel and Nicole Perlroth. 2016. *Yahoo Says 1 Billion User Accounts Were Hacked*. https://www.nytimes.com/2016/12/14/technology/yahoo-hack.html.

[15] Ines Gutzmer. 2017. *Equifax Announces Cybersecurity Incident Involving Consumer Information*. https://investor.equifax.com/news-and-events/news/2017/09-07-2017-213000628.

[16] Rob Joyce. 2016. Disrupting Nation State Hackers. USENIX Association, San Francisco, CA.

[17] Christopher Kiekintveld, Towhidul Islam, and Vladik Kreinovich. 2013. Security games with interval uncertainty. In *AAMAS*. 231–238.

[18] Christopher Kiekintveld, Janusz Marecki, and Milind Tambe. 2011. Approximation Methods for Infinite Bayesian Stackelberg Games: Modeling Distributional Payoff Uncertainty. In *AAMAS*. 1005–1012. http://dl.acm.org/citation.cfm?id=2034396.2034412

[19] Aron Laszka, Yevgeniy Vorobeychik, and Xenofon D Koutsoukos. 2015. Optimal Personalized Filtering Against Spear-Phishing Attacks.. In *AAAI*. 958–964.

[20] James Lewis. 2018. *Economic Impact of Cybercrime*. https://www.csis.org/analysis/economic-impact-cybercrime.

[21] Gordon Fyodor Lyon. 2009. *Nmap network scanning: The official Nmap project guide to network discovery and security scanning*. Insecure.

[22] Mandiant. 2013. APT1: Exposing One of China's Cyber Espionage Units. (2013).

[23] Thanh H. Nguyen, Amulya Yadav, Bo An, Milind Tambe, and Craig Boutilier. 2014. Regret-based Optimization and Preference Elicitation for Stackelberg Security Games with Uncertainty. In *AAAI*. 756–762. http://dl.acm.org/citation.cfm?id=2893873.2893991

[24] Andrea Peterson. 2015. *OPM says 5.6 million fingerprints stolen in cyberattack, five times as many as previously thought*. https://www.washingtonpost.com/news/the-switch/wp/2015/09/23/opm-now-says-more-than-five-million-fingerprints-compromised-in-breaches.

[25] Radek Pıbil, Viliam Lisỳ, Christopher Kiekintveld, Branislav Bošanskỳ, and Michal Pechoucek. 2012. Game theoretic model of strategic honeypot selection in computer networks. *Decision and Game Theory for Security* 7638 (2012), 201–220.

[26] James Pita, Richard John, Rajiv Maheswaran, Milind Tambe, Rong Yang, and Sarit Kraus. 2012. A robust approach to addressing human adversaries in security games. In *AAMAS*. 1297–1298.

[27] Mohammad Ashiqur Rahman, Mohammad Hossein Manshaei, and Ehab Al-Shaer. 2013. A game-theoretic approach for deceiving Remote Operating System Fingerprinting. *2013 IEEE Conference on Communications and Network Security (CNS)* (2013), 73–81.

[28] Aaron Schlenker, Omkar Thakoor, Haifeng Xu, Fei Fang, Milind Tambe, Long Tran-Thanh, Phebe Vayanos, and Yevgeniy Vorobeychik. 2018. Deceiving Cyber Adversaries: A Game Theoretic Approach. In *AAMAS*. 892–900. http://dl.acm.org/citation.cfm?id=3237383.3237833

[29] Aaron Schlenker, Haifeng Xu, Mina Guirguis, Chris Kiekintveld, Arunesh Sinha, Milind Tambe, Solomon Sonya, Darryl Balderas, and Noah Dunstatter. 2017. Don't Bury your Head in Warnings: A Game-Theoretic Approach for Intelligent Allocation of Cyber-security Alerts. (2017).

[30] Edoardo Serra, Sushil Jajodia, Andrea Pugliese, Antonino Rullo, and VS Subrahmanian. 2015. Pareto-optimal adversarial defense of enterprise systems. *ACM Transactions on Information and System Security (TISSEC)* 17, 3 (2015), 11.

[31] Thinkst. 2015. *Canary*. https://canary.tools/.

[32] Wei Wang and Bo Zeng. 2018. A Two-Stage Deception Game for Network Defense. In *Decision and Game Theory for Security*, Linda Bushnell, Radha Poovendran, and Tamer Başar (Eds.). Springer International Publishing, Cham, 569–582.