Security Games with Insider Threats

Derya Cansever[1]

[1] Army Research Office NC 27703, USA
derya.h.cansever.civ@mail.mil

Abstract. Many cyber-security defense strategies rely on the information asymmetry between the defender and the attacker. Examples of information asymmetry include passwords and network configuration parameters. Using private information, defenders can drastically increase the computation complexity of the attacker and render his/her attacks inefficient. Availability of insider information can alter the equilibrium and favor the attacker. This paper discusses some of the attributes of private information and describes a three-player game with a partially collaborating insider to illustrate its impact.

Keywords: Insider threat, Stackelberg games, cyber security.

1 Introduction

Protection of infrastructure and information technology systems from advanced and ever more sophisticated cyber-attacks is a major concern. These attacks, often called Advanced Persistent Threat (APT), are launched by well-funded entities and are persistent in pursuing their objectives. Moreover, they often act in a stealthy way to avoid being detected to maximize the long-term payoffs. In fact, it is well documented that cyber-attacks can remain undetected for months or even longer. APT attacks cause staggering amounts of costs to nation states as well as to corporations. Advances in technologies can in principle benefit both sides. However, attackers seem to remain one step ahead in many cases.

A main attribute of the conflict between APT, termed the Attacker, and the defending entity, termed the Defender, is the asymmetry in their respective information structures. The defender does not know when and where the attack will occur. Even with advanced network monitoring systems, attacks may occur without the defender's knowledge. The Defender has an advantage in the asymmetry of information on the actual architecture, configuration attributes and parameters of the system to be defended. This advantage can provide definitive advantage to the defender. However, if an entity in the Defender's organization is compromised and discloses some of the private information for personal gain, this potential advantage can turn into significant vulnerability. In fact, it is reported that [1] in 2018, 44% of data breaches are attributable to insiders. We discuss private information and insider threats in Section 2. In Section 3, we describe a three-player game with a compromised insider. Potential open research areas are discussed in Section 4.

2 Private information

Information asymmetry is a foundational aspect of cyber defense. Defenders use private information to make it more difficult for the Attacker to accomplish his/her goal. Examples of information asymmetry include encryption and Moving Target Defense (MTD) [2]. In symmetric encryption, lack of knowledge of the key by the Attacker makes it extremely difficult for the attacker to decrypt the text of interest. Similarly, lack of the knowledge of the pattern, or the algorithm, that governs the change in configuration parameters makes it very difficult for the Attacker to accomplish his/her goals.

The conflict between the Attacker and the Defender can be modeled as a game where each player pursues conflicting interests. If we neglect for a moment the penalties on their respective controls, the conflict can be approximated by a zero-sum game. For example, the Attacker may be trying to maximize the probability of breaching the IT system, while the Defender is attempting to minimize that probability. Let X denote the state of the system, Y_A denote the information of the Attacker, and Y_D denote the information of the Defender, α_A the strategy of the Attacker and l_D denote the strategy of the defender, and P(.) the probability of breach. Assume that P(.) is strictly concave in the actions of the Attacker and the Defender. The Defender is trying to minimize
$E\{P(X, \alpha_A, l_D \mid Y_D\}$, while the attacker is trying to maximize
$E\{E\{P(X, \alpha_A, l_D \mid Y_D\} \mid Y_A\}$. Assuming that the probability space of Y_A is coarser than the one of Y_D, $E\{P(X, \alpha_A, l_D \mid Y_D\}$ is a random variable from the point of view of the Attacker. Let us call it φ. Let us assume that Y_D has finite support. Then, for each realization of the random variable φ, the Attacker faces a different optimization problem. The Attacker will maximize the conditional expectation of φ given Y_A. Let P* denote the actual outcome of this game for a given realization of φ, and P^0 denote the expected outcome of the game relative to the Attacker. Assume that P(.) is a monotonous function of Y_D and it is strictly concave in α_A, l_D for each value of φ. Then, P* will be smaller than P^0 unless the Attacker has access to the actual realization of φ. The defender wants to make the absolute value of the difference between P* and P^0 for each possible realization of φ as large as possible, weighted by its probability of occurrence. One way to accomplish this goal is to maximize the entropy of Y^D. That is, to ensure that Y^D can take values that are wide spread from its mean, and the respective probabilities of such occurrences are not relatively small. Thus, it is in the best interest of the Defender to ensure that the entropy of its private information is maximized. This can be accomplished by choosing a complex and long password, or increasing the rate of the change of configuration parameters in MTD systems. In addition to maximize the difference between the actual and expected values of the payoff function, the Defender will also try to make the solution of this optimization problem as hard as possible to obtain for the Attacker. This can be accomplished by maximizing the Kolmogorov complexity of the optimization problem of the Attacker in computing its solution. But Kolmogorov complexity and entropy are related. In fact, the Kolmogorov complexity of an i.i.d. sequence converges to its entropy [3]. Thus, maximizing the entropy of its private information is advantageous for the defender in several aspects: make it difficult for the Attacker to compute optimal strategy, and also ensure that the expected outcome of

optimal strategy of its adversary is relatively mediocre compared with what it could have been if he/she had access to the Defender's private information. The Defender's advantage can be reduced when an agent that functions in the defending team is compromised and is willing to share parts of the Defender's private information with the Attacker for his/her own benefit. We discuss such a setting in the next Section.

3 Stealthy Attacks with Insider Information

Consider a security game obtain control of IT resources. This game is described in detail in [4]. Figure 1 describes the evolution of the game when there is no Insider in play.

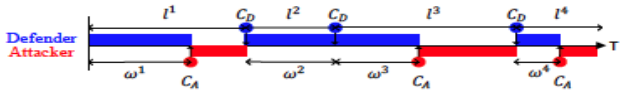

Fig. 1. A two players stealthy game. Blue circles and red circles represent defender's and attacker's actions, respectively. A blue segment denotes that the resource is under protection, and a red segment denotes that the resource is compromised.

This is is a variation of the FlipIt game, [5], where an agent that decides to make a move will have control of the resources. We assume that the defender does not know when the attacker moves, while the Attacker can find out when the Defender made the last move with a random delay ω called the "awareness time". The awareness time represents the period of time it takes for the attacker to find out that the Defender has in fact moved to control the resource. A reasonable strategy for the Defender is to move to control the resource on pints in time that are exponentially distributed to avoid learning by the adversary. The Attacker's corresponding best strategy is to act immediately after ω seconds elapsed. When there is a compromised insider in the defense team, the outcome of the game can change significantly. The insider can help the attacker in reducing the awareness time ω by partially disclosing insider information, and obtain personal benefit by doing so. On the other hand, being a part of the organization, the insider shares its revenue; hence, it may also choose to help the defender against the attacker. In both cases, however, the insider will try to hide its adversary actions from the defender. The insider can inform the Attacker to reduce the awareness time by notifying the Attacker after the Defender's last action (immediately or with a judiciously chosen delay). The sooner it notifies the Attacker, the more it gets paid by the Attacker, but also incurs a higher risk of being detected by the Defender. In this three-person game, the defender first determines and declares its strategy β. After observing β, the insider then decides whether to help the attacker or the defender. In the former case, it makes a "take-it-or-leave-it" offer $\gamma > 0$ to the attacker. In the latter case, it helps the defender by choosing a $\gamma < 0$. Finally, given β and γ, the attacker decides its strategy α. The defender determines the optimal strategy based on the above considerations, which makes the multi-person optimization problem a three level Stackelberg game. Figure 2 below shows the impact of the presence of the insider on the Defender's and Attacker's payoff functions as a function of unit cost C_D

for the Defender. With partial disclosure of private information by the insider, the Attacker and the Defender can face significant changes in their respective payoffs, depending on the values of the game parameters.

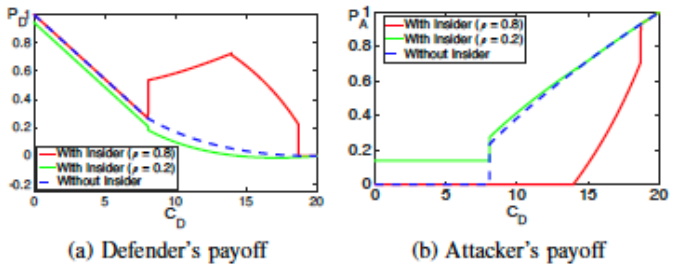

Fig. 2. Change in payoff functions under the presence of an insider.

4 Conclusions and Future Research

Increasing the entropy of the externally observable entities that are related to the state of a system appears to be good strategy for the defenders of cyber systems. Such strategies can be implemented as publicly available security policies, or as actions that amount to intentional signaling to malicious entities that are trying to learn from the actions of the defenders. Design of policies that induce uncertainty for the attackers and that also have robustness properties against actions of insiders could pave fruitful research areas. This is in contrast to the problem of reinforcement learning, where the goal is to learn about the system and to control it efficiently. Analysis of the trade-offs among exploring, mis(signaling) and exploiting functions of control policies for multi-agent systems with asymmetric information could amount to challenging and fertile research topics.

References

1. PWC Homepage, https://www.pwc.co.uk/audit-assurance/assets/pdf/insider-threat-for-google.pdf , last accessed 2020/10/12.
2. R. A. Fink, E. Gunduzhan, B. P. Benjamin, D. Cansever, M. Gralia and P. Dinsmore, "IPsec tunnels vs. identity-only obfuscation techniques for moving target networks," *MILCOM 2012 - 2012 IEEE Military Communications Conference*, Orlando, FL, 2012, pp. 1-6
3. Cover, T., Thomas, J.: Elements of Information Theory. 2nd edn. Wiley, Hoboken, New Jersey (2006).
4. X. Feng, Z. Zheng, D. Cansever, A. Swami and P. Mohapatra, "Stealthy attacks with insider information: A game theoretic model with asymmetric feedback," *MILCOM 2016 - 2016 IEEE Military Communications Conference*, Baltimore, MD, 2016, pp. 277-282
5. M. van Dijk, A. Juels, A. Oprea and R. Rivest, "FlipIt: The Game of "Stealthy Takeover", Journal of Cryptology 26(4) September 2013.