Stealthy Attacks and Observable Defenses: A Game Theoretic Model Under Strict Resource Constraints

Ming Zhang
Department of CSE
The Ohio State University

Zizhan Zheng
Department of ECE
The Ohio State University

Ness B. Shroff
Department of ECE and CSE
The Ohio State University

***Abstract*—Stealthy attacks are a major threat to cyber security. In practice, both attackers and defenders have resource constraints that could limit their capabilities. Hence, to develop robust defense strategies, a promising approach is to utilize game theory to understand the fundamental tradeoffs involved. Previous works in this direction, however, mainly focus on the single-node case without considering strict resource constraints. In this paper, a game-theoretic model for protecting a system of multiple nodes against stealthy attacks is proposed. We consider the practical setting where the frequencies of both the attack and the defense are constrained by limited resources, and an asymmetric feedback structure where the attacker can fully observe the states of nodes while largely hiding its actions from the defender. We characterize best strategies for both the defender and the attacker, and study the Nash Equilibria of the game.**

***Index Terms*—Security games, stealthy attacks**

I. Introduction

Advanced cyber attacks of increased sophistication are a major threat to cyber security [2], [6], [8]. These attacks are often launched by highly motivated entities and are persistent in compromising a system provided that the incentive is high enough. Moreover, they can be highly adaptive, e.g., trading short-term loss for long-term advantage by acting in a *stealthy* way to avoid being detected. In fact, some notorious attacks remained undetected for months or longer [2], [6]. Hence, traditional intrusion detection and prevention techniques in cyber security targeting one-shot and known attack types are insufficient in the face of long-lasting and stealthy attacks.

In this paper, we study an attacker-defender game that explicitly models stealthy attacks. In particular, we consider a system with N nodes (components) of different values, a covert attacker, and an overt defender. Over a continuous time horizon, each player determines when to make a move (i.e., an action) subject to cost for making such a move that varies over nodes. At any time, a node is either protected or compromised, and the payoff to a player depends on the amount of time that the nodes are under its control, and the total cost incurred. Critically, we assume that the attacker is stealthy and its moves are unobservable to the defender. On the other hand, the defender's moves are fully observable to the attacker. Moreover, in practice, both the attacker and the defender have resource constraints that could limit their capabilities, especially for a large system with many nodes. Ignoring such constraints can lead to either over-provisioning or under-provisioning of resources and revenue loss. In this paper, we explicitly model the resource constraints by placing an upper bound on the frequency of moves for each player. To simplify the analysis, we assume in this paper that the proper functioning of one node does not depend on other nodes. This is a reasonable setting and serves as a first-order approximation of the general setting of interdependent nodes. Although game theoretical models have been extensively applied to cyber security [4], [7], [9], [12], prior works mainly focus on attacks of known types.

Our model is inspired by the FlipIt game [13], a two-player non-zero-sum game recently proposed in response to an advanced attack on RSA Data Security [3]. In the FlipIt game, a single critical resource (a node in our model) is considered, and each player obtains control over the resource by "flipping" it subject to a cost. During the play of the game, each player obtains delayed and possibly incomplete feedback on the other player's previous moves. Several variants of the FlipIt game have been considered [10], [11]. However, neither multiple independent nodes nor explicit budget constraints are considered in prior work.

A different type of security games has also been studied in the literature [5], [12], mainly targeting physical infrastructures. Essentially a mixed strategy Stackelberg game is considered, with the defender as the leader and the attacker the follower. A key assumption is that the defender first decides upon a randomized defense policy, and the attacker then observes the randomized policy of the defender but not its realization before taking an action. While this is a useful assumption under certain circumstances, it may not hold when the attacker is highly adaptive, as in the setting allowed in this paper. For instance, since the attacker can somehow observe the defender's previous moves, it may act before the defender changes its policy to get more benefit.

We have made following contributions in this paper.

- We propose a two-player game model with multiple independent nodes, an overt defender, and a stealthy attacker where both players have strict resource constraints.
- We show that a periodic defense strategy is a best-response against a non-adaptive $i.i.d.$ attacking strategy, and vice versa.
- For the above pair of strategies, we fully characterize the set of pure strategy Nash Equilibria of the game, and show that there is at least one equilibrium.

The remainder of this paper is organized as follows. We present our game-theoretic model in Section II, and study best-response strategies of both players in Section III. The analysis of the Nash Equilibria of the game is provided in Section IV. Due to space limitation, we omit most proofs here. The reader is referred to [14] for the missing details.

II. Game Model

In this section, we discuss our two-player game model including its information structure, the actions spaces of of both the defender and the attacker, and their payoffs. Our game model extends the single node model in [11] to multiple nodes and includes a resource constraint on each player.

A. Basic Model

We consider a system with N *independent* nodes and two players, a defender that protects the whole system from attacks and an attacker, and a continuous time horizon T. A player can make a move at any time instance. The attacker incurs a cost of C_i^A per attack towards node i, and takes a period of time w_i to compromise node i. On the other hand, when the defender makes a move to protect node i, which incurs a cost of C_i^D, node i is recovered immediately even if the attack is still in process. A node i has a weight r_i that represents the benefit to the attacker per unit of time when i is compromised. To model the resource constraints of players, we place an upper bound on the average amount of resource that is available to each player at any time (to be formally defined below). As is typical in security games, we assume that the values of r_i, C_i^A, C_i^D, w_i, and the budget constraints are all common knowledge of the game, that is, they are known to both players. Without loss of generality, all nodes are assumed to be protected at time $t = 0$. Table I summarizes the notations used in the paper.

As in [11], we consider an asymmetric feedback structure where the attacker's moves are *stealthy*, while the defenders' moves are *observable*, meaning that the attacker knows the state of each node at any time, while the defender has no idea about whether a node is compromised or not. Therefore, it is reasonable to assume that the attacker will move only after it realizes that a node i has been recovered. Let $\alpha_{i,k}$ denote the time the attacker waits from the last time node i recovered, to the time when the attacker starts its k-th attack against node i. The set $\{\alpha_{i,k}\}$ is the attacker's strategy against the defender. Since the set of nodes are assumed to be independent, $\alpha_{i,k}$ are also independent with respect to i. However, they may be correlated across k as in general, the attacker can employ a time-correlated strategy. In contrast, the defender's strategy is to set the time interval between its $(k-1)$-st move and the k-th move for each node i, denoted as $X_{i,k}$.

In this paper, we focus on *non-adaptive (but possibly randomized) strategies*, that is, neither the attacker nor the defender base its moves on feedback received during the game. Therefore, the values of $\alpha_{i,k}$ and $X_{i,k}$ can be determined by the corresponding player before the game starts. And note that assuming a non-adaptive strategy is not a limitation for the defender since it does not get any feedback about the attacker's moves anyway. Interestingly it turns out not to be a big limitation on the attacker either, because we will show in Section III that, periodic defense is a best-response against any non-adaptive *i.i.d.* attacks (formally defined in Definition III.1) and vice versa. Note that when the defender's strategy is periodic, the attacker can predict the defender's moves before game starts so there is no need to be adaptive.

TABLE I: List of Notations

Symbol	Meaning
r_i	benefit per unit of time by compromising node i
w_i	attacking time for node i
C_i^A	attacker's move cost for node i
C_i^D	defender's move cost for node i
$\alpha_{i,k}$	attacker's waiting time in its k-th move for node i
$X_{i,k}$	time between the (k-1)-st and the k-th defense
B	budget to the defender
M	budget to the attacker
m_i	frequency of recovery action for node i
p_i	probability of immediate attack on node i once it recovers

B. Defender's Problem

Let L_i denote the number of defense moves against node i during T. In general L_i is a random variable. The amount of time when node i is compromised is then $T - \sum_{k=1}^{L_i} \min(\alpha_{i,k} + w_i, X_{i,k})$. Moreover, the cost for defending node i is $L_i C_i^d$. The defender's payoff is then defined as the total loss plus the total defense cost over all the nodes. Given the attacker's strategy $\{\alpha_{i,k}\}$, the defender faces the following optimization problem:

$$\begin{aligned} \max_{\{X_{i,k}, L_i\}} \; & E\left[\sum_{i=1}^{N} \frac{-\left(T - \sum_{k=1}^{L_i} \min(\alpha_{i,k} + w_i, X_{i,k})\right) \cdot r_i}{T}\right] \\ & - \frac{L_i C_i^D}{T} \\ s.t. \; & \sum_{i=1}^{N} \frac{L_i}{T} \le B \text{ w.p.1} \\ & \sum_{k=1}^{L_i} X_{i,k} \le T \text{ w.p.1 } \forall i \end{aligned} \tag{1}$$

The first constraint requires that the average number of nodes that can be protected at any time is upper bounded by a constant B. The second constraint defines the feasible set of $X_{i,k}$. Since T is given, the expectation in the objective function can be moved into the summation in the numerator.

C. Attacker's Problem

We again let L_i denote the number of defense moves against node i in T. The total cost of attacking i is then $(\sum_{k=1}^{L_i} \mathbf{1}_{\alpha_{i,k} < X_{i,k}}) \cdot C_i^A$, where $\mathbf{1}_{\alpha_{i,k} < X_{i,k}} = 1$ if $\alpha_{i,k} < X_{i,k}$ and $\mathbf{1}_{\alpha_{i,k} < X_{i,k}} = 0$ otherwise. It is important to note that when $\alpha_{i,k} \ge X_{i,k}$, the attacker actually gives up its k-th attack against node i (this is possible as the attacker can observe when the defender moves). Given the defender's strategy, the

attacker's problem can be formulated as follows, where M is an upper bound on the average number of nodes that the attacker can attack at any time instance.

$$\begin{aligned}\max_{\alpha_{i,k}} \; & E\left[\sum_{i=1}^{N} \frac{(T-\sum_{k=1}^{L_i}\min(\alpha_{i,k}+w_i, X_{i,k}))\cdot r_i}{T}\right] \\ & - E\left[\sum_{i=1}^{N} \frac{(\sum_{k=1}^{L_i}\mathbf{1}_{\alpha_{i,k}<X_{i,k}})\cdot C_i^A}{T}\right] \\ s.t. \; & E\left[\sum_{i=1}^{N}\frac{1}{T}\int_0^T v_i(t)dt\right] \le M\end{aligned} \tag{2}$$

where $v_i(t) = 1$ if the attacker is attacking node i at time t and $v_i(t) = 0$ otherwise.

Note that as in [11], we make the assumption that the attacker has to keep consuming resources when the attack is in progress instead of making an instantaneous move like the defender; hence it has a different form of budget constraint. We further have the following equation:

$$\int_0^T v_i(t)dt = \sum_{k=1}^{L_i}(\min(\alpha_{i,k}+w_i, X_{i,k}) - \min(\alpha_{i,k}, X_{i,k}))$$

III. Best Responses

In this section, we analyze the best-response strategies for the players. Our main result is that when the attacker employs a non-adaptive *i.i.d.* strategy, a periodic strategy is the best response for the defender, and vice versa.

A. Defender's Best Response

We first observe that it suffices to consider deterministic defense strategies when playing against a non-adaptive attacker.

Lemma III.1. *Suppose $X_{i,k}^*$ and L_i^* are the optimal solutions of* (1) *among all deterministic strategies, then they are also optimal among all the strategies including both deterministic and randomized strategies.*

We then show that a periodic defense is sufficient when the attacker employs a non-adaptive *i.i.d.* strategy defined below.

Definition III.1. *An attack strategy is called non-adaptive i.i.d. if it is non-adaptive, and $\alpha_{i,k}$ is independent across i and is i.i.d. across k.*

Theorem III.1. *A periodical strategy is a best response for the defender if the attacker employs a non-adaptive i.i.d. strategy.*

The main idea of the proof is to show that the defender's payoff for each node i is concave with respect to $X_{i,k}$. The optimality then follows from the KKT conditions. Intuitively, the defender tries to equalize the expected benefits the attacker could receive among its every moves in a deterministic way which gives the defender the most stable system to avoid a big loss in a short period of time.

B. Attacker's Best Response

We first analyze the attacker's best response against any deterministic defense strategies.

Lemma III.2. *When the defense strategies are deterministic, the attacker's best response among non-adaptive strategies must satisfy the following condition*

$$\alpha_{i,k}^* = \begin{cases} 0 & w.p.\ p_{i,k} \\ \ge X_{i,k} & w.p.\ 1-p_{i,k} \end{cases} \tag{3}$$

Note that when $\alpha_{i,k}$ takes the simple form given in (3), the optimization problem to the attacker becomes a continuous knapsack problem and the optimal solution can be found by a simple greedy algorithm [1].

Lemma III.2 tells us that the attacker's best response is to either attack a node immediately after it realizes the node's recovery or give up its attack until the defender's next move. Thus, the constraint M actually determines the probability that the attacker will attack immediately. If M is large enough, the attacker will never wait after defender's each move.

By utilizing the above lemma, we can show that a non-adaptive *i.i.d.* attack is sufficient against periodic defense.

Theorem III.2. *Among all non-adaptive attack strategies, a non-adaptive i.i.d. strategy is the best response against a periodic defense strategy.*

C. Simplified Game

According to Theorem III.1 and Theorem III.2, periodic defense and non-adaptive *i.i.d.* attack can form a pair of best-response strategies with respect to each other. Consider such pair of strategies. Let $m_i \triangleq \frac{L_i}{T} = \frac{1}{X_{i,k}}$, and let p_i denote the probability that $\alpha_{i,k} = 0, \forall k$. The defender's payoff then simplifies to $\sum_{i=1}^{N}\left[\left(E[\min(w_i, \frac{1}{m_i})]p_i r_i - C_i^D\right)\cdot m_i - p_i r_i\right]$. We observe that when $m_i \ge \frac{1}{w_i}$, the defender's cost becomes $m_i C_i^D$, which is minimized when $m_i = \frac{1}{w_i}$. Therefore, it suffices to consider $m_i \le \frac{1}{w_i}$. The optimization problems to the defender and the attacker can then be simplified as follows. Defender's problem:

$$\begin{aligned}\max_{m_i} & \sum_{i=1}^{N} -\left[p_i r_i - m_i(r_i w_i p_i - C_i^D)\right] \\ s.t. \; & \sum_{i=1}^{N} m_i \le B \\ & 0 \le m_i \le \frac{1}{w_i}, \forall i\end{aligned} \tag{4}$$

Attacker's problem:

$$\begin{aligned}\max_{p_i} & \sum_{i=1}^{N} p_i\left[r_i - m_i(r_i w_i + C_i^A)\right] \\ s.t. \; & \sum_{i=1}^{N} m_i w_i p_i \le M \\ & 0 \le p_i \le 1, \forall i\end{aligned} \tag{5}$$

IV. Nash Equilibria

In this section, we study the Nash Equilibria of the simplified game discussed in Section III-C, where the defender

employs a periodic strategy, and the attacker employs a non-adaptive $i.i.d.$ strategy. We fully characterize the set of pure strategy Nash Equilibria of the game and show that our game has at least one pure strategy equilibrium.

For a pair of strategies (m, p), the payoff to the defender is $U_d(m,p) = \sum_{i=1}^{N}[-p_i r_i + m_i(p_i r_i w_i - C_i^D)]$, while the payoff to the attacker is $U_a(m,p) = \sum_{i=1}^{N} p_i[r_i - m_i(r_i w_i + C_i^A)]$. A pair of strategies (m^*, p^*) is called a (pure strategy) *Nash Equilibrium (NE)* if for any pair of strategies (m, p), we have $U_d(m^*, p^*) \geq U_d(m, p^*)$ and $U_a(m^*, p^*) \geq U_a(m^*, p)$. In the following, we assume that $C_i^A > 0$ and $C_i^D > 0$. The cases where $C_i^A = 0$ or $C_i^D = 0$ or both exhibit slightly different structures, but can be analyzed using the same approach. Without loss of generality, we assume $r_i > 0$ and $\frac{C_i^D}{r_i w_i} \leq 1$ for all i. Note that if $r_i = 0$, then node i can be safely excluded from the game since there is no benefit to attack i, while if $\frac{C_i^D}{r_i w_i} > 1$, the coefficient of m_i in U_d is always negative and it is optimal not to protect node i.

Let $\mu_i(p) \triangleq p_i r_i w_i - C_i^D$ denote the coefficient of m_i in U_d, and $\rho_i(m) \triangleq \frac{r_i - m_i(r_i w_i + C_i^A)}{m_i w_i}$. Note that for a given p, the defender tends to protect more a component with higher $\mu_i(p)$, while for a given m, the attacker will attack more frequently a component with higher $\rho_i(m)$. When m and p are clear from the context, we simply let μ_i and ρ_i denote $\mu_i(p)$ and $\rho_i(m)$, respectively.

To find the set of NEs of our game, a key observation is that if there is a full allocation of defense budget B to m such that $\rho_i(m)$ is a constant for all i, any full allocation of the attack budget M gives the attacker the same payoff. Among these allocations, if there is further an assignment of p such that $\mu_i(p)$ is a constant for all i, then the defender also has no incentive to deviate from m; hence (m, p) forms an NE. The main challenge, however, is that such an assignment of p does not always exist for the entire set of nodes. Moreover, there are NEs that do not fully utilize the defense or attack budget as we show below. To characterize the set of NEs, we first prove some properties satisfied by any NE of the game. For a given strategy (m, p), we define $\mu^*(p) \triangleq \max_i \mu_i(p), \rho^*(m) \triangleq \min_i \rho_i(m), F(p) \triangleq \{i : \mu_i(p) = \mu^*(p)\}$, and $E(m,p) \triangleq \{i \in F : \rho_i(m) = \rho^*(m)\}$. We omit m and p when they are clear from the context. We then have the following properties.

Lemma IV.1. *If (m, p) is a NE, we have*

1) $\forall i \notin F, m_i = 0, p_i = 1, \rho_i = \infty$;
2) $\forall i \in F \backslash E, m_i \in [0, \frac{r_i}{w_i r_i + C_i^A}], p_i = 1$;
3) $\forall i \in E, m_i \in [0, \frac{r_i}{w_i r_i + C_i^A}], p_i \in [\frac{C_i^D}{r_i w_i}, 1]$.

Lemma IV.2. *If (m, p) forms a NE, then for $i \in E, j \in F \backslash E$ and $k \notin F$, we have $r_i w_i - C_i^D \geq r_j w_j - C_j^D > r_k w_k - C_k^D$.*

According to Lemma IV.2, to find all the equilibria of the game, it suffices to sort all the nodes in a non-increasing order of $r_i w_i - C_i^D$, and consider each F_h consisting of the first h nodes such that $r_h w_h - C_h^D > r_{h+1} w_{h+1} - C_{h+1}^D$, and each subset $E_k \subseteq F_h$ consisting of the first $k \leq h$ nodes in the list. In the following, we assume such an ordering of nodes. Consider a given pair of F and $E \subseteq F$. By Lemma IV.1 and the definitions of F and E, the following conditions are satisfied by any NE with $F(p) = F$ and $E(m,p) = E$.

$$m_i = 0, p_i = 1, \forall i \notin F; \tag{6}$$

$$m_i \in [0, \frac{r_i}{w_i r_i + C_i^A}], p_i = 1, \forall i \in F \backslash E; \tag{7}$$

$$m_i \in [0, \frac{r_i}{w_i r_i + C_i^A}], p_i \in [\frac{C_i^D}{r_i w_i}, 1], \forall i \in E; \tag{8}$$

$$\sum_{i \in F} m_i \leq B, \sum_{i \in F} m_i w_i p_i \leq M; \tag{9}$$

$$\mu_i = \mu^*, \forall i \in F; \mu_i < \mu^*, \forall i \notin F; \tag{10}$$

$$\rho_i = \rho^*, \forall i \in E; \rho_i > \rho^*, \forall i \notin E \tag{11}$$

The following theorem provides a full characterization of the set of NEs of the game.

Theorem IV.1. *Any pair of strategies (m, p) with $F(p) = F$ and $E(m,p) = E$ is a NE iff it satisfies the set of constraints (6) to (11) and one of the following constraints.*

1) $\sum_{i \in F} m_i = B$; $\rho^* = 0$;
2) $\sum_{i \in F} m_i = B$; $\rho^* > 0$; $\sum_{i \in F} m_i w_i p_i = M$;
3) $\sum_{i \in F} m_i = B$; $\rho^* > 0$; $p_i = 1, \forall i \in F$;
4) $\sum_{i \in F} m_i < B$; $\mu^* = 0$; $F = F_N$; $\rho^* = 0$;
5) $\sum_{i \in F} m_i < B$; $\mu^* = 0$; $F = F_N$; $\rho^* > 0$; $\sum_{i \in F} m_i w_i p_i = M$;
6) $\sum_{i \in F} m_i < B$; $\mu^* = 0$; $F = F_N$; $\rho^* > 0$; $p_i = 1, \forall i \in F$;

In the following, NEs that fall into each of the six cases considered above are named as Type 1 - Type 6 NEs, respectively. We next show that our game always has a NE, and may have multiple NEs of different types and different payoffs.

Theorem IV.2. *The attacker-defender game always has a pure strategy Nash Equilibrium, and may have more than one NE of different payoffs to the defender.*

Proof: The proof of the first part is given in [14]. To show the second part, consider the following example with two nodes where $r_1 = r_2 = 1, w_1 = 2, w_2 = 1, C_1^D = 1/5, C_2^D = 4/5, C_1^A = 1, C_2^A = 7/2, B = 1/3$, and $M = 1/5$. Then it is easy to check that $m = (1/6, 1/6)$ and $p = (3/20, 9/10)$ is a Type 2 NE, and $m = (1/3, 0)$ and $p = (p_1, 1)$ with $p_1 \in [1/5, 3/10]$ are all Type 1 NEs, and all these NEs have different payoffs to the defender. ■

V. Conclusion

In this paper, we propose a two-player non-zero-sum game for protecting a system of multiple nodes against a stealthy attacker where the defender's behavior is fully observable, and both players have strict resource constraints. We prove that the periodic defense and a simple non-adaptive $i.i.d.$ attack are a pair of best-response strategies with respect to each other. For this pair of strategies, we further characterize the set of Nash Equilibria of the game, and show that there is always one (and maybe more) equilibrium, for the case when the attack times are deterministic.

REFERENCES

[1] Continuous knapsack problem. http://en.wikipedia.org/wiki/Continuous_knapsack_problem.

[2] ESET and Sucuri Uncover Linux/Cdorked.A: The Most Sophisticated Apache Backdoor. http://www.eset.com/int/about/press/articles/article/eset-and-sucuri-uncover-linuxcdorkeda-apache-webserver-backdoor-the-most-sophisticated-ever-affecting-thousands-of-web-sites/, 2013.

[3] A. Coviello. Open letter to RSA customers, March 17, 2011. http://www.rsa.com/node.aspx?id=3872.

[4] T. Alpcan and T. Başar. *Network Security: A Decision and Game-Theoretic Approach*. Cambridge University Press, 2010.

[5] B. An, M. Brown, Y. Vorobeychik, and M. Tambe. Security Games with Surveillance Cost and Optimal Timing of Attack Execution. In *Proc. of AAMAS*, 2013.

[6] B. Bencsáth, G. Pék, L. Buttyán, and M. Félegyházi. The Cousins of Stuxnet: Duqu, Flame, and Gauss. *Future Internet*, 4:971–1003, 2012.

[7] A. Gueye, V. Marbukh, and J. C. Walrand. Towards a Metric for Communication Network Vulnerability to Attacks: A Game Theoretic Approach. In *Proc. of Gamenets*, 2012.

[8] Kaspersky Lab. Flame...the latest cyber-attack. http://www.kaspersky.com/flame, 2012.

[9] H. Kunreuther and G. Heal. Interdependent security. *Journal of Risk and Uncertainty*, 26(2-3), 2003.

[10] A. Laszka, G. Horvath, M. Felegyhazi, and L. Buttyán. Flipthem: Modeling targeted attacks with flipit for multiple resources. Technical report,http://www.crysys.hu/~laszka/papers/laszka2014flipthem.pdf, 2014.

[11] A. Laszka, B. Johnson, and J. Grossklags. Mitigating Covert Compromises: A Game-Theoretic Model of Targeted and Non-Targeted Covert Attacks. In *Proc. of WINE*, 2013.

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

[13] M. van Dijk, A. Juels, A. Oprea, and R. L. Rivest. FlipIt: The Game of "Stealthy Takeover". *Journal of Cryptology*, 26(4):655–713, 2013.

[14] M. Zhang, Z. Zheng, and N. B. Shroff. Stealthy attacks and observable defenses: A game theoretic model under strict resource constraints. Tehnical Report, available online at http://www.cse.ohio-state.edu/~zhanming/StealthyAttackFullVersion.pdf.