

A Game Theoretic Model for Defending Against Stealthy Attacks with Limited Resources^{*}

Ming Zhang¹, Zizhan Zheng², and Ness B. Shroff¹

¹ Department of ECE and CSE
The Ohio State University
Columbus, OH, USA
{zhang.2562, shroff.11}@osu.edu

² Department of Computer Science
University of California, Davis
Davis, CA, USA
cszheng@ucdavis.edu

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 trade-offs 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 attack and 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 the best response strategies for both attacker and defender, and study the Nash Equilibria of the game. We further study a sequential game where the defender first announces its strategy and the attacker then responds accordingly, and design an algorithm that finds a nearly optimal strategy for the defender to commit to.

Keywords: Stealthy Attacks, Resource Constraints, Game Theory

1 Introduction

The landscape of cyber security is constantly evolving in response to increasingly sophisticated cyber attacks. In recent years, *Advanced Persistent Threats* (APT) [1] is becoming a major concern to cyber security. APT attacks have several distinguishing properties that render traditional defense mechanism less effective. First, they are often launched by *incentive driven* entities with specific targets. Second, they are *persistent* in achieving the goals, and may involve multiple stages or continuous operations over a long period of time. Third, they are

^{*} This work has been funded by QNRF fund NPRP 5-559-2-227 and ARO-W911NF-15-1-0277.

highly adaptive and *stealthy*, and often operate in a “low-and-slow” fashion [7] to avoid of being detected. In fact, some notorious attacks remained undetected for months or longer [2, 6]. Hence, traditional intrusion detection and prevention techniques that target one-shot and known attack types are insufficient in the face of long-lasting and stealthy attacks.

Moreover, since the last decade, it has been increasingly realized that security failures in information systems are often caused by the misunderstanding of incentives of the entities involved in the system instead of the lack of proper technical mechanisms [5, 17]. To this end, game theoretical models have been extensively applied to cyber security [4, 9–11, 13, 16, 19]. Game theory provides a proper framework to systematically reason about the strategic behavior of each side, and gives insights to the design of cost-effective defense strategies. Traditional game models, however, fail to capture the persistent and stealthy behavior of advanced attacks. Further, they often model the cost of defense (or attack) as part of the utility functions of the players, while ignoring the strict resource constraints during the play of the game. For a large system with many components, ignoring such constraints can lead to either over-provision or under-provision of resources and revenue loss.

In this paper, we study a two-player non-zero-sum game that explicitly models stealth attacks with resource constraints. We consider a system with N *independent* nodes (or components), an attacker, and a defender. Over a continuous time horizon, the attacker (defender) determines when to attack (recapture) a node, subject to a unit cost per action that varies over nodes. At any time t , a node is either compromised or protected, depending on whether the player that makes the last move (i.e., action) towards it before t is the attacker or the defender. A player obtains a value for each node under its control per unit time, which again may vary over nodes. The total payoff to a player is then the total value of the nodes under its control over the entire time horizon minus the total cost incurred, and we are interested in the long-term time average payoffs.

To model stealthy attacks, we assume that the defender gets no feedback about the attacker during the game. On the other hand, the defender’s moves are fully observable to the attacker. This is a reasonable assumption in many cyber security settings, as the attacker can often observe and learn the defender’s behavior before taking actions. Moreover, we explicitly model their resource constraints by placing an upper bound on the frequency of moves (over all the nodes) for each player. We consider both Nash Equilibrium and Sequential Equilibrium for this game model. In the latter case, we assume that the defender is the leader that first announces its strategy, and the attacker then responds with its best strategy. The sequential setting is often relevant in cyber security, and can provide a higher payoff to the defender compared with Nash Equilibrium. To simplify the analysis, we assume that the set of nodes are independent in the sense that the proper functioning of one node does not depend on other nodes, which serves as a first-order approximation of the more general setting of interdependent nodes to be considered in our future work.

Our model is an extension of the asymmetric version of the FlipIt game considered in [15]. The FlipIt game [20] is a two-player non-zero-sum game recently proposed in response to an APT attack towards RSA Data Security [3]. In the FlipIt game, a single critical resource (a node in our model) is considered. 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. A player’s strategy is then when to move over a time horizon, and the solution of the game heavily depends on the class of strategies adopted and the feedback structure of the game. In particular, a full analysis of Nash Equilibria has only been obtained for two special cases, when both players employ a periodic strategy [20], and when the attacker is stealthy and the defender is observable as in our model [15]. However, both works consider a single node and there is no resource constraint. The multi-node setting together with the resource constraints impose significant challenges in characterizing both Nash and Sequential Equilibria. A different multi-node extension of the FlipIt game is considered in [14] where the attacker needs to compromise either all the nodes (AND model) or a single node (OR model) to take over a system. However, only preliminary analytic results are provided.

Our game model can be applied in various settings. One example is key rotation. Consider a system with multiple nodes, e.g., multiple communication links or multiple servers, that are protected by different keys. From time to time, the attacker may compromise some of the keys, e.g., by leveraging zero-day vulnerabilities and system specific knowledge, while remaining undetected from the defender. A common practice is to periodically generate fresh keys by a trusted key-management service, without knowing when they are compromised. On the other hand, the attacker can easily detect the expiration of a key (at an ignorable cost compared with re-compromising it). Both key rotation and compromise incurs a cost, and there is a constraint on the frequency of moves at each side. There are other examples where our extension of the FlipIt game can be useful, such as password reset and virtual-machine refresh [8, 15, 20].

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 in terms of the frequency of protection/attack actions across all the nodes.
- We prove that the periodic strategy is a best-response strategy for the defender against a non-adaptive *i.i.d.* strategy of the attacker, and vice versa, for general distributions of attack times.
- For the above pair of strategies, we fully characterize the set of Nash Equilibria of our game, and show that there is always one (and maybe more) equilibrium, for the case when the attack times are deterministic.
- We further consider the sequential game with the defender as the leader and the attacker as the follower. We design a dynamic programming based algorithm that identifies a nearly optimal strategy (in the sense of subgame perfect equilibrium) for the defender to commit to.

The remainder of this paper is organized as follows. We present our game-theoretic model in Section 2, and study best-response strategies of both players in Section 3. Analysis of Nash Equilibria of the game is provided in Section 4, and the sequential game is studied in Section 5. In Section 6, we present numerical result, and we conclude the paper in Section 7.

2 Game Model

In this section, we discuss our two-player game model including its information structure, the action spaces of both attacker and defender, and their payoffs. Our game model extends the single node model in [15] to multiple nodes and includes a resource constraint to each player.

2.1 Basic Model

In our game-theoretical model, there are two players and N *independent* nodes³. The player who is the lawful user/owner of the N nodes is called the defender, while the other player is called the attacker. The game starts at time $t = 0$ and goes to any time $t = T$. We assume that time is continuous. A player can make a move at any time instance subject to a cost per move. At any time t , a node is under the control of the player that makes the last move towards the node before t (see Figure 1). Each attack towards node i incurs a cost of C_i^A to the attacker, and it takes a random period of time w_i to succeed. 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. Each node i has a value r_i that represents the benefit that the attacker receives from node i per unit of time when node i is compromised.

In addition to the move cost, we introduce a strict resource constraint for each player, which is a practical assumption but has been ignored in most prior works on security games. In particular, 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 typical security games, we assume that r_i, C_i^A, C_i^D , the distribution of w_i , and the budget constraints are all common knowledge of the game, that is, they are known to both players. For instance, they can be learned from history data and domain knowledge. Without loss of generality, all nodes are assumed to be protected at time $t = 0$. Table 1 summarizes the notations used in the paper.

As in [15], we consider an asymmetric feedback model where the attacker’s moves are *stealthy*, while the defenders’ moves are *observable*. More specifically, at any time, the attacker knows the full history of moves by the defender, as well as the state of each node, while the defender has no idea about whether a node is compromised or not. Let $\alpha_{i,k}$ denote the time period the attacker waits from the latest time when node i is recovered, to the time when the attacker starts

³ The terms “components” and “nodes” are interchangeable in this paper.




Fig. 1: Game Model

Table 1: List of Notations

Symbol	Meaning
T	time horizon
N	number of nodes
r_i	value per unit of time of compromising node i
w_i	attack 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 $(k-1)$ -th and the k -th defense for node i
B	budget to the defender, greater than 0
M	budget to the attacker, greater than 0
m_i	frequency of defenses for node i
p_i	probability of immediate attack on node i once it recovers
L_i	the number of defense moves for node i

its k -th attack against node i , which can be a random variable in general. The attacker's action space is then all the possible selections of $\{\alpha_{i,k}\}$. Since the set of nodes are independent, we can assume $\alpha_{i,k}$ to be independent across i without loss of generality. However, they may be correlated across k in general, as the attacker can employ a time-correlated strategy. On the contrary, the defender's strategy is to determine the time intervals between its $(k-1)$ -th move and k -th move for each node i and k , 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 changes its strategy based 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. Note that assuming non-adaptive strategies is not a limitation for the defender since it does not get any feedback during the game anyway. Interestingly, it turns out not to be a big limitation on the attacker either. As we will show in Section 3, periodic defense is a best-response strategy against any non-adaptive *i.i.d.* attacks (formally defined in Definition 2) and vice versa. Note that when the defender's strategy is periodic, the attacker can predict defender's moves before the game starts so there is no need to be adaptive.

2.2 Defender's Problem

Consider a fixed period of time T and let L_i denote the total number of defense moves towards node i during T . L_i is a random variable in general. The total 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 (non-positive) minus the total defense cost over all the nodes. Given the attacker's strategy $\{\alpha_{i,k}\}$, the defender faces the following optimization problem:

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

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.

2.3 Attacker's Problem

We again let L_i denote the total number of defense moves towards 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 $\alpha_{i,k} \geq X_{i,k}$, 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.

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

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

On the other hand, we assume that C_i^A captures the total cost for each attack on node i , which is independent of the attack time. We further have the following equation:

$$\int_0^T v_i(t) dt = \sum_{k=1}^{L_i} (\min(\alpha_{i,k} + w_i, X_{i,k}) - \min(\alpha_{i,k}, X_{i,k})) \quad (3)$$

Putting (3) into (2) and moving the expectation inside, the attacker's problem becomes

$$\begin{aligned} \max_{\alpha_{i,k}} & \sum_{i=1}^N \frac{T \cdot r_i - E[\sum_{k=1}^{L_i} \min(\alpha_{i,k} + w_i, X_{i,k})] \cdot r_i - E[\sum_{k=1}^{L_i} P(\alpha_{i,k} < X_{i,k})] \cdot C_i^A}{T} \\ \text{s.t.} & \sum_{i=1}^N \frac{E[\sum_{k=1}^{L_i} \min(\alpha_{i,k} + w_i, X_{i,k}) - \min(\alpha_{i,k}, X_{i,k})]}{T} \leq M \end{aligned} \quad (4)$$

3 Best Responses

In this section, we analyze the best-response strategies for both players. Our main result is that when the attacker employs a non-adaptive *i.i.d.* strategy, a periodic strategy is a best response for the defender, and vice versa. To prove this result, however, we have provided characterization of best responses in more general settings. In this and following sections, we have omitted most proofs to save space. All the missing proofs can be found in our online technical report [21].

3.1 Defender's Best Response

We first show that for the defender's problem (1), an optimal deterministic strategy is also optimal in general. We then provide a sufficient condition for a deterministic strategy to be optimal against any non-adaptive attacks. Finally, we show that periodic defense is optimal against non-adaptive *i.i.d.* attacks.

Lemma 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.*

According to the lemma, it suffices to consider defender's strategies where both $X_{i,k}$ and $L_{i,k}$ are deterministic.

Definition 1. *For a given L_i , we define a set \mathcal{X}_i including all deterministic defense strategies with the following properties:*

1. $\sum_{k=1}^{L_i} X_{i,k} = T$;
2. $F_{\alpha_{i,k} + w_i}(X_{i,k}) = F_{\alpha_{i,j} + w_i}(X_{i,j}) \quad \forall k, j$,

where $F_{\alpha_{i,k} + w_i}(\cdot)$ is the CDF of r.v. $\alpha_{i,k} + w_i$.

Note that \mathcal{X}_i can be an empty set in general due to the randomness of $\alpha_{i,k} + w_i$. The following lemma shows that when \mathcal{X}_i is non-empty for all i , any strategy that belongs to \mathcal{X}_i is the defender's best deterministic strategy against a non-adaptive attacker.

Lemma 2. *For any given set of $\{L_i\}$ with $\sum_{i=1}^N \frac{L_i}{T} \leq B$, if $\mathcal{X}_i \neq \emptyset \forall i$, then any set of $\{X_{i,k}\}$ which belongs to \mathcal{X}_i , is the defender's best deterministic strategy.*

Lemma 2 gives a sufficient condition for a deterministic defense strategy to be optimal. 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 its expected loss in each period in a deterministic way, which gives the defender the most stable system to avoid a big loss in any particular period. We then show that a periodic defense is sufficient when the attacker employs a non-adaptive *i.i.d.* strategy formally defined below.

Definition 2. *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 1. *A periodical strategy is the best response for the defender if the attacker employs a non-adaptive i.i.d. strategy.*

According to the theorem, the periodic strategy gives the defender the most stable system when the attacker adopts the non-adaptive *i.i.d.* strategy. Since the attacker's waiting time $\alpha_{i,k}$ does not change with time, a fixed defense interval provides the same expected payoff between every two consecutive moves. Moreover, since the defender's problem is a convex optimization problem, the optimal defending frequency for a given attack strategy can be easily determined by solving the convex program.

3.2 Attacker's Best Response

We first analyze the attacker's best response against any deterministic defense strategies, then show that the non-adaptive *i.i.d.* strategy is the best response against periodic defense.

Lemma 3. *When defense strategies are deterministic, the attacker's best response (among non-adaptive strategies) must satisfy the following condition*

$$\alpha_{i,k}^* = \begin{cases} 0 & w.p. p_{i,k} \\ \geq X_{i,k} & w.p. 1 - p_{i,k} \end{cases} \quad (5)$$

Proof sketch: The main idea of the proof is to divide the problem (4) into $\sum_{i=1}^N L_i$ independent sub-problems, one for each node and a single period, where each subproblem has a similar target function and a budget $M_{i,k}$ where $\sum_{i=1}^N \sum_{k=1}^{L_i} M_{i,k} = M$. Due to the independence of nodes, it suffices to prove the lemma for any of these sub-problems.

Lemma 3 implies that for each node i , the attacker's best strategy is to either attack node i immediately after it realizes the node's recovery, or gives up the attack until the defender's next move. There is no incentive for the attacker to wait a small amount of time to attack a node before the defender's next move. 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. We then find the attacker's best responses when the defender employs the periodic strategy.

Theorem 2. *When the defender employs periodical strategy, the non-adaptive *i.i.d.* strategy is the attacker's best response among all non-adaptive strategies.*

3.3 Simplified Optimization Problems

According to Theorem 1 and Theorem 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 optimization problems to the defender and the attacker can then be simplified as follows.

Defender's problem:

$$\begin{aligned} \max_{m_i} \sum_{i=1}^N & \left[\left(E\left[\min\left(w_i, \frac{1}{m_i}\right)\right] p_i r_i - C_i^D \right) \cdot m_i - p_i r_i \right] \\ \text{s.t.} \quad & \sum_{i=1}^N m_i \leq B \end{aligned} \quad (6)$$

Attacker's problem:

$$\begin{aligned} \max_{p_i} \sum_{i=0}^N & p_i \cdot \left(r_i \left(1 - E\left[\min\left(w_i, \frac{1}{m_i}\right)\right] \right) \cdot m_i - C_i^A m_i \right) \\ \text{s.t.} \quad & \sum_{i=0}^N E\left[\min\left(w_i, \frac{1}{m_i}\right)\right] \cdot m_i \cdot p_i \leq M \end{aligned} \quad (7)$$

We observe that the defender's problem is a continuous convex optimization problem (see the discussion in Section 3.1), and the attacker's problem is a fractional knapsack problem. Therefore, the best response strategy of each side can be easily determined. Also, the time period T disappears in both problems.

4 Nash Equilibria

In this section, we study the set of Nash Equilibria of the simplified game as discussed in Section 3.3 where the defender employs a periodic strategy, and the attacker employs a non-adaptive *i.i.d.* strategy. We further assume that the

attack time w_i is deterministic for all i . We show that this game always has a Nash equilibrium and may have multiple equilibria of different values.

We first observe that for deterministic w_i , when $m_i \geq \frac{1}{w_i}$, the defender's payoff becomes $-m_i C_i^D$, which is maximized when $m_i = \frac{1}{w_i}$. Therefore, it suffices to consider $m_i \leq \frac{1}{w_i}$. Thus, the optimization problems to the defender and the attacker can be further simplified as follows.

For a given p , the defender aims at maximizing its payoff:

$$\begin{aligned} & \max_{m_i} \sum_{i=1}^N [m_i(r_i w_i p_i - C_i^D) - p_i r_i] \\ & \text{s.t.} \quad \sum_{i=1}^N m_i \leq B \\ & \quad 0 \leq m_i \leq \frac{1}{w_i}, \forall i \end{aligned} \tag{8}$$

On the other hand, for a given m , the attacker aims at maximizing its payoff:

$$\begin{aligned} & \max_{p_i} \sum_{i=1}^N p_i [r_i - m_i(r_i w_i + C_i^A)] \\ & \text{s.t.} \quad \sum_{i=1}^N m_i w_i p_i \leq M \\ & \quad 0 \leq p_i \leq 1, \forall i \end{aligned} \tag{9}$$

For a pair of strategies (m, p) , the payoff to the defender is $U_d(m, p) = \sum_{i=1}^N [m_i(p_i r_i w_i - C_i^D) - p_i r_i]$, 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, while if $\frac{C_i^D}{r_i w_i} > 1$, the coefficient of m_i in U_d (defined below) is always negative and there is no need 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 a component more frequently 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 whole 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 the following 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 $D(m, p) \triangleq \{i \in F : \rho_i(m) = \rho^*(m)\}$. We omit m and p when they are clear from the context.

Lemma 4. *If (m, p) is an NE, we have :*

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

Lemma 5. *If (m, p) forms an NE, then for $i \in D, j \in F \setminus D$ 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 the above lemma, to find all the equilibria of the game, it suffices to sort all the nodes by 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 $D_k \subseteq F_h$ 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 $D \subseteq F$. By Lemma 4 and the definitions of F and D , the following conditions are satisfied by any NE with $F(p) = F$ and $D(m, p) = D$.

$$m_i = 0, p_i = 1, \forall i \notin F; \quad (10)$$

$$m_i \in [0, \frac{r_i}{w_i r_i + C_i^A}], p_i = 1, \forall i \in F \setminus D; \quad (11)$$

$$m_i \in [0, \frac{r_i}{w_i r_i + C_i^A}], p_i \in [\frac{C_i^D}{r_i w_i}, 1], \forall i \in D; \quad (12)$$

$$\sum_{i \in F} m_i \leq B, \sum_{i \in F} m_i w_i p_i \leq M; \quad (13)$$

$$\mu_i = \mu^*, \forall i \in F; \quad \mu_i < \mu^*, \forall i \notin F; \quad (14)$$

$$\rho_i = \rho^*, \forall i \in D; \quad \rho_i > \rho^*, \forall i \notin D. \quad (15)$$

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

Theorem 3. *Any pair of strategies (m, p) with $F(p) = F$ and $D(m, p) = D$ is an NE iff it is a solution to one of the following sets of constraints in addition to (10) to (15).*

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. The next theorem shows that our game has at least one equilibrium and may have more than one NE.

Theorem 4. *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 [21]. 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 $M = 1/5.$ 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. \square

5 Sequential Game

In this section, we study a sequential version of the simplified game considered in the last section. In the simultaneous game we considered in the previous section, neither the defender nor the attacker can learn the opponent's strategy in advance. While this is a reasonable assumption for the defender, an advanced attacker can often observe and learn defender's strategy before launching attacks. It therefore makes sense to consider the setting where the defender first commits to a strategy and makes it public, the attacker then responds accordingly. Such a sequential game can actually provide defender higher payoff comparing to a Nash Equilibrium since it gives the defender the opportunity of deterring the attacker from moving. We again focus on non-adaptive strategies, and further assume that at $t = 0,$ the leader (defender) has determined its strategy, and the follower (attacker) has learned the defender's strategy and determined its own strategy in response. In addition, the players do not change their strategies thereafter. Our objective is to identify the best sequential strategy for the defender to commit to, in the sense of subgame perfect equilibrium [18] defined as follows. We again focus on the case where w_i is deterministic for all $i.$

Definition 3. *A pair of strategies (m^*, p^*) is a subgame perfect equilibrium of the simplified game (8) and (9) if m^* is the optimal solution of*

$$\begin{aligned}
 & \max_{m_i} \sum_{i=1}^N [m_i(r_i w_i p_i^* - C_i^D) - p_i^* r_i] \\
 & s.t. \quad \sum_{i=1}^N m_i \leq B \\
 & \quad \quad 0 \leq m_i \leq \frac{1}{w_i}, \forall i
 \end{aligned} \tag{16}$$

where p_i^* is the optimal solution of

$$\begin{aligned} \max_{p_i} \quad & \sum_{i=1}^N p_i [r_i - m_i(r_i w_i + C_i^A)] \\ \text{s.t.} \quad & \sum_{i=1}^N m_i w_i p_i \leq M \\ & 0 \leq p_i \leq 1, \forall i \end{aligned} \tag{17}$$

Note that in a subgame perfect equilibrium, p_i^* is still the optimal solution of (9) as in a Nash Equilibrium. However, defender's best strategy m_i^* is not necessarily optimal with respect to (8). Due to the multi-node setting and the resource constraints, it is very challenging to identify an exact subgame perfect equilibrium strategy for the defender. To this end, we propose a dynamic programming based algorithm that finds a nearly optimal defense strategy.

Remark 1. Since for any given defense strategy $\{m_i\}$, the attacker's problem (17) is a fractional knapsack problem, the optimal $p_i, \forall i$ has the following form: Sort the set of nodes by $\rho_i(m_i) = \frac{r_i - m_i(r_i w_i + C_i^A)}{m_i w_i}$ non-increasingly, then there is an index k such that $p_i = 1$ for the first k nodes, and $p_i \leq 1$ for the $k+1$ -th node, and $p_i = 0$ for the rest nodes. However, if $\rho_i = \rho_j$ for some $i \neq j$, the optimal attack strategy is not unique. When this happens, we assume that the attacker always breaks ties in favor of the defender, a common practice in Stackelberg security games [12].

Before we present our algorithm to the problem, we first establish the following structural properties on the subgame perfect equilibria of the game.

Lemma 6. *In any subgame perfect equilibrium (m, p) , the set of nodes can be partitioned into the following four disjoint sets according to the attack and defense strategies applied:*

1. $F = \{i | m_i > 0, p_i = 1\}$
2. $D = \{i | m_i > 0, 0 < p_i < 1\}$;
3. $E = \{i | m_i > 0, p_i = 0\}$;
4. $G = \{i | m_i = 0, p_i = 1\}$.

Moreover, they satisfy the following properties:

1. $F \cup D \cup E \cup G = \{i | i = 1, \dots, n\}$ and $|D| \leq 1$
2. $\rho_i \geq \rho_k \geq \rho_j$ for $\forall i \in F, k \in D, j \in E$

Since the set D has at most one element, we use m_d to represent $m_i, i \in D$ for simplicity, and let $\rho_d = \rho(m_d)$. If D is empty, we pick any node i in F with minimum ρ_i and treat it as a node in D .

Lemma 7. *For any given nonnegative ρ_d , the optimal solution for (16)-(17) satisfy the following properties:*

1. $r_i w_i - C_i^D > 0 \forall i \in F \cup E \cup D$
2. $m_i \leq \bar{m}_i \forall i \in F$
3. $m_j = \bar{m}_j \forall j \in E$
4. $\bar{m}_i \leq \frac{1}{w_i} \forall i$
5. $B - \sum_{i \in E} \bar{m}_i - m_d > 0$.

where $\bar{m}_i = m_i(\rho_d)$ and $m_i(\cdot)$ is the reverse function of $\rho_i(\cdot)$

Remark 2. If $\rho_d < 0$, the defender can give less budget to the corresponding node to bring ρ_d down to 0. In any case, the payoffs from nodes in set D and E are 0 since the attacker will give up attacking the nodes in set D and E . Thus, the defender has more budget to defend the nodes in set F and G which brings him more payoffs. Therefore we only need to consider nonnegative ρ_d .

Lemma 8. *For any nonnegative ρ_d , there exists an optimal solution for (16)-(17) such that $\forall i \in F$, there are at most two $m_i < \bar{m}_i$ and all the other $m_i = \bar{m}_i$*

From the above lemmas, we can establish the following results about the structure of the optimal solution for (16)-(17).

Proposition 1. *For any nonnegative ρ_d , there exists an optimal solution $\{m_i\}_{i=1}^n$ such that*

1. $\forall i \in F$, there are at most two $m_i < \bar{m}_i$ and all the other $m_i = \bar{m}_i$;
2. $m_d = \bar{m}_d$
3. $\forall i \in E$, $m_i = \bar{m}_i$;
4. $\forall i \in G$, $m_i = 0$.

According to Proposition 1, for any nonnegative ρ_d , once the set allocation is determined, the value of m_i can be immediately determined for all the nodes except the two fractional nodes in set F . Further, for the two fractional nodes, their m_i can be found using linear programming as discussed below. From these observations, we can convert (16)-(17) to (18) for any given nonnegative ρ_d , d , f_1 and f_2 .

$$\begin{aligned}
& \max_{p, m_{f_1}, m_{f_2}, E, F, G} \sum_{i \in F \setminus \{f_1, f_2\}} [\bar{m}_i (r_i w_i - C_i^D) - r_i] + \sum_{j=1}^2 [m_{f_j} (r_{f_j} w_{f_j} - C_{f_j}^D) - r_{f_j}] \\
& \quad - \sum_{i \in G} r_i - \sum_{i \in E} \bar{m}_i C_i^D + m_d (p r_d w_d - C_d^D) - p r_d \\
& \text{s.t.} \quad \sum_{i \in F \setminus \{f_1, f_2\}} \bar{m}_i + m_{f_1} + m_{f_2} + \sum_{i \in E} \bar{m}_i + m_d \leq B \\
& \quad \sum_{i \in F \setminus \{f_1, f_2\}} w_i \bar{m}_i + w_{f_1} m_{f_1} + w_{f_2} m_{f_2} + p w_d m_d \leq M \\
& \quad 0 \leq m_{f_1} \leq \bar{m}_1, 0 \leq m_{f_2} \leq \bar{m}_2, 0 \leq p \leq 1
\end{aligned} \tag{18}$$

Note that, the set allocation is part of the decision variables in (18).

We then propose the following algorithm to the defender's problem (see Algorithm 1). The algorithm iterates over nonnegative ρ_d (with a step size ρ_{step}) (lines 3-10). For each ρ_d , it iterates over all possible node d in set D , and all possible nodes f_1, f_2 with fractional assignment in set F (lines 5-8). Given ρ_d, d, f_1, f_2 , the best set allocation (together with m_i for all i and p) are determined using dynamic programming as explained below (lines 6-7), where we first assume that B, M, \bar{m}_i and w_i have been rounded to integers for all i . The loss of performance due to rounding will be discussed later.

Consider any ρ_d , node d is in set D , and nodes f_1, f_2 with fractional assignment in set F . Let $SEQ(i, b, m, d, f_1, f_2, ind)$ denote the maximum payoff of the defender considering only node 1 to node i (excluding nodes d, f_1 and f_2), for given budgets b and m for the two constraints in (18), respectively. The ind is a boolean variable that indicates whether the second constraint of (18) is tight for node 1 to i . If ind is *True*, it means all the budget m is used up for node 1 to i . ind is *False* meaning that there is still budget m available for the attacker. Here, $0 \leq b \leq B$ and $0 \leq m \leq M$. The value of $SEQ(i, b, m, d, f_1, f_2, ind)$ is determined recursively as follows. If $b < 0$ or $m < 0$, the value is set to $-\infty$. If node i is one of d, f_1 and f_2 , we simply set $SEQ(i, b, m, d, f_1, f_2, ind) = SEQ(i - 1, b, m, d, f_1, f_2, ind)$. Otherwise, we have the following recurrence equation, where the three cases refer to the maximum payoff when putting nodes i in set F, E , and G , respectively.

$$\begin{aligned}
& SEQ(i, b, m, d, f_1, f_2, ind) \\
&= \max \left\{ SEQ(i - 1, b - \bar{m}_i, m - w_i \bar{m}_i, d, f_1, f_2, ind) + \bar{m}_i(r_i w_i - C_i^D) - r_i, \right. \\
& \left. SEQ(i - 1, b - \bar{m}_i, m, d, f_1, f_2, ind) - \bar{m}_i C_i^D, SEQ(i - 1, b, m, d, f_1, f_2, ind) - r_i \right\}
\end{aligned} \tag{19}$$

Meanwhile, if ind is *False*, node i can be allocated to set E only if $r_i - \bar{m}_i(r_i w_i + C_i^A) \leq 0$. Otherwise, there is still available budget for the attacker to attack other nodes with reward greater than 0 which violates the structure of the greedy solution for (17). Also, if ind is *False*, it means m is not used up. Thus we should return $-\infty$ if ind is *False*, $i > 0$ and $m = 0$.

Moreover, we let $SEQ(0, b, m, d, f_1, f_2, ind)$ denote the maximum defense payoff when only nodes in d, f_1 , and f_2 are considered. If ind is *True*, the following linear program in (20) determines the optimal values of p, m_{f_1} and m_{f_2} for given budgets b and m :

$$\begin{aligned}
& \max_{m_{f_1}, m_{f_2}} \sum_{j=1}^2 [m_{f_j}(r_{f_j} w_{f_j} - C_{f_j}^D) - r_{f_j}] + m_d(pr_d w_d - C_d^D) - pr_d \\
& \text{s.t. } m_{f_1} + m_{f_2} + m_d \leq b \\
& \quad m_{f_1} w_{f_1} + m_{f_2} w_{f_2} \leq m \\
& \quad m_{f_1} \leq \bar{m}_{f_1}, \quad m_{f_2} \leq \bar{m}_{f_2} \\
& \quad p = \frac{m - m_{f_1} w_{f_1} - m_{f_2} w_{f_2}}{w_d m_d} \leq 1
\end{aligned} \tag{20}$$

If *ind* is *False*, we must have $p = 1$. The optimal values of m_{f_1} and m_{f_2} are determined by (21):

$$\begin{aligned} & \max_{m_{f_1}, m_{f_2}} \sum_{j=1}^2 [m_{f_j} (r_{f_j} w_{f_j} - C_{f_j}^D) - r_{f_j}] + m_d (r_d w_d - C_d^D) - r_d \\ & \text{s.t. } m_{f_1} + m_{f_2} + m_d \leq b \\ & \quad m_{f_1} w_{f_1} + m_{f_2} w_{f_2} \leq m - w_d m_d \\ & \quad m_{f_1} \leq \bar{m}_{f_1}, \quad m_{f_2} \leq \bar{m}_{f_2} \end{aligned} \quad (21)$$

Algorithm 1 Sequential Strategy for Defender

```

1: Initialize  $\rho_{step}$ 
2:  $\rho_{max} \leftarrow \min\{\rho : \sum_{i=1}^n w_i m_i(\rho) \leq M\}$ 
3: for  $\rho_d \leftarrow 0$  to  $\rho_{max}$  with step size  $\rho_{step}$  do
4:    $\bar{m}_i \leftarrow m_i(\rho_d)$  for all  $i$ 
5:   for  $d, f_1, f_2 \leftarrow 1$  to  $n$  do
6:      $val_{d, f_1, f_2} \leftarrow SEQ(n, B, M, d, f_1, f_2, True)$ 
7:      $val'_{d, f_1, f_2} \leftarrow SEQ(n, B, M, d, f_1, f_2, False)$ 
8:   end for
9:    $C_{dp}(\rho_d) \leftarrow \max_{d, f_1, f_2} \{val_{d, f_1, f_2}, val'_{d, f_1, f_2}\}$ 
10: end for
11:  $C_{alg}^* \leftarrow \max_{\rho_d} \{C_{dp}(\rho_d)\}$ 

```

Since the dynamic program searches for all the possible solutions that satisfy Proposition 1, $C_{dp}(\rho_d)$ gives us the optimal solution of (16)-(17) for any given nonnegative ρ_d . Algorithm 1 then computes the optimal solution by searching all the nonnegative ρ_d . Note that d, f_1 and f_2 can be equal to include the case that there is only one or zero node in set F . The minimum possible value of ρ is 0 (explained in Remark 2). The maximum possible value of ρ is $\min\{\rho : \sum_{i=1}^n w_i m_i(\rho) \leq M\}$. For larger ρ , the sum of all $w_i \bar{m}_i$ will be less than M . In this case, all the nodes will be in set F and $p_i = 1 \forall i$, which makes (16)-(17) a simple knapsack problem that can be easily solved.

Additionally, since the dynamic program searches over all feasible integer values, we use a simple rounding technique to guarantee it is implementable. Before the execution of $SEQ(n, B, M, d, f_1, f_2, ind)$, we set $\bar{m}_i \leftarrow \lfloor \frac{\bar{m}_i}{\delta} \rfloor$, $w_i \leftarrow \lfloor \frac{w_i}{\delta} \rfloor$ for all i and $B \leftarrow \lfloor \frac{B}{\delta} \rfloor$, $M \leftarrow \lfloor \frac{M}{\delta} \rfloor$ where δ is an adjustable parameter. Intuitively, by making δ and ρ_{step} small enough, Algorithm 1 can find a strategy that is arbitrarily close to the subgame perfect equilibrium strategy of the defender. Formally, we can establish the following result.

Theorem 5. *Let C_{alg} denote the payoffs of the strategy found by Algorithm 1, and C^* the optimal payoffs. Then for any $\epsilon > 0$, Algorithm 1 can ensure that $\frac{|C_{alg}|}{|C^*|} \leq 1 + \epsilon$ with a total time complexity of $O(\frac{n^8 B M}{\epsilon^3})$, where B and M are values before rounding.*

Note that both C_{alg} and C^* are non-positive. The details can be found in our online technical report [21].

6 Numerical Result

In this section, we present numerical results for our game models. For the illustrations, we assume that all the attack times w_i are deterministic as in Sections 4 and 5. We study the payoffs of both attacker and defender and their strategies in both Nash Equilibrium and subgame perfect equilibrium in a two-node setting, and study the impact of various parameters including resource constraints B , M , and the unit value r_i . We further study the payoffs and strategies for both players in subgame perfect equilibrium in a five-node setting, and study the impact of various parameters.

We first study the impact of the resource constraints M , B , and the unit value r_1 on the payoffs for the two node setting in Figure 2. In the figure, we have plotted both Type 1 and Type 5 NE⁴ and subgame perfect equilibrium. Type 5 NE only occurs when M is small as shown in Figure 2(a), while Type 1 NE appears when B is small as shown in Figure 2(b), which is expected since B is fully utilized in a Type 1 NE while M is fully utilized in a Type 5 NE. When the defense budget B becomes large, the summation of m_i does not necessarily equal to B and thus Type 1 NE disappears. Similarly, the Type 5 NE disappears for large attack budget M . In Figures 2(c) and 2(d), we vary the unit value of node 1, r_1 . At the beginning, the defender protects node 2 only since $w_2 > w_1$. As r_1 becomes larger and larger, the defender starts to change its strategy by protecting node 1 instead of node 2 in NE Type 1. On the other hand, since node 1 is fully protected by the defender and the defender gives up defending node 2, the attacker begins to attack node 2 with probability 1, and uses the rest budget to attack node 1 with probability less than 1, due to the high defending frequency and limited resources M . We further observe that in both the simultaneous game and the sequential game, the value of m_1 increases along with the increase of r_1 , while the value of m_2 decreases at the same time. This implies that the defender tends to protect the nodes with higher values more frequently. In addition, the subgame perfect equilibrium always bring the defender higher payoffs compared with Nash Equilibrium, which is expected.

Moreover, it is interesting to observe that under the Type 5 NE, the attacker's payoff decreases for a larger M as shown in Fig 2(a). This is because the defender's budget B is not fully utilized in Type 5 NE, and the defender can use more budget to protect both nodes when M increases. The increase of the attacker's payoff by having a larger M is canceled by the increase of the defender's move frequency m_1 and m_2 . We also note that the Type 5 NE is less preferable for the defender in Figure-2(c) when r_1 is small and favors defender as r_1 increases, which tells us that the defender may prefer different types of NEs under different scenarios and so does the attacker.

We then study the effects of varying M and r_1 on both players' payoffs and strategies in the sequential game for the five-node setting. In Figure 3(a), the parameters of all the nodes are the same except r_i . We vary the attacker's budget M from 0 to 1. When $M = 0$, the defender can set m_i for all i to arbitrary small

⁴ There are also Type 2 NE, which are omitted for the sake of clarity.

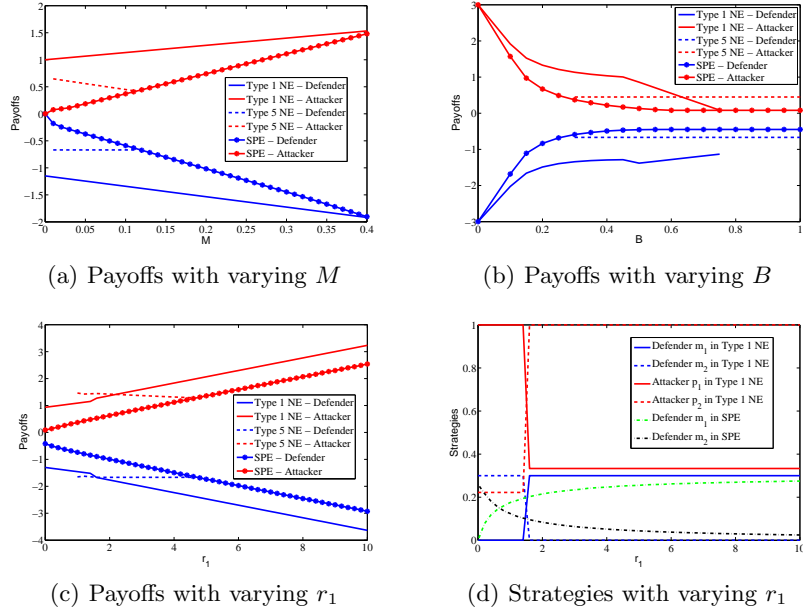


Fig. 2: The effects of varying resource constraints, where in all the figures, $r_2 = 1, w_1 = 1.7, w_2 = 1.6, C_1^D = 0.5, C_2^D = 0.6, C_1^A = 1, C_2^A = 1.5$, and $r_1 = 2$ in (a) and (b), $B = 0.3$ in (a), (c), and (d), and $M = 0.1$ in (b), (c), and (d).




Fig. 3: The effects of varying resource constraints and r_1 , where $w = [2 \ 2 \ 2 \ 2 \ 2]$, $C^D = C^A = [1 \ 1 \ 1 \ 1 \ 1]$, $B = 0.5$, $r = [5 \ 4 \ 3 \ 2 \ 1]$ in (a), $r = [r_1 \ 1 \ 1 \ 1 \ 1]$ and $M = 0.3$ in (b).

(but positive) values, so that the attacker is unable to attack any node, leading to a zero payoff for both players. As M becomes larger, the attacker's payoff increases, while the defender's payoff decreases, and the defender tends to defend the nodes with higher values more frequently, as shown in Figure 3(a)(lower). After a certain point, the defender gives up some nodes and protects higher value nodes more often. This is because with a very large M , the attacker is able to attack all the nodes with high probability, so that defending all the nodes with small m_i is less effective than defending high value nodes with large m_i . This result implies that the attacker's resource constraint has a significant impact on the defender's behavior and when M is large, protecting high value nodes more frequently and giving up several low value nodes is more beneficial for the defender compared to defending all the nodes with low frequency.

In Figure 3(b), we vary r_1 while setting other parameters to be the same for all the nodes. Since all the nodes other than node 1 are identical, they have the same m_i as shown in Figure 3(b)(lower). We observe that the defender protects node 1 less frequently when r_1 is smaller than the unit value of other nodes. When r_1 becomes larger, the defender defends node 1 more frequently, which tells us the defender should protect the nodes with higher values more frequently in the subgame perfect equilibrium when all the other parameters are the same.

7 Conclusion

In this paper, we propose a two-player non-zero-sum game for protecting a system of multiple components against a stealthy attacker where the defender's behavior is fully observable, and both players have strict resource constraints. We prove that periodic defense and non-adaptive *i.i.d.* attack are a pair of best-response strategies with respect to each other. For this pair of strategies, we 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. We further study the sequential game where the defender first publicly announces its strategy, and design an algorithm that can identify a strategy that is arbitrarily close to the subgame perfect equilibrium strategy for the defender.

References

1. Advanced persistent threat. http://en.wikipedia.org/wiki/Advanced_persistent_threat.
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. R. Anderson. Why Information Security is Hard - An Economic Perspective. In *Proc. of ACSAC*, 2001.
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. K. D. Bowers, M. E. V. Dijk, A. Juels, A. M. Oprea, R. L. Rivest, and N. Triandopoulos. Graph-based approach to deterring persistent security threats. US Patent 8813234, 2014.
8. K. D. Bowers, M. van Dijk, R. Grifn, A. Juels, A. Oprea, R. L. Rivest, and N. Triandopoulos. Defending Against the Unknown Enemy: Applying FLIPIT to System Security. In *Proc. of GameSec*, 2012.
9. L. Buttyan and J.-P. Hubaux. *Security and Cooperation in Wireless Networks: Thwarting Malicious and Selfish Behavior in the Age of Ubiquitous Computing*. Cambridge University Press, 2007.
10. 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.
11. M. Kearns and L. E. Ortiz. Algorithms for Interdependent Security Games. In *Proc. of NIPS*, 2003.
12. D. Korzhyk, Z. Yin, C. Kiekintveld, V. Conitzer, and M. Tambe. Stackelberg vs. Nash in Security Games: An Extended Investigation of Interchangeability, Equivalence, and Uniqueness. *Journal of Artificial Intelligence Research*, 41:297–327, 2011.
13. H. Kunreuther and G. Heal. Interdependent security. *Journal of Risk and Uncertainty*, 26(2-3), 2003.
14. A. Laszka, G. Horvath, M. Felegyhazi, and L. Buttyán. Flipthem: Modeling targeted attacks with flipit for multiple resources. In *Proc. of GameSec*, 2014.
15. 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.
16. M. H. Manshaei, Q. Zhu, T. Alpcan, and T. Başar. Game Theory Meets Network Security and Privacy. *ACM Computing Surveys*, 2012.
17. T. Moore and R. Anderson. Economics and internet security: a survey of recent analytical, empirical and behavioral research. <ftp://ftp.deas.harvard.edu/techreports/tr-03-11.pdf>, 2011.
18. M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, 1994.
19. M. Tambe. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, 2011.
20. 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.
21. M. Zhang, Z. Zheng, and N. B. Shroff. A game theoretic model for defending against stealthy attacks with limited resources. Technical Report, available online at <http://arxiv.org/abs/1508.01950>.