A Signaling Game Model for Moving Target Defense

Xiaotao Feng*, Zizhan Zheng†, Derya Cansever‡, Ananthram Swami§ and Prasant Mohapatra¶
*Department of Electrical and Computer Engineering, University of California, Davis, USA,
†Department of Computer Science, Tulane University, New Orleans, USA,
‡US Army CERDEC, USA
§U.S. Army Research Laboratory, USA,
¶Department of Computer Science, University of California, Davis, USA,
Email: {xtfeng, pmohapatra}@ucdavis.edu, zzheng3@tulane.edu, {derya.h.cansever.civ, ananthram.swami.civ}@mail.mil

***Abstract*—Incentive-driven advanced attacks have become a major concern to cyber-security. Traditional defense techniques that adopt a passive and static approach by assuming a fixed attack type are insufficient in the face of highly adaptive and stealthy attacks. In particular, a passive defense approach often creates information asymmetry where the attacker knows more about the defender. To this end, moving target defense (MTD) has emerged as a promising way to reverse this information asymmetry. The main idea of MTD is to (continuously) change certain aspects of the system under control to increase the attacker's uncertainty, which in turn increases attack cost/complexity and reduces the chance of a successful exploit in a given amount of time. In this paper, we go one step beyond and show that MTD can be further improved when combined with information disclosure. In particular, we consider that the defender adopts a MTD strategy to protect a critical resource across a network of nodes, and propose a Bayesian Stackelberg game model with the defender as the leader and the attacker as the follower. After fully characterizing the defender's optimal migration strategies, we show that the defender can design a signaling scheme to exploit the uncertainty created by MTD to further affect the attacker's behavior for its own advantage. We obtain conditions under which signaling is useful, and show that strategic information disclosure can be a promising way to further reverse the information asymmetry and achieve more efficient active defense.**

I. Introduction

Advanced cyber-attacks have become a major concern to security engineers these days. One example is advanced persistent threats (APT), an emerging class of continuous and stealthy hacking processes launched by highly motivated entities with specific targets in mind [1]. These attacks are highly persistent and adaptive in achieving their goals. For instance, they may operate in a stealthy way to avoid detection and obtain long-term advantages. Traditional defense techniques targeting one-shot attacks with known types are insufficient in the face of these more advanced attacks with unique behavioral patterns. To this end, optimal decision and game theoretic approaches have been introduced into cyber-security in recent years to better reason about the strategic behavior of the attacker (and the defender). While these models provide valuable insights on designing more efficient defense strategies, most of them adopt a *passive defense* approach without exploring the attacker's cognitive or resource constraint.

Passive defense is insufficient in the face of advanced attacks. In particular, it creates information asymmetry where the attacker knows more about the defender than the defender knows about the attacker, which is an important obstacle to achieving more efficient cyber-defense. This is especially true in the case of APT, where the attacker may take time to observe and predict the defender's strategy before taking its action. To this end, Moving Target Defense (MTD) [2], [3] is emerging as a promising way to achieve *active defense*. The main idea of MTD is to (continuously) change certain aspects of the system under control to increase the attacker's uncertainty, which in turn increases attack cost/complexity and reduces the chance of a successful exploit in a given amount of time. To date, MTD has been studied in various contexts, including cloud computing [4], [5], web applications [6], [7], and game theoretic approaches have been applied to devise efficient MTD strategies [8], [9], [10].

In this paper, we go one step beyond by showing that MTD can be further improved when combined with information disclosure. In particular, we show that the defender may design a signaling scheme that exploits the uncertainty created by MTD to modify the attacker's behavior for its own advantage. We envision that strategic information disclosure can be a promising way to further reverse the information asymmetry and achieve more efficient active defense.

More specifically, we consider that a defender protects a security sensitive resource across a network of nodes to make it difficult for the attacker to identify the real location of the resource. We model this setting as a Bayesian Stackelberg game where the defender first determines a *randomized* strategy in terms of the conditional probabilities of moving the resource from one node to the other for every pair of nodes. We consider a single resource for the sake of simplicity. Our model can be readily generalized to the case of multiple independent resources. The defender first commits to a strategy and declares it (but not its realization) to the attacker at the beginning of the game. The attacker then decides whether to attack or not and which node to attack based on this prior information and its attack cost. We assume that the attacker can attack at most one node each time. Crucially, there is a migration cost associated with nodes, which needs to be taken into account by the defender. Building upon this model, we then introduce a signaling scheme to further improve the defender's payoff. The idea is that after the defender samples an action from its declared migration strategy, it may send a signal to the attacker that discloses part of its realized strategy to affect the attacker's posterior belief and the corresponding attack behavior.

To design the signaling scheme, we adopt the influential Bayesian Persuasion model [11], which is a variant of signaling games with commitment, and is also more tractable than traditional signaling games. We derive the set of subgame perfect equilibria of the basic MTD game as well as the enhanced game when signaling is applied, and identify complete conditions under which signaling is useful. We observe that signaling can be a useful tool for the defender to deter the attacker under a broad setting of system parameters and can sometimes significantly improve the defender's payoff compared to a pure migration strategy.

We make the following contributions in this paper:

- We show that MTD can be further improved through strategic information disclosure to the attacker;
- We propose a Bayesian Stackelberg game that models the joint migration and signaling strategies for the defender in the face of a strategic and rational attacker.
- We thoroughly investigate the proposed game model and characterize the subgame perfect equilibria of the game under a pure migration setting and when both migration and signaling are applied.
- We derive several insights from equilibrium analysis and numerical study, and make suggestions to the defender on realizing a more efficient MTD.

The remainder of the paper is organized as follows. We discuss the related work in Section II and propose the game model in Section III. Detailed analysis of strategies and equilibria derivation are presented in Section IV and Section V, respectively. The performance of optimal strategies under different system settings is evaluated via numerical study in Section VI. Finally, we conclude the paper in Section VII.

II. Related Work

A. Moving Target Defense

One main idea of Moving Target Defense (MTD) is to hinder the attacker from discovering vulnerabilities or critical resource of the system. Generally, MTD schemes can be categorized into five different domains [12]: (1) *Dynamic Network*: where network properties are continuously modified [13], (2) *Dynamic Platforms*: where the computing platform properties are changed [14], (3) *Dynamic Runtime Environment*: where the environment in which the application operates is randomized [15], (4) *Dynamic Software*: where the application code is shifted while ensuring that the functionality is unaffected [6], (5) *Dynamic Data*: where the internal or external representation of an application's data is changed [16]. In the context of game theory, each domain has substantial literature related to our work. An approach which introduces MTD into the FlipIt game model [17] has been studied in [18] where the attacker randomly selected a server to attack (which might take some time) and the defender could neutralize the attack by protecting that server (with a probability). A dynamic platform MTD which models an attacker who has feedback on defender's moves has been investigated in [19]. [10] proposed a game model in which the defender adopted diversity defense to prevent the attacker from finding the vulnerability of a server. [6] proposed a repeated Bayesian game to model the switching technologies of web applications and derived an effective strategy while considering the cost of switching between different web-stack configurations.

Rather than focusing on a specific application, we propose a general game framework which is applicable to most scenarios in the five domains mentioned above. In particular, the resource in our model can represent a critical server, a software object such as a web application, etc.. The nodes can be physical or virtual machines, types of data format, types of application technique, etc.. The migration cost can represent the expenses on changing network properties, computing platforms, programming techniques, etc.. The network topology can be considered as the internal restriction of migration. The main focus of this paper is to understand the impact of signaling in MTD. Therefore, we develop a simple but generic and flexible two-player MTD model which is appropriate for rigorously analyzing the complicated strategies.

B. Signaling Game

Another aspect of our model is the signaling game model which has been extensively studied in the literature. Bayesian Persuasion model [11] captures the general concept of the signaling game. The main contribution of the Bayesian Persuasion model is in answering the question: when and how does the sender (the defender in our model) exploit the asymmetric information to persuade the receiver (the attacker in our model). [20] investigated the optimal information disclosure problems in the signaling game. [21] examined the computational complexity problem of the Bayesian Persuasion model. Deception games can be considered as another application of the signaling model. [22], [23] studied a two-player game in which one player increases the other player's uncertainty by either disguising a normal node as a honeypot or disguising a honeypot as a normal node. The common feature of the above work is that the state of the system is not decided by the sender (defender). In our model, the defender can not only observe the system states but also employ strategies to change the system states. [24] investigated the general case where the sender has additional private information. However, [24] did not solve the defender's optimal strategies in an explicit form. By introducing a concrete MTD model, we solve the defender's optimal strategies and characterize the equilibria of the defender-attacker game under general system settings.

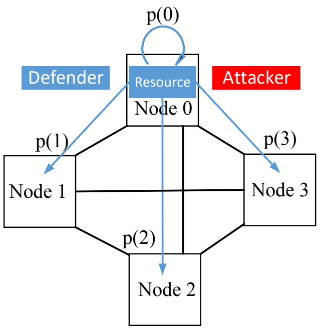

Figure 1: An Example of a 4-Node System

III. Signaling Game on Moving Target Defense

In this section, we propose a two-player game model in which the defender adopts a moving target defense scheme, *e.g.*, a migration strategy, to thwart a strategic attacker. MTD creates uncertainty to the attacker, which can be further exploited by the defender to improve her performance through strategic information disclosure. To this end, we adopt the influential Bayesian Persuasion approach to design an effective signaling scheme over MTD. In the following, we first provide the formal definition of the game including the action spaces and objective functions of the players (Section III-A). We then describe an example primarily to show how Bayesian Persuasion, when designed properly, can help reshape the interactions between the attacker and the defender, leading to a reduced loss for the defender (Section III-C).

A. Game Model

We consider a system in which a strategic attacker can observe the defender's actions and then move in a rational way. The system consists of four components: a critical resource, a fully connected network with n $(n \geq 2)$ nodes, one attacker and one defender. The goal of the attacker is to maximize its payoff by comprising the resource, while the goal of the defender is to protect the resource at a minimum cost. Figure 1 illustrates an example of a 4-node system. The interplay between the defender and the attacker forms a two-player game.

In this game, the defender adopts a MTD strategy by migrating the resource across the network to make it difficult for the attacker to identify the real location of the resource. Since the attacker may observe the defender's actions by monitoring network traffic, we consider a Bayesian Stackelberg game by allowing the defender to commit to and declare a randomized migration strategy. Knowing this strategy (but not its realization), the attacker then determines which node to attack or not attack. The defender can further employ a signaling strategy by sending a message to the attacker to affect its behavior. We consider a one stage game where the defender migrates the resource once and the attacker takes one corresponding action.

At the beginning of the game, the resource is installed at node 0 (randomly selected). The defender could migrate the resource from node 0 to one of the $n-1$ nodes in the network, or keep the resource at node 0. Denote $\mathbf{p} = \{p(i)|p(i) \in [0,1], i = 0, 1, n-1, \sum_{i=0}^{n-1} p(i) = 1\}$ as the defender's randomized migration strategy where $p(0)$ is the probability that the defender keeps the resource at node 0 and $p(i)$ $(i = 1, 2, ..., n-1)$ is the probability that the defender moves the resource to node i. In the rest of paper, for simplicity of exposition, we say that the defender moves the resource to node 0 (even though there is no migration) if the defender keeps the resource at the original node 0.

Given $\mathbf{p}$, the attacker will either attack the node i whose $p(i) = \max \mathbf{p}$ or not attack. Let π denotes the attacker's binary strategy, where

$$\pi = \begin{cases} 1 & \text{attack,} \\ 0 & \text{not attack.} \end{cases}$$

The attacker's expected payoff is then defined as a function of $\mathbf{p}$ and π as follows:

$$\mathcal{U}(\mathbf{p}, \pi) = (\max \mathbf{p} - c_a) \cdot \pi \tag{1}$$

where $c_a \geq 0$ is the attack cost, which is assumed to be a constant that is known to the defender. Since $\max \mathbf{p} \leq 1$, it is reasonable to assume that $c_a < 1$.

On the other hand, the defender's cost is captured by the following linear quadric form:

$$\mathcal{C}(\mathbf{p}, \pi) = (\max \mathbf{p} \cdot \pi) + c_d \left(1 - p(0)\right)^2 \tag{2}$$

where $c_d \geq 0$ is the unit migration cost, which is a constant known to the defender. The first term in (2) depicts the expected loss from attack, which is related to both $\mathbf{p}$ and π. The second term depicts the migration cost, which is related to $\mathcal{P}$ only. The quadratic form captures the resource requirement of moving the resource, *e.g.*, network bandwidth and energy, which is often nonlinear. In fact, we have also studied a linear migration cost model, *e.g.*, $\mathcal{C}(\mathbf{p}, \pi) = (\max \mathbf{p} \cdot \pi) + c_d \left(1 - p(0)\right)$. We have proved that if the defender has a linear migration cost function, signaling strategies cannot improve the defender's cost.

Due to randomized of migration strategy, the attacker is uncertain about the real destination. This fact can be further utilized by the defender. In particular, in the beginning of the game, in addition to $\mathbf{p}$, the defender could also commit to a signalling scheme Σ (to be defined) and declare it to the attacker. After the real migration destination is sampled from $\mathbf{p}$, the defender can then send a signal according to Σ conditioning on the sampled destination. Assuming the attacker is rational, it can then use the Bayes' rule to get a posterior belief and make a move accordingly. The attacker might also ignore the signal and make decision upon the prior belief only.

More specifically, we consider the following form of signaling adopted from Bayesian Persuation. Let $\sigma(m|i)$ denotes the conditional probability that the defender tells attacker it will move to m when its true destination is i. The defender's signaling scheme is then defined by the following matrix:

$$\Sigma = \begin{pmatrix} \sigma(0|0) & \sigma(1|0) & \cdots & \sigma(n-1|0) \\ \sigma(0|1) & \sigma(1|1) & \cdots & \sigma(n-1|1) \\ \vdots & \vdots & \ddots & \vdots \\ \sigma(0|n-1) & \sigma(1|n-1) & \cdots & \sigma(n-1|n-1) \end{pmatrix}$$

Thus, $(\mathbf{p}, \Sigma)$ form a pair of joint migration and signaling strategy. Given $(\mathbf{p}, \Sigma)$, the probability that the defender sends a signal indicating movement to node m can be determined as:

$$q_m = \sum_{i=0}^{n-1} p(i)\sigma(m|i) \tag{3}$$

When a signal "m" is received, the attacker uses Bayes' rule to obtain a posterior belief $q(i|m) = \frac{p(i)\sigma(i|m)}{q_m}$, i.e., the probability that the defender migrates the resource to node i conditioned on the signal "m" received. The following matrix summarizes the attacker's posterior belief for each signal.

$$\mathcal{Q} = \begin{pmatrix} q(0|0) & q(1|0) & \cdots & q(n-1|0) \\ q(0|1) & q(1|1) & \cdots & q(n-1|1) \\ \vdots & \vdots & \ddots & \vdots \\ q(0|n-1) & q(1|n-1) & \cdots & q(n-1|n-1) \end{pmatrix}$$

When a signaling scheme is adopted, the attacker can base its attack decision on the signal received. Define the attacker's strategy against $(\mathbf{p}, \Sigma)$ as $\Pi = \{\pi(m)|\pi(m) \in \{0, 1\}, m = 0, 1, ..., n-1\}$, where

$$\pi(m) = \begin{cases} 1 & \text{attack node } i \text{ whose } q(i|m) = \max \mathcal{Q}(m,:), \\ 0 & \text{not attack.} \end{cases}$$

where $\mathcal{Q}(m,:)$ denotes the elements in the $(m+1)$-th row of $\mathcal{Q}$.

The attacker's payoff when receiving signal "m" is:

$$\mathcal{U}\left(\mathcal{Q}, \pi(m)\right) = \{\max \mathcal{Q}(m,:) - c_a\} \cdot \pi(m) \tag{4}$$

Then the attacker's expected payoff is:

$$\begin{aligned} \widetilde{\mathcal{U}}\left((\mathbf{p}, \Sigma), \Pi\right) &= \sum_{m=0}^{n-1} q_m \mathcal{U}\left(\max \mathcal{Q}(m,:), \pi(m)\right) \\ &= \sum_{m=0}^{n-1} \sum_{i=0}^{n-1} p(i)\sigma(m|i)\mathcal{U}\left(\mathcal{Q}, \pi(m)\right) \end{aligned} \tag{5}$$

The defender's expected payoff as a function of the attacker's action is:

$$\begin{aligned} \widetilde{\mathcal{C}}\left((\mathbf{p}, \Sigma), \Pi\right) &= \sum_{m=0}^{n-1} q_m \max \mathcal{Q}(m,:)\pi(m) + c_d \left(1 - p(0)\right)^2 \\ &= \sum_{m=0}^{n-1} \sum_{i=0}^{n-1} p(i)\sigma(m|i) \max \mathcal{Q}(m,:)\pi(m) \\ &\quad + c_d \left(1 - p(0)\right)^2 \end{aligned} \tag{6}$$

As is common in the security game literature, we assume that when multiple strategies give the attacker the same payoff, the tie is broken in the favor of the defender.

Table I summarizes the notation used in this paper.

Table I: List of Parameters

Symbol	Meaning
n	number of nodes
$p(i)$	probability of moving the resource to node i
$\mathbf{p}$	defender's migration strategy, $\mathbf{p} = \{p(i)\}$
$\sigma(m\|i)$	probability of sending "m" conditional on migrating the resource to node i
Σ	defender's signal pattern, $\Sigma = \{\sigma(m\|i)\}$
q_m	probability of the defender sending "m"
$q(i\|m)$	probability of the defender moving the resource to node i conditional on sending "m"
$\mathcal{Q}$	attacker's posterior believe, $\mathcal{Q} = \{q(i\|m)\}$
π	attacker's strategy when the defender doesn't send signals
$\pi(m)$	attacker's action on receiving signal "m"
Π	attacker's strategy $\Pi = \{\pi(m)\}$
c_d	migration cost
c_a	attacking cost

B. Discussion of the Model

We consider a simplified model for moving target defense. In particular, we consider a single stage game and assume a homogeneous migration cost. This setting serves as a good starting point as it helps us to obtain an explicit form of the joint migration and signaling schemes and makes it easy to identify the impact of signaling in the context of MTD.

We consider a Stackelberg game with the defender as the leader and the attacker as the follower. Stackelberg games have been extensively used for modeling cyber-security scenarios as they naturally capture the fact that a targeted attacker may first observe the defender's action and then make a move [25]. As the leader of the game, the optimality of the defender's strategy relies on the assumptions that (1) the attacker correctly identifies the randomized strategy of the defender; and (2) the attacker is rational and will respond to the defender's strategy as expected. In our setting, it further means that the signals from the defender should be correctly received and responded. As we show in Sections IV and V, the optimal migration strategy boils down to determining the value of $p(0)$ while the optimal signaling matrix Σ contains at most four different values. Thus, our strategy can be communicated to the attacker at very low cost.

C. A Motivating Example

Before we present the results for the general case, we give a small example in this section to illustrate how a signaling strategy can improve the defender's performance.

Consider a network with three physical nodes 0, 1 and 2, where the resource is installed at node 0 at present. We compare three migration strategies of the defender with the system parameters set as $c_d = \frac{3}{2}$ and $c_a = \frac{1}{3}$:

1) $\mathbf{p} = \{1, 0, 0\}$.
   $\mathcal{C} = 1 + 0 = 1, \mathcal{U} = 1 - \frac{1}{3} = \frac{2}{3}$;
2) $\mathbf{p} = \{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\}$.
   $\mathcal{C} = 0 + (1 - \frac{1}{3})^2 \times \frac{3}{2} = \frac{2}{3}, \mathcal{U} = 0$ (not attack);
3) $\mathbf{p} = \{\frac{1}{2}, \frac{1}{4}, \frac{1}{4}\}$.

$\mathcal{C} = \frac{1}{2} + (1-\frac{1}{2})^2 \times \frac{3}{2} = \frac{7}{8}, \mathcal{U} = \frac{1}{2} - \frac{1}{3} = \frac{1}{6}$.

Obviously, the defender's minimum cost is $\frac{2}{3}$ when she adopts a strategy $\mathcal{P} = \{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\}$. However, the defender could improve her performance by announcing her moving destination! It can be proved that there is a unique optimal joint strategy to the defender, where $\mathbf{p} = \{\frac{1}{2}, \frac{1}{4}, \frac{1}{4}\}$ and

$$\Sigma = \begin{pmatrix} \frac{1}{2} & \frac{1}{4} & \frac{1}{4} \\ 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & \frac{1}{2} & \frac{1}{2} \end{pmatrix}$$

Under this strategy, the probability of sending signal "0", "1" and "2" are: $q_0 = \frac{1}{4}$, $q_1 = q_2 = \frac{3}{8}$, respectively. In this case, the attacker's posterior belief $\mathcal{Q}$ is:

$$\mathcal{Q} = \begin{pmatrix} 1 & 0 & 0 \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{pmatrix}$$

Thus, the attacker will attack node 0 when he receives "0" with probability 1. The attacker will not attack when he receives "1" or "2" since $\max \mathcal{Q}(1,:) = \max \mathcal{Q}(2,:) = \frac{1}{3} \leq c_a$. Thus, the attacker's payoff is: $\mathcal{U} = \frac{1}{4} \times (1-\frac{1}{3}) = \frac{1}{6}$, while the defender's cost is reduced to: $\mathcal{C} = \frac{1}{4} + (1-\frac{1}{2})^2 \times \frac{3}{2} = \frac{5}{8} < \frac{2}{3}$. Notice, the attacker will accept the signals since his payoff is better than if he ignores the signals. In the rest of the paper, we will study the general problem of how the defender jointly designs the migration and signaling strategy.

IV. Optimal Pure Migration Strategy

To begin with, we analyze the case when the defender only adopts a migration strategy in this section. The analysis will be extended to the joint migration and signaling setting in the next section.

A. Subgame Perfect Equilibrium $(\mathbf{p}^*, \pi^*)$

In this subsection, we characterize the subgame perfect equilibrium of the two-player Bayesian Stackelberg game when the defender adopts $\mathbf{p}$.

Definition 1. *A pair of strategies $(\mathbf{p}^*, \pi^*)$ form a subgame perfect equilibrium if*

- *The defender's cost $\mathcal{C}(\mathbf{p}, \pi)$ is optimized at $\mathbf{p} = \mathbf{p}^*$ over every possible responses from the attacker, and*
- *Given $\mathbf{p}^*$, the attacker's payoff $\mathcal{U}(\mathbf{p}^*, \pi)$ is optimized at $\pi = \pi^*$.*

We solve the game via a backward induction. First, according to (1), the attacker's best strategy can be found straightforwardly as a function of $\mathbf{p}$:

$$\pi^* = \begin{cases} 1 & \text{if } c_a < \max \mathbf{p}, \\ 0 & \text{if } c_a \geq \max \mathbf{p}. \end{cases}$$

Therefore, the attacker's action is only related to the maximum element in the defender's strategy $\mathbf{p}$. The defender's cost function in (2) can then be expressed as the function of $p(0)$ and $\max \mathbf{p}$ as follows:

$$\mathcal{C}\left(p(0), \max \mathbf{p}\right) = \begin{cases} \max \mathbf{p} + c_d \left(1 - p(0)\right)^2 & \max \mathbf{p} > c_a, \\ c_d \left(1 - p(0)\right)^2 & \max \mathbf{p} \leq c_a. \end{cases}$$

The following lemma gives a necessary condition for a strategy $\mathbf{p}$ to be optimal. (See proof in the technical report [26].)

Lemma 1. *If $\mathbf{p}$ is an optimal migration strategy, $p(0)$ is the maximum element in $\mathbf{p}$.*

Lemma 1 also implies that if $\mathbf{p}$ is the defender's optimal migration strategy, then $p(0) \geq \frac{1}{n}$. It follows that the defender's cost function (2) can be written as a function of $p(0)$ and π:

$$\mathcal{C}\left(p(0), \pi\right) = p(0) \cdot \pi + c_d \left(1 - p(0)\right)^2 \tag{7}$$

where,

$$\pi = \begin{cases} 1 & \text{if } p(0) > c_a, \\ 0 & \text{if } p(0) \leq c_a. \end{cases} \tag{8}$$

Moreover, if the defender's cost is optimized at some $\mathbf{p}$ with $p(0) = p^*(0)$, then it is also optimized for any $\mathbf{p}'$ with $p'(0) = p^*(0), p'(i) \leq p'(0), \sum_{i=1}^{n-1} p'(i) = 1 - p'(0), i = 1, 2, ..., n-1$. Thus, the problem of optimizing $\mathbf{p}$ can be simplified to optimize $p(0)$. Also, since the defender should disclose her strategy, $p(0)$ can be communicated at very low cost compared to communicating $\mathbf{p}$.

Therefore, the defender's strategy can be written as the following optimization problem with variable $p(0)$.

$$\begin{aligned} \min \quad & \mathcal{C}\left(p(0), \pi\right) \\ \text{s.t.} \quad & p(0) \in [\frac{1}{n}, 1]. \end{aligned} \tag{9}$$

The theorem below fully characterizes the set of subgame perfect equilibria $(p^*(0), \pi^*)$ of the migration game. (See proof in the technical report [26].)

Theorem 1. *Given c_a, c_d and n,*

1) *If $c_a < \frac{1}{n}$ and $c_d \leq \frac{n}{2(n-1)}$, $(p^*(0), \pi^*) = (\frac{1}{n}, 1)$*
   The equilibrium payoffs are:
   $\mathcal{C}^* = \frac{1}{n} + c_d \left(1 - \frac{1}{n}\right)^2$, $\mathcal{U}^* = \frac{1}{n} - c_a$
2) *If $c_a \geq \frac{1}{n}$ and $c_d \leq \frac{1+\sqrt{2c_a - c_a^2}}{2(1-c_a)^2}$, $(p^*(0), \pi^*) = (c_a, 0)$*
   The equilibrium payoffs are:
   $\mathcal{C}^* = c_d \left(1 - c_a\right)^2$, $\mathcal{U}^* = 0$
3) *Otherwise, $(p^*(0), \pi^*) = (1 - \frac{1}{2c_d}, 1)$*
   The equilibrium payoffs are:
   $\mathcal{C}^* = 1 - \frac{1}{4c_d}$, $\mathcal{U}^* = 1 - \frac{1}{2c_d} - c_a$

B. Discussion on $(p^*(0), \pi^*)$

Figure 2 provides a graphical illustration of the set of equilibria in Theorem 1. c_d and c_a form Quadrant I of a two-dimensional Cartesian system. Blue dash lines and $c_a = 1$ separate Quadrant I into three areas which correspond to the three equilibria in Theorem 1. We have the following insights from this figure:

1) By increasing the number of nodes in the network, the defender could almost avoid the equilibrium $(\frac{1}{n}, 1)$ which leads to the highest cost for the defender compared to any other equilibria with the same c_d.

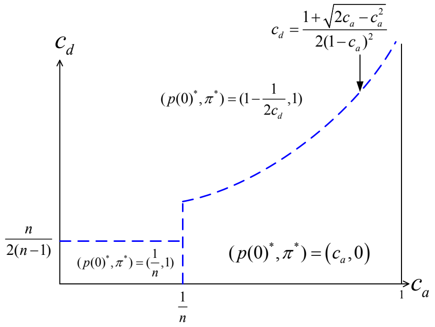

Figure 2: Illustration of Theorem 1

2) For a fixed n, the defender always benefits from reducing c_d since both $(\frac{1}{n}, 1)$ and $(c_a, 0)$ ensure lower cost with the same c_a.
3) Even with an extremely large c_d and a tiny c_a, *e.g.*, a poorly connected network and a very powerful attacker, there always exists a pure migration strategy that incurs less cost than no migration ($p(0) = 1$).

V. Optimal Joint Migration and Signaling Strategy

In this section, we investigate the joint migration and signaling strategy in our game. We would like to answer three questions in this section: 1) What is the optimal joint strategy? 2) When does a joint strategy improve a pure migration strategy? 3) How much improvement can a joint strategy achieve?

A. Subgame Perfect Equilibrium $((\mathbf{p}^\circ, \Sigma^\circ), \Pi^\circ)$

We characterize the subgame perfect equilibrium of the two-player game when the defender adopts $(\mathbf{p}, \Sigma)$.

Definition 2. *A group of strategies $((\mathbf{p}^\circ, \Sigma^\circ), \Pi^\circ)$ form a subgame perfect equilibrium if*

- *The defender's expected payoff $\widetilde{\mathcal{C}}((\mathbf{p}, \Sigma), \Pi)$ is optimized at $(\mathbf{p}, \Sigma) = (\mathbf{p}^\circ, \Sigma^\circ)$ over every possible responses from the attacker, and*
- *Given $(\mathbf{p}^\circ, \Sigma^\circ)$, for each signal "$m$" realized by $(\mathbf{p}^\circ, \Sigma^\circ)$, the attacker's expected payoff $\widetilde{\mathcal{U}}((\mathbf{p}^\circ, \Sigma^\circ), \Pi)$ is optimized at $\Pi = \Pi^\circ$.*

According to (5), the attacker's best strategy is straightforward when given a signal "m" and $(\mathbf{p}, \Sigma)$ as follows,

$$\pi^\circ(m) = \begin{cases} 1 & \text{if } c_a < \max \mathcal{Q}(m, :), \\ 0 & \text{if } c_a \geq \max \mathcal{Q}(m, :). \end{cases}$$

The defender's problem is to jointly design an optimal $(\mathbf{p}^\circ, \Sigma^\circ) = \text{argmin}_{\mathbf{p}, \Sigma} \widetilde{\mathcal{C}}((\mathbf{p}, \Sigma), \Pi)$. It can be written as the following optimization problem with variables $\{p(i), \sigma(m|i) : i, m = 0, 1, 2, ...n-1\}$.

$$\begin{aligned} \min \quad & \widetilde{\mathcal{C}}((\mathbf{p}, \Sigma), \Pi) \\ \text{s.t.} \quad & \sum_{m=0}^{n-1} \sigma(m|i) = 1, && \text{for } i = 0, 1, ..., n-1. \\ & \sigma(m|i) \geq 0, && \text{for } i, m = 0, 1, ..., n-1. \\ & p(i) \geq 0, && \text{for } i = 0, 1, ..., n-1. \\ & \widetilde{\mathcal{U}}((\mathbf{p}, \Sigma), \Pi) \geq \mathcal{U}(\mathbf{p}, \pi^*(\mathbf{p})). \end{aligned} \tag{10}$$

The last condition implies that the attacker accepts the signaling strategy if and only if his expected payoff is greater than the payoff without signals under the same prior belief $\mathbf{p}$.

The defender should decide to adopt either an optimal migration strategy $\mathbf{p}^*$ or an optimal joint migration and signaling strategy $(\mathbf{p}^\circ, \Sigma^\circ)$. Since we have already solved $\mathbf{p}^*$ in Section IV, it is sufficient to find those $(\mathbf{p}^\circ, \Sigma^\circ)$ that satisfy $\widetilde{\mathcal{C}}((\mathbf{p}^\circ, \Sigma^\circ), \Pi^\circ) < \mathcal{C}(\mathbf{p}^*, \pi^*)$. We will first show in what conditions there exists such $(\mathbf{p}^\circ, \Sigma^\circ)$, and then solve the explicit form of the $(\mathbf{p}^\circ, \Sigma^\circ)$. (See proof in the technical report [26].)

Lemma 2. *If $c_a < \frac{1}{n}$ or $c_a \geq 1 - \sqrt{\frac{1}{2c_d}}$, for any joint strategy $(\mathbf{p}, \Sigma)$, $\widetilde{\mathcal{C}}((\mathbf{p}, \Sigma), \Pi) \geq \mathcal{C}(\mathbf{p}^*, \pi^*)$.*

From Lemma 2, it is sufficient to solve (10) under the condition $\frac{1}{n} \leq c_a < 1 - \sqrt{\frac{1}{2c_d}}$. The result is summarized in the following theorem, which is the main result of the paper.

Theorem 2. *Given c_a, c_d and n, if $\frac{1}{n} \leq c_a < 1 - \sqrt{\frac{1}{2c_d}}$, $\widetilde{\mathcal{C}}((\mathbf{p}, \Sigma), \Pi)$ is optimized at $(\mathbf{p}^\circ, \Sigma^\circ)$, where:*

$$\mathbf{p}^\circ = \{1 - \frac{1}{2c_d(1-c_a)}, \underbrace{\frac{1}{2c_d(1-c_a)(n-1)}, ..., \frac{1}{2c_d(1-c_a)(n-1)}}_{n-1}\} \tag{11}$$

$$\Sigma^\circ = \begin{pmatrix} \frac{p_0 - c_a}{p_0(1-c_a)} & \frac{c_a(1-p(0))}{p(0)(1-c_a)(n-1)} & \cdots & \frac{c_a(1-p(0))}{p(0)(1-c_a)(n-1)} \\ 0 & \frac{1}{n-1} & \cdots & \frac{1}{n-1} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \frac{1}{n-1} & \cdots & \frac{1}{n-1} \end{pmatrix} \tag{12}$$

$\widetilde{\mathcal{U}}((\mathbf{p}^\circ, \Sigma^\circ), \Pi)$ is optimized at $\Pi = \Pi^\circ = \{1, \underbrace{0, ..., 0}_{n-1}\}$. $((\mathbf{p}^\circ, \Sigma^\circ), \Pi^\circ)$ form a subgame perfect equilibrium.

Proof. As the full proof is long, we only provide a proof sketch here. We prove this theorem in five steps:

1) If $\widetilde{\mathcal{C}}((\mathbf{p}, \Sigma), \Pi)$ is optimized at $(\mathbf{p}^\circ, \Sigma^\circ)$, the attacker attacks if and only if he receives one certain signal. We prove this statement by recursion: if a joint strategy $((\mathbf{p}, \Sigma), \Pi)$ exists such that the attacker attacks when he receives k out of n signals, then there exists another strategy such that attacker attacks when he receives $k-1$ out of n signals, and the later strategy has a better performance than the former one for the defender.

2) If $p(i) = \text{argmax}\mathbf{p}^\circ$, then the attacker attacks if and only if he receives "i".
3) $p(0) = \max \mathbf{p}^\circ$. This is because changing order of the elements in $\mathbf{p}$ does not effect the first term in (6) (if we switch the $i-$th and $j-$th elements in $\mathbf{p}$, we should also switch $i-$th and $j-$th rows and $i-$th and $j-$th columns), however, if and only if $p(0) = \max \mathbf{p}$, the second term in (6) achieves the minimum value.
4) Solving $\mathbf{p}^\circ$ and Σ°. According to the three statements above, $\mathbf{p}^\circ = \{p(0), \underbrace{\frac{1-p(0)}{n-1}, ..., \frac{1-p(0)}{n-1}}_{n-1}\}$, $\sigma(0|i) = \frac{1-\sigma(0|0)}{n-1}$, $\sigma(i|0) = \frac{1-\sigma(0|0)}{n-1}$ and $\sigma(m|i) = \frac{1-\sigma(0|i)}{n-1}$, $m, i = 1, 2, ..., n-1$. The defender's expected payoff is:

$$\begin{aligned}\widetilde{\mathcal{C}}\left((\mathbf{p}, \Sigma), \Pi\right) =& \sum_{i=0}^{n-1} p(i)\sigma(0|i)q(0|0) + c_d\left(1 - p(0)\right)^2 \\ =& p(0)\sigma(0|0) + c_d\left(1 - p(0)\right)^2 \end{aligned} \tag{13}$$

Then the optimization problem in (10) can be simplified with only two variables $p(0)$ and $\sigma(0|0)$ as shown below,

$$\begin{array}{ll} \min & p(0)\sigma(0|0) + c_d\left(1 - p(0)\right)^2 \\ \text{s.t.} & p(0) > c_a, \\ & \sigma(0|0) > 0, \\ & p(0)\sigma(0|0)(1 - c_a) \geq p(0) - c_a. \end{array} \tag{14}$$

If and only if $\frac{1}{n} \leq c_a < 1 - \sqrt{\frac{1}{2c_d}}$, $p(0) = 1 - \frac{1}{2c_d(1-c_a)}$, $\sigma(0|0) = \frac{p_0 - c_a}{p_0(1-c_a)}$ are the unique solutions of the problem above. Otherwise, there is no solution for (14). This step also prove the statement in Lemma 2: if $c_a \geq 1 - \sqrt{\frac{1}{2c_d}}$, for any joint strategy $(\mathbf{p}, \Sigma)$, $\widetilde{\mathcal{C}}\left((\mathbf{p}, \Sigma), \Pi\right) \geq \mathcal{C}(\mathbf{p}^*, \pi^*)$.
5) $\Pi^\circ = \{1, \underbrace{0, ..., 0}_{n-1}\}$. Without signals, $p(0) = 1 - \frac{1}{2c_d(1-c_a)} = \max \mathbf{p}^\circ > c_a$, then $\pi^\circ = 1$, the attacker will attack node 0 with an expected payoff:

$$\mathcal{U}(\mathbf{p}^\circ, 1) = 1 - \frac{1}{2c_d(1 - c_a)} - c_a \tag{15}$$

With signals, the attacker forms his posterior belief $\mathcal{Q}$ as follows:

$$\mathcal{Q} = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ c_a & \frac{1-c_a}{n-1} & \cdots & \frac{1-c_a}{n-1} \\ \vdots & \vdots & \ddots & \vdots \\ c_a & \frac{1-c_a}{n-1} & \cdots & \frac{1-c_a}{n-1} \end{pmatrix}$$

$$\max \mathcal{Q}(m, :) = \begin{cases} 1, & \text{if } m = 0 \\ c_a, & \text{otherwise} \end{cases}$$

Thus, the attacker will attack node 0 if and only if he receives signal "0". Also, the attacker's payoff is:

$$\begin{aligned}\widetilde{\mathcal{U}}^\circ =& q_0\mathcal{U}\left(\mathcal{Q}(0, :), \pi^\circ(0)\right) \\ =& \left(1 - \frac{1}{2c_d(1 - c_a)^2}\right)(1 - c_a) \\ =& 1 - \frac{1}{2c_d(1 - c_a)} - c_a \end{aligned} \tag{16}$$

$\widetilde{\mathcal{U}}^\circ = \mathcal{U}(\mathbf{p}^\circ, 1)$, the attacker will accept the defender's signals. When the defender adopts $(\mathbf{p}^\circ, \Sigma^\circ)$,
This concludes the proof of Theorem 2.

□

From Theorem 2 we can observe that the optimal migration strategy $\mathbf{p}^\circ$ and signaling matrix Σ contain only two and four values, respectively. Thus, our strategy can be communicated to the attacker at very low cost.

The following corollary shows the corresponding equilibrium payoffs under $((\mathbf{p}^\circ, \Sigma^\circ), \Pi^\circ)$. It also indicates that under the condition $\frac{1}{n} \leq c_a < 1 - \sqrt{\frac{1}{2c_d}}$, the defender's cost is always reduced by the joint strategy. (See proof in the technical report [26].)

Corollary 1. *Given c_a, c_d and n, if $\frac{1}{n} \leq c_a < 1 - \sqrt{\frac{1}{2c_d}}$, the defender's equilibrium cost is:*

$$\begin{aligned}\widetilde{\mathcal{C}}^\circ = & c_d(1 - c_a)^2 - c_d\left[1 - c_a - \frac{1}{2c_d(1 - c_a)}\right]^2 \\ & < \mathcal{C}(\mathbf{p}^*, \pi^*) \end{aligned} \tag{17}$$

The attacker's equilibrium payoff is:

$$\widetilde{\mathcal{U}}^\circ = 1 - \frac{1}{2c_d(1 - c_a)} - c_a \tag{18}$$

According to Theorem 2 and Corollary 1, if and only if $\frac{1}{n} \leq c_a < 1 - \sqrt{\frac{1}{2c_d}}$, $\mathcal{C}\left((\mathbf{p}^\circ, \Sigma^\circ), \Pi^\circ\right) < \mathcal{C}(\mathbf{p}^*, \pi^*)$.

B. Discussion on Subgame Perfect Equilibria

In this subsection, we discuss the overall equilibrium strategies as studied in Section IV and Section V. According to Theorem 1 and Theorem 2, the subgame perfect equilibria when the defender adopts either a pure migration strategy or a joint strategy can be summarized as follows:

1) $((\mathbf{p}^\circ, \Sigma^\circ), \Pi^\circ)$, where $\mathbf{p}^\circ$, Σ° and Π° are defined in Theorem 2, if $\frac{1}{n} \leq c_a < 1 - \sqrt{\frac{1}{2c_d}}$.
2) $(p(0)^*, \pi^*) = (c_a, 0)$, if $c_a \geq \frac{1}{n}$ and $c_d \leq \frac{1}{2(1-c_a)^2}$.
3) $(p(0)^*, \pi^*) = (\frac{1}{n}, 1)$, if $c_a < \frac{1}{n}$ and $c_d \leq \frac{n}{2(n-1)}$.
4) $(p(0)^*, \pi^*) = (1 - \frac{1}{2c_d}, 1)$, if $c_a < \frac{1}{n}$ and $c_d \geq \frac{n}{2(n-1)}$.

Figure 3 provides a graphical illustration of the set of subgame perfect equilibria in both cases. c_a and c_d form Quadrant I of a two-dimensional Cartesian system, red dash lines and $c_a = 1$ separate Quadrant I into four areas which correspond to the four equilibria above. Comparing this to Figure 2, the area

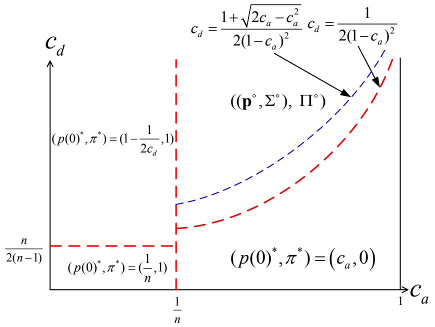
$c_d = \frac{1+\sqrt{2c_a - c_a^2}}{2(1-c_a)^2}$
$c_d = \frac{1}{2(1-c_a)^2}$
c_d
$((\mathbf{p}^\circ, \Sigma^\circ),\ \Pi^\circ)$
$(p(0)^*, \pi^*) = (1-\frac{1}{2c_d}, 1)$
$\frac{n}{2(n-1)}$
$(p(0)^*, \pi^*) = (\frac{1}{n}, 1)$
$(p(0)^*, \pi^*) = (c_a, 0)$
$\frac{1}{n}$
1
c_a

Figure 3: Illustration of subgame perfect equilibria when the defender adopts a joint migration and signaling strategy

where the joint strategy is beneficial consists of two parts: the area bounded by $c_d = \frac{1}{2(1-c_a)^2}$, $c_d = \frac{1+\sqrt{2c_a - c_a^2}}{2(1-c_a)^2}$ and $c_a = \frac{1}{n}$, and the area bounded by $c_d = \frac{1+\sqrt{2c_a - c_a^2}}{2(1-c_a)^2}$ and $c_a = \frac{1}{n}$.

From this figure, we have the following suggestions to the defender to ensure a more secure system:

1) The defender can benefit from the joint strategy in almost every system settings by increasing the number of nodes in the network. For those settings that do not satisfy $\frac{1}{n} \le c_a < 1 - \sqrt{\frac{1}{2c_d}}$, the defender could adopt $p^*(0) = c_a$ to avoid being attacked. This property implies that the defender can always obtain benefit from signaling in a large network.
2) The signaling strategy is beneficial even with an extremely large c_d. In fact, the signaling strategy has a relative better performance in the large c_d area.
3) For a fixed n, the defender cannot always benefit from the signaling strategy by increasing or reducing c_d. However, if the defender could increase c_a, *e.g.*, increasing attack complexity or reducing the chance of a successful exploit in a given amount of time, the defender could always adopt a useful joint strategy.

VI. Numerical Results

In this section, we examine our proposed model via numerical study under different system scenarios and configurations.

A. Defender Adopts a Pure Migration Strategy

We first present the equilibrium payoffs for both players when the defender adopts a migration strategy as studied in Theorem 1. The red and blue curves in Figure 4 represent the defender's equilibrium cost and the attacker's equilibrium payoff, respectively.

- Scenario 1: $c_a \ge \frac{1}{n}$, $c_a = 0.3$, $n = 4$, $c_d \in [0, 3]$. Figure 4a shows that a larger c_d is always harmful to the defender. However, the attacker is not always beneficial from a smaller c_d since the defender could prevent the resource from being attacked when $c_d \le \frac{1+\sqrt{2c_a - c_a^2}}{2(1-c_a)^2}$.
- Scenario 2: Defender and attacker have the same cost parameters, $c_a = c_d = 0.3$, $n \in [2, 8]$. From Figure 4b we find that the defender's equilibrium cost drops dramatically when c_a exceeds $\frac{1}{n}$. However, further increasing n cannot provide more benefit to the defender. Therefore, adding more nodes is not always beneficial to the defender.

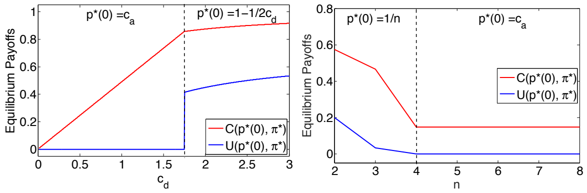

(a) $c_a = 0.3$, $n = 4$

(b) $c_a = c_d = 0.3$

Figure 4: Equilibrium Payoffs $\mathcal{C}^*$ and $\mathcal{U}^*$ v.s. c_d and n.

B. Defender Adopts a Joint Migration and Signaling Strategy

In this subsection, we compare the defender's equilibrium cost when the defender adopts a joint strategy as mentioned in Section V with the case when the defender adopts a pure migration strategy. The red and blue curves in Figure 5 represent the defender's equilibrium cost with joint strategy and pure migration strategy, respectively.

- Scenario 1: The defender has a low migration cost: *e.g.*, $c_d = 1$, $n = 4$, and $c_a \in [0, 1]$. In Figure 5a, the defender's equilibrium cost has been slightly reduced in a small range through joint strategy. Therefore, if employing the signaling strategy has a cost, *e.g.*, a cost on information disclosure or communication, the defender might adopt a pure migration strategy.
- Scenario 2: The defender has a high migration cost $c_d = 6$, $n = 4$, and $c_a \in [0, 1]$. The defender's cost has been reduced as much as 0.3 as shown in Figure 5b. Also, the joint strategy is efficient over a large range of c_a from 0.2 to 0.7. This implies that the joint strategy has better performance especially when the migration cost is high.
- Scenario 3: $c_a \ge \frac{1}{n}$, $n = 5$, $c_a = 0.3$, and $c_d \in [0, 10]$. From Figure 5c we can see that even with an extremely high migration cost, the joint strategy is always beneficial to the defender.
- Scenario 4: Fix $c_d = 1.2$. Vary n from 2 to 8, and set $c_a = \frac{1}{n}$. We consider the problem that if the attacker has a very small attacking cost, does the defender have an efficient strategy to fight against such an attacker? We setup a configuration that fixes the product of c_a and n so that we can reduce c_a by increasing n. The x-axis in Figure 5d represents both n varying from 2 to 8 and c_a varying from 0.5 to 0.125. With decreasing of attacking cost, the defender's equilibrium cost gets worse if she adopts a pure migration strategy. However, the defender could maintain a relatively stable cost through a joint strategy. Therefore, we conclude that no matter how small the attacking cost is, the defender could always construct an efficient defense system to keep a relatively stable cost in a large network.

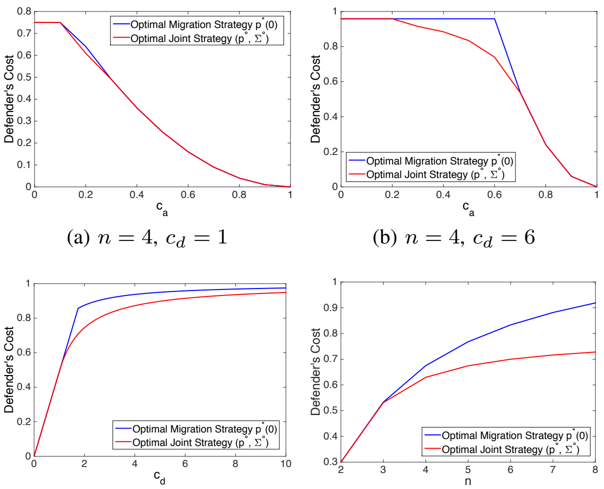

(a) $n = 4$, $c_d = 1$

(b) $n = 4$, $c_d = 6$

(c) $n = 5$, $c_a = 0.3$

(d) $c_d = 1.2$, $c_a = \frac{1}{n}$

Figure 5: Defender's Equilibrium Cost under Different System Settings

VII. Conclusion

We propose a two-player Bayesian Stackelberg game that models the joint migration and signaling strategies for the defender in the face of a strategic and rational attacker. By rigorous investigation, we show that MTD can be improved through strategic information disclosure. We fully characterize the subgame perfect equilibria of the game under a pure migration setting and when both migration and signaling are applied. Through theoretical analysis and numerical study of the proposed model, we have derived several insights and made suggestions for more efficient MTD.

VIII. Acknowledgement

The effort described in this article was partially sponsored by the U.S. Army Research Laboratory Cyber Security Collaborative Research Alliance under Contract Number W911NF-13-2-0045. The views and conclusions contained in this document are those of the authors, and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

[1] E. Cole, *Advanced persistent threat: understanding the danger and how to protect your organization*. Newnes, 2012.

[2] M. Carvalho and R. Ford, "Moving-target defenses for computer networks," *IEEE Security & Privacy*, vol. 2, no. 12, pp. 73–76, 2014.

[3] S. Jajodia, A. K. Ghosh, V. Subrahmanian, V. Swarup, C. Wang, and X. S. Wang, "Moving target defense ii," *Application of game Theory and Adversarial Modeling. Series: Advances in Information Security*, vol. 100, 2013.

[4] W. Peng, F. Li, C.-T. Huang, and X. Zou, "A moving-target defense strategy for cloud-based services with heterogeneous and dynamic attack surfaces," in *2014 IEEE International Conference on Communications (ICC)*. IEEE, 2014, pp. 804–809.

[5] A. D. Keromytis, R. Geambasu, S. Sethumadhavan, S. J. Stolfo, J. Yang, A. Benameur, M. Dacier, M. Elder, D. Kienzle, and A. Stavrou, "The meerkats cloud security architecture," in *32nd International Conference on Distributed Computing Systems Workshops*. IEEE, 2012, pp. 446–450.

[6] S. G. Vadlamudi, S. Sengupta, S. Kambhampati, M. Taguinod, Z. Zhao, A. Doupé, and G.-J. Ahn, "Moving target defense for web applications using bayesian stackelberg games," *arXiv preprint arXiv:1602.07024*, 2016.

[7] M. Taguinod, A. Doupé, Z. Zhao, and G.-J. Ahn, "Toward a moving target defense for web applications," in *Information Reuse and Integration (IRI), 2015 IEEE International Conference on*. IEEE, 2015, pp. 510–517.

[8] Q. Zhu and T. Başar, "Game-theoretic approach to feedback-driven multi-stage moving target defense," in *Proc. of GameSec*, 2013, pp. 246–263.

[9] P. K. Manadhata, "Game theoretic approaches to attack surface shifting," in *Moving Target Defense II*. Springer, 2013, pp. 1–13.

[10] E. Al-Shaer, "Toward network configuration randomization for moving target defense," in *Moving Target Defense*. Springer, 2011, pp. 153–159.

[11] E. Kamenica and M. Gentzkow, "Bayesian persuasion," *The American Economic Review*, vol. 101, no. 6, pp. 2590–2615, 2011.

[12] H. Okhravi, T. Hobson, D. Bigelow, and W. Streilein, "Finding focus in the blur of moving-target techniques," *IEEE Security & Privacy*, vol. 12, no. 2, pp. 16–26, 2014.

[13] J. H. Jafarian, E. Al-Shaer, and Q. Duan, "Openflow random host mutation: transparent moving target defense using software defined networking," in *Proceedings of the first workshop on Hot topics in software defined networks*. ACM, 2012, pp. 127–132.

[14] B. Salamat, T. Jackson, G. Wagner, C. Wimmer, and M. Franz, "Runtime defense against code injection attacks using replicated execution," *IEEE Transactions on Dependable and Secure Computing*, vol. 8, no. 4, pp. 588–601, 2011.

[15] K. Onarlioglu, L. Bilge, A. Lanzi, D. Balzarotti, and E. Kirda, "G-free: defeating return-oriented programming through gadget-less binaries," in *Proceedings of the 26th Annual Computer Security Applications Conference*. ACM, 2010, pp. 49–58.

[16] A. Nguyen-Tuong, D. Evans, J. C. Knight, B. Cox, and J. W. Davidson, "Security through redundant data diversity," in *2008 IEEE International Conference on Dependable Systems and Networks With FTCS and DCC (DSN)*. IEEE, 2008, pp. 187–196.

[17] M. Van Dijk, A. Juels, A. Oprea, and R. L. Rivest, "Flipit: The game of "stealthy takeover"," *Journal of Cryptology*, vol. 26, no. 4, pp. 655–713, 2013.

[18] A. Prakash and M. P. Wellman, "Empirical game-theoretic analysis for moving target defense," in *Proceedings of the Second ACM Workshop on Moving Target Defense*, 2015, pp. 57–65.

[19] K. M. Carter, J. F. Riordan, and H. Okhravi, "A game theoretic approach to strategy determination for dynamic platform defenses," in *Proceedings of the First ACM Workshop on Moving Target Defense*. ACM, 2014, pp. 21–30.

[20] L. Rayo and I. Segal, "Optimal information disclosure," *Journal of Political Economy*, vol. 118, no. 5, pp. 949–987, 2010.

[21] S. Dughmi and H. Xu, "Algorithmic bayesian persuasion," in *In Proceedings of the Forty-eighth Annual ACM Symposium on Theory of Computing, STOC 16*, 2016.

[22] J. Pawlick and Q. Zhu, "Deception by design: evidence-based signaling games for network defense," *arXiv preprint arXiv:1503.05458*, 2015.

[23] T. E. Carroll and D. Grosu, "A game theoretic investigation of deception in network security," *Security and Communication Networks*, vol. 4, no. 10, pp. 1162–1172, 2011.

[24] H. Xu, R. Freeman, V. Conitzer, S. Dughmi, and M. Tambe, "Signaling in bayesian stackelberg games," in *Proc. of AAMAS*, 2016, pp. 150–158.

[25] Z. Yin, D. Korzhyk, C. Kiekintveld, V. Conitzer, and M. Tambe, "Stackelberg vs. nash in security games: Interchangeability, equivalence, and uniqueness," in *Proc. of AAMAS*, 2010, pp. 1139–1146.

[26] X. Feng, Z. Zheng, P. Mohapatra, D. Cansever, and A. Swami, "A Signaling Game Model for Moving Target Defense," Technical Report, available online at http://spirit.cs.ucdavis.edu/pubs/tr/fenginfo17.pdf.