FlipLeakage: A Game-Theoretic Approach to Protect Against Stealthy Attackers in the Presence of Information Leakage

Sadegh Farhang(✉) and Jens Grossklags

College of Information Sciences and Technology,
The Pennsylvania State University, University Park, PA, USA
{farhang,jensg}@ist.psu.edu

Abstract. One of the particularly daunting issues in the cybersecurity domain is *information leakage* of business or consumer data, which is often triggered by multi-stage attacks and advanced persistent threats. While the technical community is working on improved system designs to prevent and mitigate such attacks, a significant residual risk remains that attacks succeed and may not even be detected, i.e., they are *stealthy*.

Our objective is to inform security policy design for the mitigation of stealthy information leakage attacks. Such a policy mechanism advises system owners on the optimal timing to reset defense mechanisms, e.g., changing cryptographic keys or passwords, reinstalling systems, installing new patches, or reassigning security staff.

We follow a game-theoretic approach and propose a model titled *FlipLeakage*. In our proposed model, an attacker will incrementally and stealthily take ownership of a resource (e.g., similar to advanced persistent threats). While her final objective is a complete compromise of the system, she may derive some utility during the preliminary phases of the attack. The defender can take a costly recovery move and has to decide on its optimal timing.

Our focus is on the scenario when the defender can only partially eliminate the foothold of the attacker in the system. Further, the defender cannot undo any information leakage that has already taken place during an attack. We derive optimal strategies for the agents in FlipLeakage and present numerical analyses and graphical visualizations.

1 Introduction

Security compromises which cause *information leakage* of business or consumer data are a particularly challenging problem in the cybersecurity domain. Affected businesses frequently struggle to recover once the consequences of a breach become apparent such as a competitor outpacing them in a race for the next innovation, or data troves appearing on cybercriminal marketplaces and eventually impacting consumer confidence. For example, data about small and medium-sized businesses suggests that approximately 60 % fail within six months after a data breach [3].

Q. Zhu et al. (Eds.): GameSec 2016, LNCS 9996, pp. 195–214, 2016.
DOI: 10.1007/978-3-319-47413-7_12

Businesses struggle for multiple reasons to prevent information leakage. In particular, the increasing prevalence of well-motivated, technically capable, and well-funded attackers who are able to execute sophisticated multi-stage attacks and advanced persistent threats (APT) poses significant challenges to prevent information leakage. Such attacks may take time to execute, but they will eventually succeed with high likelihood. In a recent talk, the Chief of Tailored Access Operations, National Security Agency, characterized the mindset of these attackers in the following way: "We are going to be persistent. We are going to keep coming, and coming, and coming [12]."

Further, carefully orchestrated attacks as employed during corporate, cyber-criminal or nation-state sponsored cyber-espionage and sabotage (see Stuxnet [4]) change our understanding of the likelihood to reliably detect stealthy attacks before it is too late. Estimates for how long attacks remain undetected are dire. For example, a recent presentation by the CEO of Microsoft suggested that the time until detection of a successful attack is on average over 200 days [21].

All of these observations emphasize the need to reason about the suitable response to stealthy attacks which cause continued information leakage. We know that perfect security is too costly; and even air-gaped systems are vulnerable to insider risks or creative technical approaches. Another mitigation approach is to limit the impact of attacks by resetting system resources to a presumed safe state to lower the chances of a perpetual undetected leak. However, in most scenarios such actions will be costly. For example, they may impact productivity due to system downtime or the need to reissue cryptographic keys, passwords or other security credentials. As such, determining the best schedule to reset defense mechanisms is an economic question which needs to account for monetary and productivity costs, strategic and stealthy attacker behavior, and other important facets of information leakage scenarios such as the effectiveness of the system reset. To address this combination of factors, we propose a new game-theoretic model called *FlipLeakage*.

In our proposed model, an attacker has to engage in a sustained attack effort to compromise the security of a system. Our approach is consistent with two scenarios. On the one hand, the attacker may conduct surveillance of the system to collect information that will enable a security compromise, e.g., by pilfering traffic for valuable information, or by gathering information about the system setup. On the other hand, the attacker may incrementally take over parts of a system, such as user accounts, parts of a cryptographic key, or collect business secrets to enable further attack steps. In both scenarios, persistent activity and the accumulated information will then enable the attacker to reach her objective to compromise the system and to acquire the primary business secret; if the defender does not interfere by returning the system to a presumed safe state.

In Fig. 1, we provide an initial abstract representation of the studied strategic interaction between an attacker and a defender. The attacker initiates sustained attack efforts at t_1, t_2, and t_3 right after the defender's moves, where each time she also starts gaining information about the system. After accumulating sufficient information about the system, the attacker will be able to compromise it. The attacker's benefit until the security compromise is completed is represented as a triangle, which represents the value of the leaked information during the attack execution. After the compromise, the attacker continues to receive benefits from the compromised system which is represented as a rectangle.

The defender can take a recovery action (to reset the resource to a presumed safe state) and can thereby stop the attack. In our model, we consider the scenario when the defender only partially eliminates the foothold of the attacker in the system. In Fig. 1 those defensive moves occur at t_1, t_2, and t_3. Further, the defender cannot undo any information leakage that has already taken place during an attack.

In our model, we focus on periodic defensive moves for the defender. That means the time between any two consecutive moves is assumed the same motivated by practical observations for security policy updates of major software vendors such as Microsoft and Oracle which we will discuss in detail in Sect. 3. Within this context, we aim to determine the defender's best periodic defensive strategies when the moves of the attacker are unobservable to the defender, i.e., the attacks succeed to be stealthy. At the same time, we assume that the attacker can observe the defender's moves. The latter assumption rests on two observations. On the one hand, the attacker will be cut off from access to a partially compromised system when a recovery move takes place. On the other hand, many defensive moves may actually be practically observable for attackers, e.g., when a patch for a software system becomes available which makes a particular attack strategy impractical. The scenario under investigation is a security game of timing, e.g., we are studying *when* players should move to act optimally.

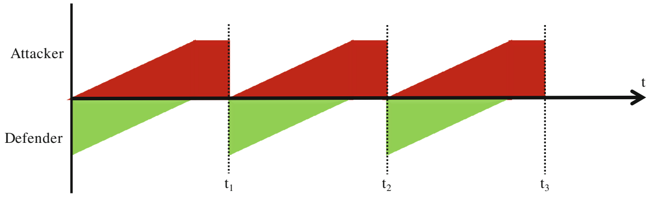

Fig. 1. FlipLeakage is a two-player game between an attacker and a defender competing with each other to control a resource. t_1, t_2, and t_3 represent the defender's move times. During the time when the attacker launches her attack, she incrementally benefits from information leakage which is shown as red triangles. (Color figure online)

In the following, we provide a brief summary overview over our contributions.

- We develop a game-theoretic model titled *FlipLeakage*. In our model, an attacker will incrementally take ownership of a resource (e.g., similar to advanced persistent threats). While her final objective is a complete compromise of the system, she may derive some utility during the preliminary phases of the attack. The defender can take a costly periodic mitigation move and has to decide on its optimal periodic timing.

- We consider the scenario when the defender only partially eliminates the foothold of the attacker in the system. Further, the defender cannot undo any information leakage that has already taken place during an attack.
- We derive optimal strategies for the agents in our model and present numerical analyses and graphical visualizations. One of our findings corroborates an intuition: the higher the defensive cost, the slower the defender's periodic move rhythm. Moreover, our numerical observations imply that the defender moves faster when the attacker's average time to totally compromise the defender's system is lower.

In the presence of stealthy attacks and information leakage, defenders have to set a schedule for updating and resetting their defense mechanisms without any feedback about the occurrence of attacks. This poses significant challenges for the design of new methods to mitigate such attacks. The objective of our theoretical model is to provide a systematic approach for the defender's best schedule to reset his system to a presumed safe state to lower the chances of a perpetually undetected leak. As such, our work provides important steps towards building a rigorous model for an optimal defender's response to these unknowns.

Roadmap: The rest of our paper is organized as follows. We discuss related work in Sect. 2. In Sect. 3, we develop the FlipLeakage model followed by payoff calculations in Sect. 4. We analyze our proposed model in Sect. 5. In Sect. 6, we present numerical examples. Finally, we conclude our paper in Sect. 7.

2 Related Work

Game theory is widely used in cybersecurity and privacy scenarios to study interdependencies [7,10,13,27], and dynamic interactions between defenders and attackers of varying complexity [5,17,19]. One recently emphasized aspect of security games is the consideration of *when* to act to successfully mitigate attacks. In particular, the issue of optimally timing defensive actions to successfully thwart stealthy attacks has attracted attention in the cybersecurity domain with the introduction of the *FlipIt* game [2,29] which broadens the *games of timing* literature initiated in the cold-war era [1,28]. In what follows, we provide a brief description of the FlipIt game as well as theoretical follow-up research.

FlipIt is a two-player game between a defender and an attacker competing with each other to control a resource which generates a payoff to the owner of the resource. Moves to take over the resource, i.e., *flips*, are costly [2,29]. In [29], the authors studied the FlipIt game with different choices of strategy profiles and aimed to calculate dominant strategies and Nash equilibria of the game in different situations. Pham and Cid [26] extended the FlipIt game by considering that players have the ability to check the state of the resource before their moves.

Feng et al. [6] and Hu et al. [9] modified the FlipIt game by considering insiders in addition to external adversaries. Zhang et al. [31] studied the FlipIt game with resource constraints on both players. Pawlick et al. extended the FlipIt game with characteristics of signaling games [25]. Wellman and Prakash developed a discrete-time model with multiple, ordered states in which attackers may compromise a server through cumulative acquisition of knowledge rather than in a one-shot takeover [30].

The original FlipIt paper assumed that the players compete with each other for *one* resource. Laszka et al. [14] addressed this limitation by modeling multiple contested resources in a game called *FlipThem*. Other authors extended this game by considering a threshold for the number of contested resources which need to be compromised to achieve the attacker's objective [18]. In a similar way, a variation of the game has been proposed with multiple defenders [24]. Laszka et al. [15,16] studied timing issues when the attacker's moves are non-instantaneous. Moreover, they considered that the defender's moves are non-covert and the attacker's type can be targeting and non-targeting. Johnson et al. [11] investigate the role of time in dynamic environments where an adversary discovers vulnerabilities based on an exogenous vulnerability discovery process and each vulnerability has its corresponding survival time.

Complementing these theoretical analyses, Nochenson and Grossklags [22] as well as Grossklags and Reitter [8] study human defensive players when they interact with a computerized attacker in the FlipIt framework.

Our work differs from the previous FlipIt literature regarding two key considerations. First, we take into account the problem of information leakage and propose a more realistic game-theoretic framework for defender's best time to update his defense mechanism. We propose a model in which an attacker will incrementally take ownership of a resource. Note that the attacker's goal is to compromise the defender's system completely, but she may acquire already some benefit during the initial steps of her attack. Second, we consider the possibility of the defender's defense strategy not being able to completely eliminate the attacker's foothold in the system. As a result, our work overcomes several significant simplifications in the previous literature which limited their applicability to realistic defense scenarios.

3 Model Definition

In this section, we provide a description of the FlipLeakage model which is a two-player game between a defender ($\mathcal{D}$) and an attacker ($\mathcal{A}$). We use the term **resource** for the defended system, but also for the target of the attack which will leak information during the attack and after the successful compromise. The attack progresses in a **stealthy** fashion. However, the defender can regain partial control over a compromised resource by taking a defensive recovery move (e.g., a variety of system updates).

In the FlipLeakage model, we emphasize the following aspects which we will discuss below: (1) **uncertainty** about the time of compromising the defender's resource entirely, (2) process of **information leakage**, (3) **quality** of defensive moves, (4) **strategies** of both players, and (5) other parameters which are necessary for our model.

Uncertainty About Attack Launch and Success Timings: In FlipLeakage, the defender is the owner of the resource at the beginning of the game. The resource is in a secure state, when it is completely controlled by the defender. However, due to the stealthy nature of many practically deployed attacks, e.g., related to cyber-espionage and advanced persistent threats, it is reasonable to assume that the defender cannot acquire any information about the time when an attack is launched as well as its success [21].

In contrast, we assume that the attacker can observe the time of a defender's move. One motivating practical example for this consideration is that many software companies publicly announce the arrival of new patches for previously discovered vulnerabilities. Hence, an attacker could infer when a certain system weakness is not available anymore. It follows that we model asymmetry with respect to knowledge between the two players.

Furthermore, we differentiate between the time of launching an attack and the time of an attack's full effectiveness (i.e., the resource is completely compromised). It is worth mentioning that the value of this time difference is not known to both the defender and the attacker. Hence, this time difference is represented by a random variable $t_{\mathcal{A}}$ with probability density function $f_{\mathcal{A}}(t_{\mathcal{A}})$. The value of $t_{\mathcal{A}}$ depends on many factors such as the defender's defense strategy and the attacker's ability to compromise the defender's system.

The gap between these two factors can be interpreted as the attacker requiring a nontrivial amount of time and effort to control the resource completely, e.g., to gather leaked information from the resource and to conduct subsequent attack steps. Further, the time of launching an attack can be understood as the time that the attacker starts to gather information from the defender to execute the attack successfully (e.g., by conducting surveillance of the system setup or pilfering traffic to collect information that will enable a security compromise). For simplicity, we assume that the value of $t_{\mathcal{A}}$ is chosen according to a random variable, but it is constant during each round of the attack. For future work, we are going to consider the case where the values of $t_{\mathcal{A}}$ are different for each round of the attack. Note that we assume that other important parameters of the game are common knowledge between the players. The extension of the framework to uncertainty about game-relevant parameters is subject of future work

Process of Information Leakage: After initiation of the attack move, the attacker's reward until a complete compromise is accomplished is based on the percentage of the whole resource which is currently controlled by the attacker. For this purpose, we consider a function $g_{\mathcal{A}}(t)$ (which is increasing on the range $[0, 1]$). $g_{\mathcal{A}}(t)$ can also be interpreted as the normalized amount of leaked information accessible to the attacker over time which can be used by her to improve her attack effectiveness. Recall that the time of completing an attack successfully is represented by a random variable $t_{\mathcal{A}}$. It follows that the function $g_{\mathcal{A}}(t)$ should be dependent on $t_{\mathcal{A}}$. In doing so, we define a general function $g_{\mathcal{A}}(t)$ reaching to 1 (i.e., the amount at which the attacker would control the whole resource completely) at one unit of time. We represent, as an example, a simple version of this function in the left-hand side of Fig. 2. To represent the described dependency, we use then the function $g_{\mathcal{A}}(t/t_{\mathcal{A}})$ for the reward calculation for

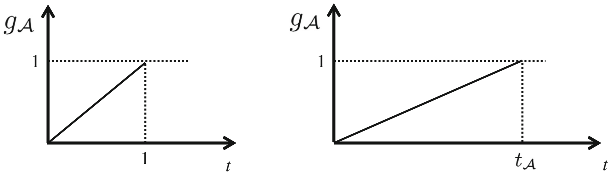

Fig. 2. The attacker's reward function during the time of completing an attack successfully depends on $t_{\mathcal{A}}$. To show this dependency in our model, we define a function as shown on the left-hand side of this figure with one unit of time to reach 1. The figure on the right-hand side is $g_{\mathcal{A}}\left(t/t_{\mathcal{A}}\right)$ representing this dependence.

the attacker during the time of completing the attack successfully, i.e., as shown on the right-hand side of Fig. 2.

Defense Quality: In FlipLeakge, we consider the quality of the defender's recovery action (or alternatively the ability of the attacker to transfer information from a previous attack to the next attempt). That is, the defender's recovery action does not guarantee regaining complete control over the resource, so that the attacker has an initial advantage (during the next attack attempt) and retains a foothold in the system. In other words, the defender's defense strategy cannot entirely eliminate previous attacks' effects. Motivating examples for this imperfect recovery model are *advanced persistent threats*. These attacks are typically driven by human staff who intelligently make use of any available and gathered information during the next multi-stage attack step which may include an initial compromise, foothold establishment, reconnaissance, etc. In this scenario, any recovery move by the defender will frequently only partially remove the attacker from the system, or at the very least cannot eliminate any information advantage by the attacker. In the FlipLeakage game, we introduce a new random variable, i.e., α with range $[0,1]$, to represent the fraction of retained control over the previously compromised resource by the attacker after the defender's recovery move.

In the worst case, the defender's recovery move does not impact the level of the resource being controlled by the attacker (i.e., $\alpha = 1$). In contrast, $\alpha = 0$ represents the situation when the defender's recovery is perfect. Then, the attacker has to start with a zero level of knowledge during her next attack. We model α as a continuous random variable with PDF $f_{\alpha}(.)$ in which α chooses values between zero and one, i.e., $\alpha \in [0,1]$. Note that in the FlipLeakage model, the attacker never starts with a higher level than the level attained in the most recent compromise attempt, i.e., we assume that defense moves are not counterproductive. For simplicity, we assume that the random variable α takes its value after the first attack and it remains constant during the game. For future work, we will consider the case where the values of α are completely independent from each other in each step of the attack.

Players' Strategies: In FlipLeakage, we assume that the defender moves according to periodic strategies, i.e., the time interval between two consecutive moves is identical and denoted by $\delta_{\mathcal{D}}$. In what follows, we provide two examples to show that in practice, several major software vendor organizations update their security policies in a periodic manner to underline the practical relevance of this assumption.

The first example that we take into account are Microsoft's security policy updates which are known as *Patch Tuesday*, i.e., according to [20], "Microsoft security bulletins are released on the second Tuesday of each month." We visualize the time differences among security updates from March 14th, 2015, until March 12th, 2016, which is shown in Fig. 3(a). In this figure, the vertical axis represents the number of security updates for each update instance. On the horizontal axis, 0 represents the first security update we take into account which took place on March 14th, 2015. Based on this figure, Microsoft security policy updates are almost perfectly periodic. We only observe two dates with out-of-schedule security updates. These two security updates are corresponding to an update for Internet Explorer and a vulnerability in a Microsoft font driver which allowed remote code execution.

Another example are Oracle's critical patch updates. These updates occur in January, April, July, and October of each year. To visualize the time differences between updates, which are shown in Fig. 3(b), we consider Oracle's critical patch updates from 13 July, 2013, to January 19, 2016, based on available information at [23]. We calculate the time differences between two consecutive patch updates in terms of days and divided this number by 30 in order to calculate an approximate difference in months. In this figure, 1 along the vertical axis represents the occurrence of a patch update. We observe that Oracle's policy for critical patch updates is almost periodic.[1]

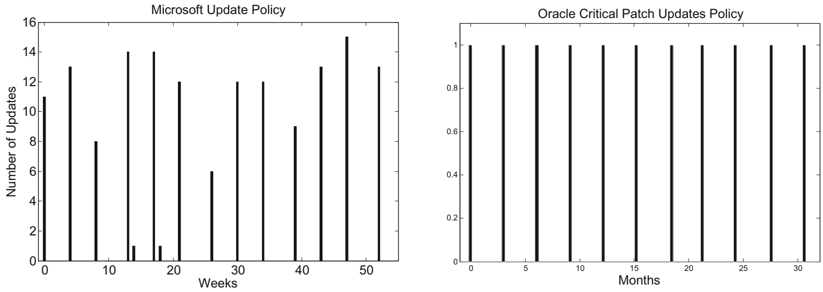

(a) Microsoft policy for security updates (b) Oracle critical patch updates

Fig. 3. In practice, many organizations update their system according to periodic strategies. As examples, we provide two organizations: (1) Microsoft and (2) Oracle.

[1] Note that in our model, we do not consider the case where a software vendor has the ability to conduct out-of-schedule security updates. We are going to consider this issue in future work.

In the FlipLeakage model, we assume that the attacker moves right after the defender. We follow with this assumption the results of [16] who showed that in the scenario of a defender with a periodic strategy, the best strategy for the attacker, who has the ability to observe the defender's defense strategy, is to move right after the defender.

Other Parameters: The cost of the defender's recovery moves and the attacker's attack moves are represented by $c_{\mathcal{D}}$ and $c_{\mathcal{A}}$, respectively, and we assume that they do not change over time. Examples of the defender's moves are changes of passwords, reinstallations of systems, and the application of new patches. Taking steps to incrementally infer cryptographic keys, brute-force passwords, or to inject malware are examples of the attacker's moves.

Once the attacker controls the resource completely, she receives an immediate reward which is represented by a constant value $I_{\mathcal{A}}$. The rationale behind the introduction of this parameter is that once the attacker infers the defender's secret such as a cryptographic key, she can, for example, decrypt secret messages which she has collected.

For the time that the attacker (defender) controls the resource completely, we assume that the defender's (attacker's) reward is equal to zero and the attacker (defender) receives $B_{\mathcal{A}}$ ($B_{\mathcal{D}}$) per unit of time controlling the resource. For example, these incremental earnings for the attacker represent newly arriving messages which can be decrypted with the compromised key. Note that the resource is controlled by the attacker completely after a successful attack and before the next recovery move by the defender.

4 Payoff Model

In this section, we develop the payoff functions for the FlipLeakage model based on what we presented in Sect. 3.

The time required to execute an attack successfully is defined by a continuous random variable with PDF $f_{\mathcal{A}}$. We consider one of the realizations of this random variable as $t_{\mathcal{A}}$. Moreover, the time between two consecutive defender's moves is represented by $\delta_{\mathcal{D}}$. Based on the $t_{\mathcal{A}}$ realization, we have two possible cases, i.e., $t_{\mathcal{A}} \geq \delta_{\mathcal{D}}$ and $t_{\mathcal{A}} < \delta_{\mathcal{D}}$. In what follows, we consider each of these two cases separately and then combine them according to the probability of each case to propose the payoff function.

Case 1: $t_{\mathcal{A}} < \delta_{\mathcal{D}}$

In this case, the attacker can complete her attack before the defender's recovery move. Hence, she receives the immediate reward for compromising the resource completely, i.e., $I_{\mathcal{A}}$, as well as the reward for controlling the resource completely, i.e., $B_{\mathcal{A}}$.

In our model, we assume that the attacker's control over the resource does not fall to zero right after the defender's recovery move. As discussed in Sect. 3, we have introduced a new parameter, α, and described the resulting changes to players' payoffs. For $t_{\mathcal{A}} < \delta_{\mathcal{D}}$, the attacker controls the resource completely before the next recovery move by the defender. Then, right after the defender's move, the attacker controls a fraction α of the resource. For the remainder of the resource to be taken over, i.e., $(1-\alpha)$, the attacker can gain control based on $g_{\mathcal{A}}\left(t/t_{\mathcal{A}}\right)$. Hence, the attacker's benefit for this period is then based on $\alpha + (1-\alpha)\, g_{\mathcal{A}}\left(t/t_{\mathcal{A}}\right)$. The attacker's payoff is as follows:

$$u_{\mathcal{A}}^{1}(t_{\mathcal{A}},\alpha,\delta_{\mathcal{D}}) = \frac{\displaystyle\int_{0}^{t_{\mathcal{A}}} \left(\alpha + (1-\alpha)\, g_{\mathcal{A}}\left(\frac{t}{t_{\mathcal{A}}}\right)\right) dt + I_{\mathcal{A}} + B_{\mathcal{A}}(\delta_{\mathcal{D}} - t_{\mathcal{A}}) - c_{\mathcal{A}}}{\delta_{\mathcal{D}}}. \tag{1}$$

In the above equation, based on our discussion in Sect. 3, the first term in the numerator represents the attacker's benefit due to information leakage. Note that the utility function is divided by $\delta_{\mathcal{D}}$, since this function is the average attacker's payoff over time.

Since the defender's move time is greater than the attacker's time of completing an attack successfully, the defender only receives a partial benefit during the period when the attacker is in the process of completing her attack. Therefore, the defender's payoff is as follows:

$$u_{\mathcal{D}}^{1}(t_{\mathcal{A}},\alpha,\delta_{\mathcal{D}}) = \frac{\displaystyle\int_{0}^{t_{\mathcal{A}}} \left(1 - \left(\alpha + (1-\alpha)\, g_{\mathcal{A}}\left(\frac{t}{t_{\mathcal{A}}}\right)\right)\right) dt - c_{\mathcal{D}}}{\delta_{\mathcal{D}}}. \tag{2}$$

Both payoff functions, i.e., Eqs. 1 and 2, are a function of $t_{\mathcal{A}}$ which is a random variable with PDF $f_{\mathcal{A}}$ as well as $\delta_{\mathcal{D}}$. Therefore, we need to calculate the expected value of both payoff functions. Note that these expected payoff functions are conditional, i.e., they are a function of a random variable $t_{\mathcal{A}}$ given that $t_{\mathcal{A}} < \delta_{\mathcal{D}}$. The conditional expected payoffs for these two functions are calculated as follows:

$$u_{\mathcal{A}}^{1}(\alpha,\delta_{\mathcal{D}}) = \frac{\displaystyle\int_{0}^{\delta_{\mathcal{D}}} u_{\mathcal{A}}^{1}(t_{\mathcal{A}},\alpha,\delta_{\mathcal{D}}) f_{\mathcal{A}}(t_{\mathcal{A}}) dt_{\mathcal{A}}}{\displaystyle\int_{0}^{\delta_{\mathcal{D}}} f_{\mathcal{A}}(t_{\mathcal{A}}) dt_{\mathcal{A}}}, \tag{3}$$

$$u_{\mathcal{D}}^{1}(\alpha,\delta_{\mathcal{D}}) = \frac{\displaystyle\int_{0}^{\delta_{\mathcal{D}}} u_{\mathcal{D}}^{1}(t_{\mathcal{A}},\alpha,\delta_{\mathcal{D}}) f_{\mathcal{A}}(t_{\mathcal{A}}) dt_{\mathcal{A}}}{\displaystyle\int_{0}^{\delta_{\mathcal{D}}} f_{\mathcal{A}}(t_{\mathcal{A}}) dt_{\mathcal{A}}}. \tag{4}$$

Defender's and attacker's payoffs are both functions of α and $\delta_{\mathcal{D}}$. Finally, the probability of $t_{\mathcal{A}} < \delta_{\mathcal{D}}$ is calculated as follows:

$$P[t_{\mathcal{A}} < \delta_{\mathcal{D}}] = \int_{0}^{\delta_{\mathcal{D}}} f_{\mathcal{A}}(t_{\mathcal{A}}) dt_{\mathcal{A}}. \tag{5}$$

Case 2: $t_{\mathcal{A}} \geq \delta_{\mathcal{D}}$

In contrast to the previous case, the attacker cannot get the immediate reward as well as the benefit from controlling the resource completely. In this case, the attacker only reaches $g_{\mathcal{A}}\left(\delta_{\mathcal{D}}/t_{\mathcal{A}}\right)$ level of control over the resource upon the defender's recovery move, and her reward is then equal to $\alpha g_{\mathcal{A}}\left(\delta_{\mathcal{D}}/t_{\mathcal{A}}\right)$ right after the defender's move. The attacker gains her control for the rest of the resource, i.e., $(1-\alpha)$, based on $g_{\mathcal{A}}\left(t/t_{\mathcal{A}}\right)$. Hence, during the time between two consecutive defender's moves, the attacker's benefit is equal to $\alpha g_{\mathcal{A}}\left(\delta_{\mathcal{D}}/t_{\mathcal{A}}\right) + (1-\alpha)\, g_{\mathcal{A}}\left(t/t_{\mathcal{A}}\right)$. Note that the upper integral bound changes into $\delta_{\mathcal{D}}$ from $t_{\mathcal{A}}$ compared to the previous case.

$$u^2_{\mathcal{A}}(t_{\mathcal{A}}, \alpha, \delta_{\mathcal{D}}) = \frac{\displaystyle\int_0^{\delta_{\mathcal{D}}} \left(\alpha g_{\mathcal{A}}\left(\frac{\delta_{\mathcal{D}}}{t_{\mathcal{A}}}\right) + (1-\alpha)\, g_{\mathcal{A}}\left(\frac{t}{t_{\mathcal{A}}}\right)\right) dt - c_{\mathcal{A}}}{\delta_{\mathcal{D}}}. \tag{6}$$

The defender's payoff function is almost equivalent to Eq. 2 except the upper bound for the integral is changed into $\delta_{\mathcal{D}}$. Hence, the defender's payoff is as follows:

$$u^2_{\mathcal{D}}(t_{\mathcal{A}}, \alpha, \delta_{\mathcal{D}}) = \frac{\displaystyle\int_0^{\delta_{\mathcal{D}}} \left(1 - \left(\alpha g_{\mathcal{A}}\left(\frac{\delta_{\mathcal{D}}}{t_{\mathcal{A}}}\right) + (1-\alpha)\, g_{\mathcal{A}}\left(\frac{t}{t_{\mathcal{A}}}\right)\right)\right) dt - c_{\mathcal{D}}}{\delta_{\mathcal{D}}}. \tag{7}$$

Both players' payoffs are functions of $t_{\mathcal{A}}$, α, and $\delta_{\mathcal{D}}$. We take the conditional expectation over parameter $t_{\mathcal{A}}$ in order to calculate the average payoffs with respect to $t_{\mathcal{A}}$ for this condition. The resulting equations are:

$$u^2_{\mathcal{A}}(\alpha, \delta_{\mathcal{D}}) = \frac{\displaystyle\int_{\delta_{\mathcal{D}}}^{\infty} u^2_{\mathcal{A}}(t_{\mathcal{A}}, \alpha, \delta_{\mathcal{D}}) f_{\mathcal{A}}(t_{\mathcal{A}}) dt_{\mathcal{A}}}{\displaystyle\int_{\delta_{\mathcal{D}}}^{\infty} f_{\mathcal{A}}(t_{\mathcal{A}}) dt_{\mathcal{A}}}, \tag{8}$$

$$u^2_{\mathcal{D}}(\alpha, \delta_{\mathcal{D}}) = \frac{\displaystyle\int_{\delta_{\mathcal{D}}}^{\infty} u^2_{\mathcal{D}}(t_{\mathcal{A}}, \alpha, \delta_{\mathcal{D}}) f_{\mathcal{A}}(t_{\mathcal{A}}) dt_{\mathcal{A}}}{\displaystyle\int_{\delta_{\mathcal{D}}}^{\infty} f_{\mathcal{A}}(t_{\mathcal{A}}) dt_{\mathcal{A}}}. \tag{9}$$

Furthermore, the probability that the required time by the attacker to compromise the resource entirely is greater than the time between two consecutive recovery moves is given by:

$$P[t_{\mathcal{A}} \geq \delta_{\mathcal{D}}] = \int_{\delta_{\mathcal{D}}}^{\infty} f_{\mathcal{A}}(t_{\mathcal{A}}) dt_{\mathcal{A}}. \tag{10}$$

By taking into account the probability of occurrence of each condition as well as their corresponding payoffs, we can calculate the defender's and the attacker's payoff functions which are represented by the following equations, respectively.

$$u_{\mathcal{D}}(\alpha, \delta_{\mathcal{D}}) = P[t_{\mathcal{A}} \geq \delta_{\mathcal{D}}] u^2_{\mathcal{D}}(\alpha, \delta_{\mathcal{D}}) + P[t_{\mathcal{A}} < \delta_{\mathcal{D}}] u^1_{\mathcal{D}}(\alpha, \delta_{\mathcal{D}}), \tag{11}$$

$$u_{\mathcal{A}}(\alpha, \delta_{\mathcal{D}}) = P[t_{\mathcal{A}} \geq \delta_{\mathcal{D}}]u^2_{\mathcal{A}}(\alpha, \delta_{\mathcal{D}}) + P[t_{\mathcal{A}} < \delta_{\mathcal{D}}]u^1_{\mathcal{A}}(\alpha, \delta_{\mathcal{D}}). \tag{12}$$

In the above equation, each player's payoff is a function of α and $\delta_{\mathcal{A}}$. As mentioned before, α is a random variable whose range is in $[0,1]$ with PDF $f_\alpha(.)$. Therefore, we can calculate the expected value of the defender's and the attacker's payoff functions with respect to α being represented in the following equations, respectively.

$$u_{\mathcal{D}}(\delta_{\mathcal{D}}) = \int_0^1 u_{\mathcal{D}}(\alpha, \delta_{\mathcal{D}}) f_\alpha(\alpha) d\alpha, \tag{13}$$

$$u_{\mathcal{A}}(\delta_{\mathcal{D}}) = \int_0^1 u_{\mathcal{A}}(\alpha, \delta_{\mathcal{D}}) f_\alpha(\alpha) d\alpha. \tag{14}$$

5 Analytical Results

In the previous section, we have developed the general payoff functions for the FlipLeakage model. Our payoff calculations are general and can be applied to many cybersecurity problems and we did not quantify any of the parameters being used in our model. For our analyses in this paper, we quantify $g_{\mathcal{A}}(.)$, $f_{\mathcal{A}}(.)$, and $f_\alpha(.)$, but we believe that the concrete functions we use still allow for meaningful insights about the stealthy information leakage scenarios. The instantiations of the other parameters in our proposed models would be specific to the concrete scenario under consideration, e.g., the corresponding cost for each player as well as the benefits.

To model the time of the attacker completing her attack successfully, we consider an *exponential* distribution with rate parameter $\lambda_{\mathcal{A}}$. The rationale behind choosing an exponential distribution for the random variable $t_{\mathcal{A}}$ is the memoryless feature of this distribution. Due to the memoryless condition, if the defender knows that his system is not compromised entirely at a specific time, it does not give any further information to the defender about the time of the next potential compromise. Moreover, the exponential distribution is a widely accepted candidate to model waiting times for event-driven models. The exponential distribution with rate parameter $\lambda_{\mathcal{A}}$ is as follows:

$$f_{\mathcal{A}}(t_{\mathcal{A}}) = \begin{cases} \lambda_{\mathcal{A}} e^{-\lambda_{\mathcal{A}} t_{\mathcal{A}}} & \text{if } t_{\mathcal{A}} \geq 0 \\ 0 & \text{if } t_{\mathcal{A}} < 0. \end{cases} \tag{15}$$

Moreover, for the random variable $\alpha \in [0,1]$, we consider the uniform distribution, since the defender does not have any knowledge about the ability of the attacker to use previously leaked information and, accordingly, all values are possible with the same probability. The uniform distribution, $f_\alpha(.)$, is represented in Eq. 16.

$$f_\alpha(\alpha) = \begin{cases} 1 & \text{if } 0 \leq \alpha \leq 1 \\ 0 & \text{Otherwise.} \end{cases} \tag{16}$$

The attacker's reward function during the time to launch her attack successfully can be represented by a linear function:

$$g_{\mathcal{A}}\left(\frac{t}{t_{\mathcal{A}}}\right) = \begin{cases} \frac{t}{t_{\mathcal{A}}} & \text{if } 0 \leq t \leq t_{\mathcal{A}} \\ 0 & \text{Otherwise.} \end{cases} \tag{17}$$

In the following, we provide our lemmas and theorem based on our payoff calculation and the specification described above. First, the defender's and the attacker's best responses are stated in Lemma 1 and Lemma 2, respectively. Then, we propose the Nash equilibrium of the game being stated in Theorem 1.

Lemma 1. *The defender's best response is as follows:*

- *The defender plays a periodic strategy with period $\delta^{\star}_{\mathcal{D}}$ which is the solution of Eq. 18, if the corresponding payoff is non-negative, i.e., $u_{\mathcal{D}}(\delta^{\star}_{\mathcal{D}}) \geq 0$, and it yields a higher payoff compared to other solutions of Eq. 18.*

$$\begin{aligned} BR_{\mathcal{D}} = e^{-\lambda_{\mathcal{A}}\delta_{\mathcal{D}}} \left(\frac{1}{4} - \frac{3}{4}\lambda_{\mathcal{A}}\delta_{\mathcal{D}} + \frac{3}{4}\lambda_{\mathcal{A}} + \frac{1}{4\lambda_{\mathcal{A}}\delta^2_{\mathcal{D}}}\right) \\ + \frac{1}{\delta^2_{\mathcal{D}}}\left(c_{\mathcal{D}} - \frac{1}{4\lambda_{\mathcal{A}}}\right) - \frac{3}{4}\lambda_{\mathcal{A}}\Gamma(0, \lambda_{\mathcal{A}}\delta_{\mathcal{D}}) = 0. \end{aligned} \tag{18}$$

- *The defender drops out of the game (i.e., the player does not move anymore) if Eq. 18 has no solution for $\delta_{\mathcal{D}}$.*
- *The defender drops out of the game if the solutions of Eq. 18 yield a negative payoffs, i.e., $u_{\mathcal{D}}(\delta_{\mathcal{D}}) < 0$.*

Note that in Lemma 1, $\Gamma(0, \lambda_{\mathcal{A}}\delta_{\mathcal{D}})$ represents a Gamma function which is defined as follows:

$$\Gamma(s, x) = \int_x^{\infty} t^{s-1}e^{-t}dt. \tag{19}$$

Proof of Lemma 1 is provided in Appendix A.1.

Lemma 1 exhibits how we should calculate the defender's time between his two consecutive moves. As we see in Eq. 18, the defender's best response is a function of $c_{\mathcal{D}}$ and $\lambda_{\mathcal{A}}$.

Lemma 2 describes the attacker's best response in the FlipLeakage game.

Lemma 2. *In the FlipLeakage game model, the attacker's best response is:*

- *The attacker moves right after the defender if $c_{\mathcal{A}} < M(\delta)$ where*

$$\begin{aligned} M(\delta_{\mathcal{D}}) = \frac{3}{4}\delta_{\mathcal{D}}\lambda_{\mathcal{A}}\Gamma(0, \delta_{\mathcal{D}}\lambda_{\mathcal{A}}) + I_{\mathcal{A}} + B_{\mathcal{A}}\delta_{\mathcal{D}} + \frac{3}{4\lambda_{\mathcal{A}}} + B_{\mathcal{A}}\left(\delta_{\mathcal{D}} + \frac{1}{\lambda_{\mathcal{A}}}\right)e^{-\delta_{\mathcal{D}}\lambda_{\mathcal{A}}} \\ - \left(I_{\mathcal{A}} + B_{\mathcal{A}}\delta_{\mathcal{D}} + \frac{3}{4}\left(\delta_{\mathcal{D}} + \frac{1}{\lambda_{\mathcal{A}}}\right)\right)e^{-\delta_{\mathcal{D}}\lambda_{\mathcal{A}}} - \frac{B_{\mathcal{A}}}{\lambda_{\mathcal{A}}}. \end{aligned} \tag{20}$$

- *The attacker drops out of the game if $c_{\mathcal{A}} > M(\delta)$.*

– *Otherwise, i.e., $c_{\mathcal{A}} = M(\delta)$, dropping out of the game and moving right after the defender are both the attacker's best responses.*

The proof of Lemma 2 is provided in Appendix A.2. This lemma identifies conditions in which the attacker should move right after the defender, not move at all, and be indifferent between moving right after the defender and not moving at all. Note that the attacker's decision depends on $c_{\mathcal{A}}$, $\delta_{\mathcal{D}}$, $\lambda_{\mathcal{A}}$, $I_{\mathcal{A}}$, and $B_{\mathcal{A}}$.

The following theorem describes the Nash equilibria of the FlipLeakage game based on our described lemmas.

Theorem 1. *The FlipLeakage game's pure Nash equilibria can be described as follows.*

A. If Eq. 18 has a solution, i.e., $\delta^{\star}_{\mathcal{D}}$, yielding the highest positive payoff for the defender compared to other solutions (if other solutions exist), then the following two outcomes apply:

1- If $c_{\mathcal{A}} \leq M(\delta_{\mathcal{D}})$, then there is a unique pure Nash equilibrium in which the defender moves periodically with period $\delta^{\star}_{\mathcal{D}}$ and the attacker moves right after the defender.

2- If $c_{\mathcal{A}} > M(\delta_{\mathcal{D}})$, then there exists no pure Nash equilibrium.

B. If Eq. 18 does not have a solution or the solutions of this equation yield a negative payoff for the defender, i.e., $u_{\mathcal{D}}(\delta_{\mathcal{D}}) < 0$, then there exists a unique pure Nash equilibrium in which the defender does not move at all and the attacker moves once at the beginning of the FlipLeakage game.

The proof of Theorem 1 is provided in Appendix A.3.

In this theorem, in the first case, the defender's cost is lower than his benefit when he moves according to the solution of Eq. 18 and the attacker's cost is lower than Eq. 20. Hence, the attacker moves right after the defender's periodic move. In the second case, if the defender moves periodically, it is not beneficial for the attacker to move at all. Therefore, it is better for the defender to not move at all. But, if the defender does not move at all, the attacker can move once at the beginning of the game and control the resource for all time. However, as a result, the defender should move in order to hinder this situation. Because of this strategic uncertainty, in this scenario a Nash equilibrium does not exist. The third case represents the situation where the defender's benefit is lower than his cost for defending the resource. Then, it is beneficial for him to not move at all, and because of that the attacker has to move only once at the beginning of the game.

6 Numerical Illustrations

In this section, we provide selected numerical illustrations for our theoretical findings. First, we represent the defender's best response curves, i.e., Eq. 18, as well as the defender's payoff for different defender's cost values, i.e., $c_{\mathcal{D}}$, which are depicted in Fig. 4. Then, we illustrate the defender's best responses for different values of $c_{\mathcal{D}}$ and $\lambda_{\mathcal{A}}$ in Fig. 5.

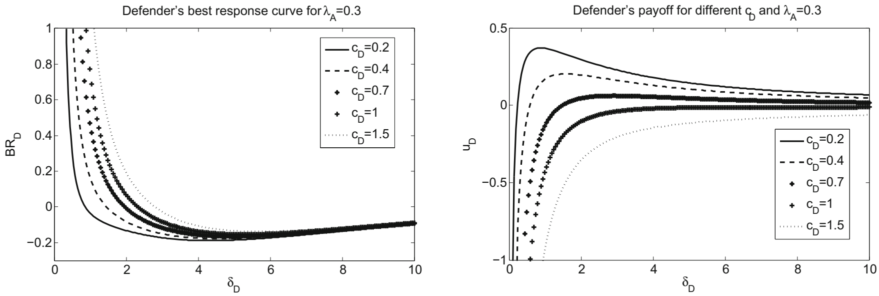

(a) Defender's best response curve for different values of $c_{\mathcal{D}}$.

(b) Defender's payoff as function of $\delta_{\mathcal{D}}$.

Fig. 4. The defender's best response curves and the corresponding payoff functions for different values of $c_{\mathcal{D}}$ are represented in (a) and (b), respectively. These two figures depict the situation that Eq. 18 has a solution, but the corresponding payoff may be negative.

We plot Eq. 18, i.e., the defender's best response curve, for different values of $c_{\mathcal{D}}$, i.e., $c_{\mathcal{D}} = \{0.2, 0.4, 0.7, 1, 1.5\}$, and $\lambda_{\mathcal{A}} = 0.3$ in Fig. 4(a). We illustrate the defender's payoff for these values in Fig. 4(b), as well. For all of these different $c_{\mathcal{D}}$s, Eq. 18 has a solution. But as we see in Fig. 4(b), the defender's payoffs are negative for $c_{\mathcal{D}} = \{0.7, 1, 1.5\}$ for all values of $\delta_{\mathcal{D}}$. Therefore, the defender will drop out of the game given these defense costs. For lower values of $c_{\mathcal{D}}$, i.e., $c_{\mathcal{D}} = \{0.2, 0.4\}$, the defender's best responses are to move periodically with period 0.8711 and 1.4681, respectively. This also provides us with the intuition that the higher the defender's costs are, the slower will be the defender's moves. To examine this intuition, we calculate the defender's best responses for different values of $c_{\mathcal{D}}$.

Figure 5(a) represents the defender's best response for different values of defense costs in which $\lambda_{\mathcal{A}} = 0.5$. This figure corroborates our intuition that the higher the defense costs are, the slower will be the defender's move period. When the cost of defense is high, the defender's best response is to drop out of the game which is represented as $\delta^{\star}_{\mathcal{D}} = 0$ in Fig. 5(a).

We are also interested to see the relation between $\lambda_{\mathcal{A}}$ and $\delta_{\mathcal{D}}$. We represent this relation in Fig. 5(b). It is worth mentioning that an exponential distribution with parameter $\lambda_{\mathcal{A}}$ has mean being equal to $1/\lambda_{\mathcal{A}}$. In the FlipLeakage game, a higher value of $\lambda_{\mathcal{A}}$ means that the attacker will successfully compromise the defender's system faster on average which is corresponding to $1/\lambda_{\mathcal{A}}$. Figure 5(b) represents the defender's best response for different values of $\lambda_{\mathcal{A}}$ for specific defender's cost, i.e., $c_{\mathcal{D}} = 0.3$. This figure shows that the faster the attacker can completely compromise the defender's system on average, the faster will be the defender's periodic move. In other words, the defender moves faster when the attacker's average time to successfully compromise the defender's system is faster. But if the attacker's average time to successfully compromise the

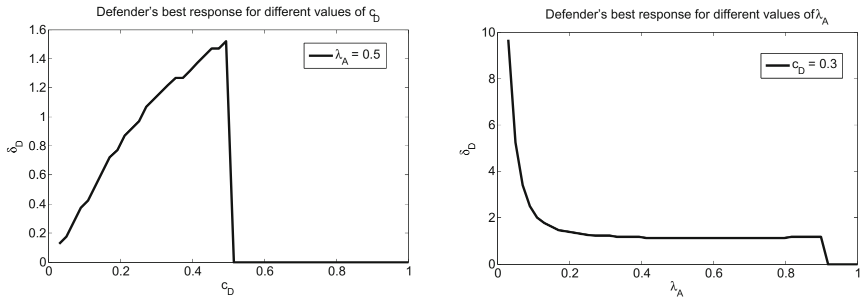

(a) The defender's best response with respect to his cost

(b) The defender's best response with respect to $\lambda_{\mathcal{A}}$

Fig. 5. The impact of $c_{\mathcal{D}}$ and $\lambda_{\mathcal{A}}$ on the defender's best response

defender's system is too fast, the rational choice for the defender is to drop out of the game.

7 Conclusion

In this paper, we have proposed a novel theoretical model to provide guidance for the defender's optimal defense strategy when faced with a stealthy information leakage threat. In our model, an attacker will incrementally take ownership of a resource (e.g., as observed during advanced persistent threats). While her final objective is a complete compromise of the system, she may derive some utility during the preliminary phases of the attack. The defender can take a costly mitigation move and has to decide on its optimal timing.

In the FlipLeakage game model, we have considered the scenario when the defender only partially eliminates the foothold of the attacker in the system. In this scenario, the defender cannot undo any information leakage that has already taken place during an attack. We have derived optimal strategies for the agents in this model and present numerical analyses and graphical visualizations.

We highlight two observations from our numerical analyses which match well with intuition. First, the higher the defender's cost, the slower is the defender's periodic move. The second observation is that the faster the attacker's average time to compromise the defender's system completely (i.e., higher $\lambda_{\mathcal{A}}$), the faster is the defender's periodic move. In addition, our model also allows for the determination of the impact of less-than-optimal strategies, and comparative statements regarding the expected outcomes of different periodic defensive approaches in practice, when information about the attacker and her capabilities is extremely scarce. As this problem area is understudied but of high practical significance, advancements that allow a rigorous reasoning about defense moves against stealthy attackers are of potentially high benefit.

In future work, we aim to conduct theoretical and numerical analyses using insights from data about practical information leakage scenarios. However, our current study is an important first step to reason about frequently criticized system reset policies to prevent information leakage in high-value systems. Reset policies have to provide an expected utility in the absence of concrete evidence due to the stealthiness of attacks which can be challenging to articulate. Our work also illustrates the positive deterrence function of system reset policies from a theoretical perspective. Further, we aim to consider a more general case in which the values of $t_{\mathcal{A}}$ and α are different in each step of the attack. In future work, we will also consider the case where a defender (e.g., a software vendor) has the ability to provide out-of-schedule security updates besides the periodic one.

Acknowledgments. We appreciate the comments from the anonymous reviewers. An earlier version of this paper benefited from the constructive feedback from Aron Laszka. All remaining errors are our own.

A Proof

A.1 Proof of Lemma 1

Based on our payoff calculation, i.e., Eq. 13, as well as the quantified parameters, i.e., $g_{\mathcal{A}}(.)$, $f_{\mathcal{A}}(.)$, and $f_{\alpha}(.)$, the defender's payoff is:

$$u_{\mathcal{D}}(\delta_{\mathcal{D}}) = \frac{1}{\delta_{\mathcal{D}}}\left(\frac{1}{4\lambda_{\mathcal{A}}}\left(1-e^{-\lambda_{\mathcal{A}}\delta_{\mathcal{D}}}\right)-c_{\mathcal{D}}\right)+\frac{3}{4}e^{-\lambda_{\mathcal{A}}\delta_{\mathcal{D}}}-\frac{3}{4}\lambda_{\mathcal{A}}\delta_{\mathcal{D}}\Gamma(0,\lambda_{\mathcal{A}}\delta_{\mathcal{D}}) \tag{21}$$

To find the maximizing time between two consecutive defender's moves (if there exist any), we take the partial derivative of Eq. 21 with respect to $\delta_{\mathcal{D}}$ and solve it for equality to 0 as follows:

$$\begin{aligned}\frac{\partial u_{\mathcal{D}}}{\partial \delta_{\mathcal{D}}} &= -\frac{1}{\delta_{\mathcal{D}}^2}\left(\frac{1}{4\lambda_{\mathcal{A}}}-c_{\mathcal{D}}-\frac{1}{4\lambda_{\mathcal{A}}}e^{-\lambda_{\mathcal{A}}\delta_{\mathcal{D}}}\right)+\frac{1}{4}e^{-\lambda_{\mathcal{A}}\delta_{\mathcal{D}}}\\ -\frac{3}{4}\lambda_{\mathcal{A}}\delta_{\mathcal{D}}e^{-\lambda_{\mathcal{A}}\delta_{\mathcal{D}}}&-\frac{3}{4}\lambda_{\mathcal{A}}\Gamma(0,\lambda_{\mathcal{A}}\delta_{\mathcal{D}})+\frac{3}{4}\lambda_{\mathcal{A}}e^{-\lambda_{\mathcal{A}}\delta_{\mathcal{D}}}=0\end{aligned} \tag{22}$$

Note that Eq. 18 is neither increasing nor decreasing on $\delta_{\mathcal{D}}$. Therefore, we have three possibilities for the above equation: (1) no solution, (2) one solution, and (3) more than one solution. When there is no solution, the defender's best response is to drop out of the game. In the case of one solution, the defender moves periodically with $\delta_{\mathcal{D}}$, i.e., the solution of Eq. 18 if the resulting payoff is non-negative. When there is more than one solution, the defender plays periodically with the solution with the highest non-negative payoff. Otherwise, the defender drops out of the game. □

A.2 Proof of Lemma 2

In order to calculate the attacker's payoff, we first calculate the following based on Eq. 12.

$$\begin{aligned} u_{\mathcal{A}}(\alpha, \delta_{\mathcal{D}}) = \frac{1}{\delta_{\mathcal{D}}} \Bigg(\left(\frac{1+\alpha}{2} - B_{\mathcal{A}} \right) \left(\frac{1}{\lambda_{\mathcal{A}}} - \frac{1}{\lambda_{\mathcal{A}}} e^{-\lambda_{\mathcal{A}} \delta_{\mathcal{D}}} - \delta_{\mathcal{D}} e^{-\lambda_{\mathcal{A}} \delta_{\mathcal{D}}} \right) \\ + (B_{\mathcal{A}} \delta_{\mathcal{D}} + I_{\mathcal{A}}) \left(1 - e^{-\lambda_{\mathcal{A}} \delta_{\mathcal{D}}}\right) + \frac{1+\alpha}{2} \delta_{\mathcal{D}} \lambda_{\mathcal{A}} \Gamma(0, \delta_{\mathcal{D}} \lambda_{\mathcal{A}}) - c_{\mathcal{A}} \Bigg). \end{aligned} \tag{23}$$

According to Eq. 14, the attacker's payoff is as follows.

$$\begin{aligned} u_{\mathcal{A}}(\delta_{\mathcal{D}}) = \frac{1}{\delta_{\mathcal{D}}} \Bigg(\left(\frac{3}{4} - B_{\mathcal{A}} \right) \left(\frac{1}{\lambda_{\mathcal{A}}} - \frac{1}{\lambda_{\mathcal{A}}} e^{-\lambda_{\mathcal{A}} \delta_{\mathcal{D}}} - \delta_{\mathcal{D}} e^{-\lambda_{\mathcal{A}} \delta_{\mathcal{D}}} \right) \\ + (B_{\mathcal{A}} \delta_{\mathcal{D}} + I_{\mathcal{A}}) \left(1 - e^{-\lambda_{\mathcal{A}} \delta_{\mathcal{D}}}\right) + \frac{3}{4} \delta_{\mathcal{D}} \lambda_{\mathcal{A}} \Gamma(0, \delta_{\mathcal{D}} \lambda_{\mathcal{A}}) - c_{\mathcal{A}} \Bigg). \end{aligned} \tag{24}$$

The attacker moves right after the defender if her payoff is positive, i.e., $u_{\mathcal{A}}(\delta_{\mathcal{D}}) > 0$. If the attacker's payoff is negative, her reward is lower than her cost. Then, a rational player does not have any incentive to actively participate in the game. Hence, the attacker drops out of the game. If $u_{\mathcal{A}}(\delta_{\mathcal{D}}) = 0$, the attacker is indifferent between moving right after the defender or dropping out of the game. By considering Eq. 24 and $u_{\mathcal{A}}(\delta_{\mathcal{D}}) \geq 0$, we can derive Eq. 20. □

A.3 Proof of Theorem 1

In Lemma 1, we have provided the best response for the defender. The defender has two choices: periodic move or dropping out of the game. Similarly, according to Lemma 2, the attacker has two choices for her best response: she moves right after the defender or drops out of the game. Note that Nash equilibrium is a mutual best response.

In doing so, we first consider the case where the defender's best response is to drop out of the game (this means that Eq. 18 does not have any solution(s) giving non-negative payoff(s)). Therefore, the attacker's best choice is to move only once at the beginning of the game.

The other choice for the defender, according to Lemma 1, is to move periodically when Eq. 18 has a solution which yields a positive payoff. By calculating $\delta_{\mathcal{D}}^{\star}$ using this equation, we insert this value to Eq. 20 and compare it with $c_{\mathcal{A}}$. Based on Lemma 2, the attacker has two possible choices. First, if $c_{\mathcal{A}} \leq M(\delta_{\mathcal{D}})$, the attacker will initiate her attack right after the defender's move. Hence, the Nash equilibrium is to move periodically from the defender side and the attacker should initiate her attack right after the defender's move. Second, if $c_{\mathcal{A}} > M(\delta_{\mathcal{D}})$, the attacker will drop out of the game. In this case, the best response for the defender is to never move. Since he controls the resource all the time without spending any cost. But, if the defender never moves, then it is beneficial for the attacker to move at the beginning of the game. Hence, this situation is not a Nash equilibrium. □

References

1. Blackwell, D.: The noisy duel, one bullet each, arbitrary accuracy. Technical report, The RAND Corporation, D-442 (1949)
2. Bowers, K.D., Dijk, M., Griffin, R., Juels, A., Oprea, A., Rivest, R.L., Triandopoulos, N.: Defending against the unknown enemy: applying FlipIt to system security. In: Grossklags, J., Walrand, J. (eds.) GameSec 2012. LNCS, vol. 7638, pp. 248–263. Springer, Heidelberg (2012). doi:10.1007/978-3-642-34266-0_15
3. Experian, Small business doesn't mean small data: Experian data breach resolution advises small businesses to be prepared for a data breach (2013). https://www.experianplc.com/media/news/
4. Falliere, N., Murchu, L., Chien, E.: W32.Stuxnet Dossier. Technical report, Symantec Corp., Security Response (2011)
5. Farhang, S., Manshaei, M.H., Esfahani, M.N., Zhu, Q.: A dynamic Bayesian security game framework for strategic defense mechanism design. In: Poovendran, R., Saad, W. (eds.) GameSec 2014. LNCS, vol. 8840, pp. 319–328. Springer, Heidelberg (2014). doi:10.1007/978-3-319-12601-2_18
6. Feng, X., Zheng, Z., Hu, P., Cansever, D., Mohapatra, P.: Stealthy attacks meets insider threats: a three-player game model. In: Proceedings of MILCOM (2015)
7. Grossklags, J., Christin, N., Chuang, J.: Secure or insure? A game-theoretic analysis of information security games. In: Proceedings of the 17th International World Wide Web Conference, pp. 209–218 (2008)
8. Grossklags, J., Reitter, D.: How task familiarity and cognitive predispositions impact behavior in a security game of timing. In: Proceedings of the 27th IEEE Computer Security Foundations Symposium (CSF), pp. 111–122 (2014)
9. Hu, P., Li, H., Fu, H., Cansever, D., Mohapatra, P.: Dynamic defense strategy against advanced persistent threat with insiders. In: Proceedings of the 34th IEEE International Conference on Computer Communications (INFOCOM) (2015)
10. Johnson, B., Grossklags, J., Christin, N., Chuang, J.: Uncertainty in interdependent security games. In: Alpcan, T., Buttyán, L., Baras, J.S. (eds.) GameSec 2010. LNCS, vol. 6442, pp. 234–244. Springer, Heidelberg (2010). doi:10.1007/978-3-642-17197-0_16
11. Johnson, B., Laszka, A., Grossklags, J.: Games of timing for security in dynamic environments. In: Khouzani, M.H.R., Panaousis, E., Theodorakopoulos, G. (eds.) GameSec 2015. LNCS, vol. 9406, pp. 57–73. Springer, Heidelberg (2015). doi:10.1007/978-3-319-25594-1_4
12. Joyce, R.: Disrupting nation state hackers (2016). https://www.youtube.com/watch?v=bDJb8WOJYdA
13. Laszka, A., Felegyhazi, M., Buttyan, L.: A survey of interdependent information security games. ACM Comput. Surv. **47**(2), 23:1–23:38 (2014)
14. Laszka, A., Horvath, G., Felegyhazi, M., Buttyán, L.: FlipThem: modeling targeted attacks with FlipIt for multiple resources. In: Poovendran, R., Saad, W. (eds.) GameSec 2014. LNCS, vol. 8840, pp. 175–194. Springer, Heidelberg (2014). doi:10.1007/978-3-319-12601-2_10
15. Laszka, A., Johnson, B., Grossklags, J.: Mitigating covert compromises. In: Chen, Y., Immorlica, N. (eds.) WINE 2013. LNCS, vol. 8289, pp. 319–332. Springer, Heidelberg (2013). doi:10.1007/978-3-642-45046-4_26
16. Laszka, A., Johnson, B., Grossklags, J.: Mitigation of targeted and non-targeted covert attacks as a timing game. In: Das, S.K., Nita-Rotaru, C., Kantarcioglu, M. (eds.) GameSec 2013. LNCS, vol. 8252, pp. 175–191. Springer, Heidelberg (2013). doi:10.1007/978-3-319-02786-9_11

17. Laszka, A., Johnson, B., Schöttle, P., Grossklags, J., Böhme, R.: Secure team composition to Thwart insider threats, cyber-espionage. ACM Trans. Internet Technol. **14**(2–3), 19:1–19:22 (2014)
18. Leslie, D., Sherfield, C., Smart, N.P.: Threshold FlipThem: when the winner does not need to take all. In: Khouzani, M.H.R., Panaousis, E., Theodorakopoulos, G. (eds.) GameSec 2015. LNCS, vol. 9406, pp. 74–92. Springer, Heidelberg (2015). doi:10.1007/978-3-319-25594-1_5
19. Manshaei, M.H., Zhu, Q., Alpcan, T., Bacşar, T., Hubaux, J.-P.: Game theory meets network security and privacy. ACM Comput. Surv. **45**(3), 25:1–25:39 (2013)
20. Microsoft, Microsoft security bulletin. https://technet.microsoft.com/en-us/security/bulletin/dn602597.aspx
21. Nadella, S.: Enterprise security in a mobile-first, cloud-first world (2015). http://news.microsoft.com/security2015/
22. Nochenson, A., Grossklags, J.: A behavioral investigation of the FlipIt game. In: 12th Workshop on the Economics of Information Security (WEIS) (2013)
23. Oracle, Oracle critical patch updates. http://www.oracle.com/technetwork/topics/security/alerts-086861.html
24. Pal, R., Huang, X., Zhang, Y., Natarajan, S., Hui, P.: On security monitoring in SDNS: a strategic outlook. Technical report
25. Pawlick, J., Farhang, S., Zhu, Q.: Flip the cloud: cyber-physical signaling games in the presence of advanced persistent threats. In: Khouzani, M.H.R., Panaousis, E., Theodorakopoulos, G. (eds.) GameSec 2015. LNCS, vol. 9406, pp. 289–308. Springer, Heidelberg (2015). doi:10.1007/978-3-319-25594-1_16
26. Pham, V., Cid, C.: Are we compromised? Modelling security assessment games. In: Grossklags, J., Walrand, J. (eds.) GameSec 2012. LNCS, vol. 7638, pp. 234–247. Springer, Heidelberg (2012). doi:10.1007/978-3-642-34266-0_14
27. Pu, Y., Grossklags, J.: An economic model and simulation results of app adoption decisions on networks with interdependent privacy consequences. In: Poovendran, R., Saad, W. (eds.) GameSec 2014. LNCS, vol. 8840, pp. 246–265. Springer, Heidelberg (2014). doi:10.1007/978-3-319-12601-2_14
28. Radzik, T.: Results and problems in games of timing. In: Fergusons, T.S., Shapleys, L.S., MacQueen, J.B. (eds.) Statistics, Probability and Game Theory: Papers in Honor of David Blackwell. Lecture Notes-Monograph Series, vol. 30, pp. 269–292. Institute of Mathematical Statistics, Hayward (1996)
29. Van Dijk, M., Juels, A., Oprea, A., Rivest, R.: FlipIt: the game of "stealthy takeover". J. Cryptol. **26**(4), 655–713 (2013)
30. Wellman, M.P., Prakash, A.: Empirical game-theoretic analysis of an adaptive cyber-defense scenario (preliminary report). In: Poovendran, R., Saad, W. (eds.) GameSec 2014. LNCS, vol. 8840, pp. 43–58. Springer, Heidelberg (2014). doi:10.1007/978-3-319-12601-2_3
31. Zhang, M., Zheng, Z., Shroff, N.: Stealthy attacks and observable defenses: a game theoretic model under strict resource constraints. In: Proceedings of the IEEE Global Conference on Signal and Information Processing (GlobalSIP), pp. 813–817 (2014)