FlipThem: Modeling Targeted Attacks with FlipIt for Multiple Resources

Aron Laszka[1], Gabor Horvath[2], Mark Felegyhazi[2], and Levente Buttyán[2]

[1] Institute for Software Integrated Systems (ISIS)
Vanderbilt University, Nashville, USA
[2] Department of Networked Systems and Services (HIT)
Budapest University of Technology and Economics (BME), Budapest, Hungary

Abstract. Recent high-profile targeted attacks showed that even the most secure and secluded networks can be compromised by motivated and resourceful attackers, and that such a system compromise may not be immediately detected by the system owner. Researchers at RSA proposed the FlipIt game to study the impact of such stealthy takeovers. In the basic FlipIt game, an attacker and a defender fight over a single resource; in practice, however, systems typically consist of multiple resources that can be targeted. In this paper, we present FlipThem, a generalization of FlipIt to multiple resources. To formulate the players' goals and study their best strategies, we introduce two control models: in the AND model, the attacker has to compromise all resources in order to take over the entire system, while in the OR model, she has to compromise only one. Our analytical and numerical results provide practical recommendations for defenders.

Keywords: FlipIt, game theory, advanced persistent threats, targeted attacks, attacker-defender games

1 Introduction

In recent years, the world witnessed a series of high-profile targeted attacks against various targets [4,19,7,8,2,5,14,13]. These attacks showed that even the most secure and secluded networks can be compromised, and they induced an interesting discussion in the security industry and in the research community alike. An important lesson that the security community can learn from these incidents is that we must revisit some of the most fundamental assumptions which our systems rely on for security. In particular, one must make the assumption that motivated and resourceful attackers can fully compromise a system and gain access to its resources, and this may not be immediately detected by the system owner. The new challenge is to design security mechanisms that minimize the damage that such determined attackers can cause.

In order to help to address this challenge, researchers at RSA – which itself was a victim of a successful targeted attack in 2011 [18] – developed a game-theoretic modeling framework, called FlipIt [3,1]. FlipIt is an attacker-defender game designed to study the problem of stealthy takeover of control over

a critical resource. In `FlipIt`, control over the critical resource is obtained by "flipping" it for a certain cost, and the players receive benefits proportional to the total time that they control the resource. The payoff of each player is, therefore, determined by the difference between the benefit of controlling the resource and the cost of flipping it. Naturally, the goal of the players is to maximize their payoffs.

This is a simple, yet powerful model to study the strategic interaction of attackers and designers of security policies and mechanisms. Moreover, the basic model can be extended in different directions. For instance, in the basic `FlipIt` game, the players flip the resource without being able to observe who was in control before the flip. This model is ideal to study the security of a resource with off-line properties, such as passwords or cryptographic keys. In [16], Pham and Cid extend the basic model by giving the players the option to test if they control the resource before making a move, and use this extended model to study periodic security assessments and their positive effects. In [12,11], Laszka et al. propose and study another variation of the model, in which the defender's moves are non-stealthy, while the attacker's moves are non-instantaneous. Finally, researchers have also studied the `FlipIt` game in behavioral experiments, where human participants played against computerized opponents [15,17,6], which complement the theoretical work by showing the difficulty of finding optimal choices in games of timing.

In this paper, we propose a new generalization of the `FlipIt` game, which, to the best of our knowledge, has not been considered yet in the academic literature. Namely, we extend the basic `FlipIt` model, where the attacker and the defender fight over a single resource, to multiple resources. Accordingly, we call our generalized model the `FlipThem` game. In practice, compromising a system often requires more than attacking just a single component of it. Typically, successful takeovers consist of multiple steps, aiming at gradually escalating the privileges obtained by the attacker until he obtains full administrative access to the system. During this process, the attacker must gain control over a subset of available resources (e.g., he may be required to break a password *and* exploit a software vulnerability in an application). Hence, our model is closer to reality than the original `FlipIt` game, and, as we show in this paper, it is still amenable to mathematical analysis.

More specifically, we make the following contributions in this paper:

- We extend the `FlipIt` game to multiple resources. To formulate the players' goals, we introduce two control models: the AND and the OR control model. In the AND control model, the attacker needs to compromise all resources in order to take over the entire system, whereas in the OR control model, the attacker needs to control at least one resource (out of many available) to take over the entire system. More complex requirements on combinations of resources to be compromised for a successful take-over can be constructed by appropriate combination of these basic control models.
- As a first step to derive good multi-resource `FlipThem` strategies, we introduce two combinations of single-resource `FlipIt` strategies, namely the

independent and the synchronized combinations. In the independent case, the player flips each resource independently of the other resources, whereas in the synchronized case, the player always flips all resources together. We study and compare these two combinations, and derive analytical results for the players' gains.
- As a next step, to represent more complex multi-resource strategies, we introduce the Markov strategy class, where the decision to flip a resource (or set of resources) at a given time depends only on the times elapsed since the previous flips of the resources. We show how the best-response Markov strategy can be computed using a linear program. Using this linear program, we compare various defender strategies based on the resulting benefit for the defender.
- Finally, based on our analytical and numerical results, we provide practical recommendations for defenders. These recommendations can readily be used in practice where the assumptions of the `FlipThem` game apply.

It is important to note that, while the idea of generalizing `FlipIt` to multiple resources may seem straightforward, the exact mathematical treatment of `FlipThem` is not trivial at all. The reason for this is that `FlipThem` is more than just the collection of independent `FlipIt` instances. In general, the attacker and/or defender strategies in `FlipThem` do not handle the different resources independently from each other, and this dependence among the resources results in complex optimization problems when solving the game.

The organization of this paper is the following. In Section 2, we summarize the `FlipIt` game and the most important conclusions drawn in related work. In Section 3, we introduce `FlipThem`, the generalization of `FlipIt` for multiple resources. In Section 4, we show how single-resource `FlipIt` strategies can be combined into multi-resource strategies and compute the players' benefits for various combinations. In Section 5, we introduce the Markov strategy class and show how a best-response Markov strategy can be computed using a linear program. Finally, in Section 6, we discuss the implications of our results and provide practical recommendations for defenders.

2 The `FlipIt` Game

In this section, we summarize the `FlipIt` game and the most important conclusions drawn in related work. It is important to get familiar with the key concepts and notation of the original `FlipIt` game to understand our results for the multiple resources case. Table 1 contains the most important differences in notation between the original `FlipIt` game and our `FlipThem` game. Note that the assumptions of the `FlipIt` game are very different from those of the previous work in the field of game theory for security. For a detailed comparison between `FlipIt` and previous work, we refer the reader to [3].

`FlipIt` [3,1] is a two-player, non-zero-sum game modeling stealthy takeovers, in which both players are trying take control of a single resource. One of the players is called the *defender* (denoted by D), while the other player is called the

Table 1. List of Symbols

Symbol	Description
	`FlipIt`
c^i	player i's flipping cost
β^i	" asymptotic benefit rate
γ^i	" " gain rate
α^i	" " flip rate
Z^i	random variable representing the time since the last flip of player i
	`FlipThem`
N	number of resources
c_r^i	player i's flipping cost for resource r
α_r^i	" asymptotic flip rate for resource r
Z_r^i	rand. var. representing the time since the last flip of player i on resource r

attacker (denoted by A). The game starts at time $t = 0$ and continues indefinitely (that is, $t \to \infty$). In general, time can be both continuous and discrete, with most results being applicable to both cases. At any time instance, player i may choose to take control of the resource by "flipping" it, which costs her c^i. Then, the resource remains under the control of player i until the other player flips it. Consequently, at any given time instance, the resource is controlled by either one or the other player. The interesting aspect of the `FlipIt` game is that neither of the players knows who is in control. As a result, the players occasionally make unnecessary flips (i.e., flip the resource when it is already under their control) since they have to execute their flips "blindly". For an illustration of the game, see Figure 1.

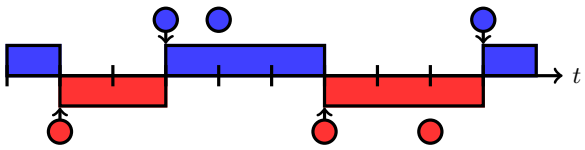

Fig. 1. An illustration of the `FlipIt` game with discrete flip timing. Blue and red disks represent the defender's and attacker's flips. Takeovers, that is, flips changing the player controlling the resource, are indicated by arrows. Blue and red shaded rectangles represent control of the resource by the defender and the attacker, respectively.

The state of the resource is represented by the time-dependent variables C^A and C^D: $C^A(t) = 1$ when the attacker controls the resource, and 0 otherwise; $C^D(t)$ is vice versa (i.e., $C^D(t) = 1 - C^A(t)$). Since the players can (and, as we will soon see, should) employ randomized strategies, both $C^D(t)$ and $C^A(t)$ are random variables. The variables $C^D(t)$ and $C^A(t)$ can be also expressed using the times elapsed since the last flips made by the players as

$$C^D(t) = I_{Z^D(t) \leq Z^A(t)} \qquad \text{and} \qquad C^A(t) = I_{Z^D(t) > Z^A(t)} \,, \tag{1}$$

where Z^i is the time elapsed since the last flip of player i and I is the indicator function.

Player i's *asymptotic gain rate* γ^i is defined as the average fraction of time the resource is controlled by player i. Formally,

$$\gamma^i = \liminf_{t\to\infty} \frac{\int_0^t C^i(\tau)d\tau}{t} \ . \tag{2}$$

Note that player i's asymptotic gain is equal to the probability that the resource is controlled by player i at a random time instance. Formally,

$$\gamma^i = \Pr\left[C^i = 1\right] \ . \tag{3}$$

Player i's *asymptotic flip rate* α^i is defined as the average number of flips made by player i in a unit of time. Formally,

$$\alpha^i = \liminf_{t\to\infty} \frac{n^i(t)}{t} \ , \tag{4}$$

where $n^i(t)$ denotes the number of flips made by player i up to time t. Finally, player i's game-theoretic utility, called player i's *asymptotic benefit* β^i, is defined as the average fraction of time the resource is controlled by the player minus the average cost of flips. Formally,

$$\beta^i = \gamma^i - c^i\alpha^i \ . \tag{5}$$

Since takeovers are assumed to be stealthy in the `FlipIt` game, players do not automatically know when the other player has last moved. However, when a player makes a move (i.e., flips the resource), she might be able to receive some feedback. For example, when an attacker compromises a system, she may learn when the defender last updated the system (that could be attributed as a flip action), and use this information to plan her next move. In [3], three models are introduced for *feedback* received *during the game*:

- Non-adaptive (NA): The player does not receive any feedback when she moves.
- Last move (LM): The player learns the exact time of the other player's last flip.
- Full history (FH): The player learns the complete history of flips made by the other player.

Besides receiving feedback during the game, a player might also be able to receive information before the game starts. For example, an attacker might learn the defender's flip strategy and exploit this knowledge. In [3], two models are introduced for *information* received by a player *before the game starts*:

- Rate of Play (RP): The player knows the asymptotic flip rate α of the other player.
- Knowledge of Strategy (KS): Besides the asymptotic flip rate, the player knows additional information about the other player's strategy. For example, the player may know that the other player employs a renewal process

to generate her flip sequence, and may also know the probability density function of the process. However, it is always assumed that the randomness of the other player's strategy remains secret; consequently, the player cannot know which realization of the renewal process will be used.

In our analysis of defender's strategies in Section 5, we assume a strong attacker model meaning that the attacker always has the Knowledge of Strategy. We assume that the attacker knows everything, except the randomness part of the defender's strategy. This complies with Kerckhoff's principle on security without obscurity.

2.1 Strategies

In this subsection, we summarize the most important strategies and the corresponding results from [3]. For a detailed analysis of these and some other strategies, we refer the interested reader to [3].

In this paper, we focus on non-adaptive strategies, which do not require feedback received by the player during the game. The rationale behind this is that

- defenders rarely know the exact strategies of the attackers (or even the identities of the attackers) in practice; thus, they have to use strategies that do not rely on feedback,
- defenders can choose randomized strategies that schedule their subsequent flips such that even an FH attacker has no more advantage than random guessing (see exponential strategy below), and
- in case of high-importance computer systems, attackers might have limited feedback options if they want to operate stealthily.

A renewal strategy is a non-adaptive strategy in which the time intervals between consecutive flips are generated by a renewal process. More formally, time intervals between consecutive moves are independent and identically distributed random variables, chosen according to a probability density function f. Renewal strategies include (but are not limited to) *periodic strategies* and *non-arithmetic renewal strategies*, which we discuss below.

A player can also choose to drop out of the game (i.e., never flip the resource), which is a rational decision if her expected benefit is less than zero for every strategy choice available to her. This can happen when her opponent's flipping cost is much lower and her opponent can afford to flip the resource extremely fast.

Periodic $\mathcal{P}$: A strategy is periodic if the time intervals between consecutive flips are constant, denoted by δ. It is assumed that a periodic strategy has a *random phase*, that is, the time of the first flip is chosen uniformly at random from $[0, \delta]$. A *periodic strategy with random phase* is characterized by the fixed time interval δ between consecutive flips. It is easy to see that the flip rate of a periodic strategy is $\alpha = \frac{1}{\delta}$. The periodic strategy of rate α is denoted by P_α, and the class of all periodic strategies is denoted by $\mathcal{P}$.

Periodic is probably the strategy most widely used in practice as most systems require passwords, cryptographic keys, etc. to be changed at regular inter-

vals, for example, every thirty days or every three months. In [3], it was shown that the periodic strategy strongly dominates all other renewal strategies if the other player uses a periodic or non-arithmetic renewal strategy. Thus, the periodic strategy is a good choice for an attacker who plays against a non-adaptive (NA) defender.

However, due to its completely deterministic nature[3], the periodic strategy is a very poor choice for defenders who face an attacker observing the last move of the defender (LM attacker). An LM attacker can learn the exact time of the defender's next flip, and schedule her own flip to be immediately after that. Consequently, if flipping costs are of the same order of magnitude, an attacker can keep the resource permanently under her control (with negligible interrupts from the defender). Therefore, a defender facing an LM attacker has two options: if her flipping cost is much lower than that of the attacker, she can flip fast enough to force the attacker to drop out; otherwise, she has to use a randomized strategy, such as the following ones.

Non-arithmetic renewal $\mathcal{R}$: A renewal process is called *non-arithmetic* if there is no positive real number $d > 0$ such that interarrival times are all integer multiples of d. The renewal strategy generated by the non-arithmetic renewal process with probability density function f is denoted by R_f, and the class of all non-arithmetic renewal strategies is denoted by $\mathcal{R}$.

The class of non-arithmetic renewal strategies is very broad as there are an infinite number of possible probability density functions, even for a given flip rate. Of these probability density functions, the exponential is the most important one in the `FlipIt` game.

Exponential $\mathcal{E}$: An *exponential* (or *Poisson*) strategy is a non-arithmetic renewal strategy generated by a Poisson process. Formally, the interarrival times of the process follow an exponential distribution: $f(\tau) = \lambda e^{-\lambda\tau}$, where λ is the parameter characterizing the distribution. The flip rate of this strategy is simply $\alpha = \lambda$. The exponential strategy with rate λ is denoted by E_λ, and the class of all exponential strategies is denoted by $\mathcal{E}$.

The exponential strategy is of key importance, because the exponential distribution is the only *memoryless* continuous probability distribution. The memoryless property means that the conditional probability that we have to wait more than τ_1 time before the next flip, given that the time elapsed since the last flip is τ_2, is independent of τ_2. This implies that, if a defender uses an exponential strategy, an LM (or even an FH) attacker cannot learn *any* information regarding the time of the defender's next flip. Consequently, the exponential strategy is a good choice for a defender facing an LM attacker.

[3] The random phase ensures that an NA opponent cannot determine the flip times of the player; however, if the opponent learns the exact time of at least one flip made by the player, she is able to determine the time of every flip.

3 The FlipThem Game: FlipIt on Multiple Resources

In this section, we generalize the FlipIt game for multiple resources as follows. There are N resources, identified by integer numbers $1, \ldots, N$. Each resource can be flipped individually and, as a result, becomes controlled by the flipping player. The cost of flipping resource r for player i is c_r^i. Each resource has to be flipped individually; i.e., if a player chooses to flip multiple resources at the same time, she still has to pay the flipping cost for each resource that she flips.

The goal of the attacker is to control the system of resources, while the goal of the defender is to prevent the attacker from doing so. The criterion for the attacker controlling the system can be defined in multiple ways, which makes the generalization non-straightforward: as we will later see, different formulations can lead to opposite results. In this paper, we study two elementary control models (see Figure 2 for an illustration):

- All resources [AND]: The attacker controls the system only if she controls *all* resources. Formally,

$$C^A(t) \;=\; Z_1^D(t) > Z_1^A(t) \;\wedge \ldots \wedge\; Z_N^D(t) > Z_N^A(t)\;. \tag{6}$$

  This models scenarios where the attacker has to compromise every resource in order to compromise her target.
- One resource [OR]: The attacker controls the system if she controls *at least one* resource. Formally,

$$C^A(t) \;=\; Z_1^D(t) > Z_1^A(t) \;\vee \ldots \vee\; Z_N^D(t) > Z_N^A(t)\;. \tag{7}$$

  This models scenarios where the attacker only has to compromise a single resource in order to compromise her target.

Similarly to the basic FlipIt game, the players receive benefits proportional to the time that they are controlling the system minus their costs of flipping the resources. More complex control models can be built by combining the AND and OR models in appropriate ways, but the study of that is left for future work.

Notice that, for non-adaptive strategies, the two control models are completely symmetric: the benefit of one player in one model is equivalent to the benefit of the other player in the other model. Consequently, for non-adaptive strategies, it suffices to compute the benefits only in one control model (the AND model in our paper) as the formulas for the other model can be derived readily.

In the following sections, we introduce and study various FlipThem (i.e., multi-resource) strategies, compute the resulting asymptotic benefits, and discuss which strategies should be chosen by the players. First, in Section 4, we study combinations of multiple single-resource strategies. Then, in Section 5, we propose a novel multi-resource strategy class, called the *Markov strategy* class.

4 Combining Single-Resource Strategies

One of the challenges posed by FlipThem lies in the potentially complex structure of the strategies, which can use elaborate rules to exploit the dependence

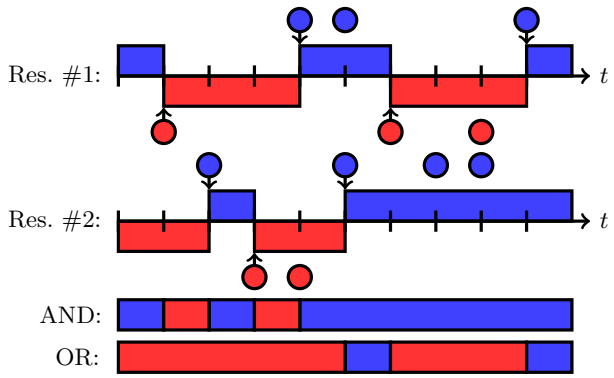

Fig. 2. An illustration of the `FlipThem` game with the AND and OR control models (see Figure 1 for graphical notations).

among the resources. A possible way of finding well-performing, yet analytically tractable multi-resource strategies is to combine multiple single-resource strategies that are known to perform well in the basic `FlipIt` game. In this section, we propose and study two combinations:

- *Independent:* The player flips each resource independently of the other resources. More specifically, the player uses N independent single-resource strategies (i.e., processes), one for each resource, with each one having its own flip rate α_r^i. The asymptotic benefit of a player i using the independent combination is $\beta^i = \gamma^i - \sum_{r=1}^{N} c_r^i \alpha_r^i$.
- *Synchronized:* The player always flips all resources together. More specifically, the player uses only one single-resource strategy (i.e., process) for all of the resources, with a single flip rate α^i. The asymptotic benefit of a player i using the synchronized combination is $\beta^i = \gamma^i - \alpha^i \sum_{r=1}^{N} c_r^i$.

Since the AND and OR control models are symmetric, we only compute the asymptotic gains in the AND model in this paper. Formulas for the asymptotic gains in the OR model can be derived from our results readily. Furthermore, since the defender's asymptotic gain γ^D can be computed from the attacker's asymptotic gain γ^A using the simple formula $\gamma^D = 1 - \gamma^A$, we only compute the asymptotic gain of the attacker.

The proofs of the formulas can be found in the extended online version of this paper [10]. Here, we first show the more general results for the strategy class $\mathcal{R} \cup \mathcal{P}$ (Table 2); then, we analyze the game for the classes $\mathcal{E}$ and $\mathcal{P}$ (Table 3 and Figure 3).

Table 2 shows the attacker's asymptotic gain for various multi-resource strategies chosen by the defender and the attacker. The $\mathcal{R} \cup \mathcal{P}$ in the first and third column indicates that we assume that the players use combinations of either non-arithmetic renewal ($\mathcal{R}$) or periodic ($\mathcal{P}$) single-resource strategies. The combinations used by the defender and the attacker are in the second and fourth

Table 2. Asymptotic Gain for Various Combinations of Single-Resource Strategies

Defender single-resource strategies	comb.	Attacker single-resource strategies	comb.	Attacker's gain γ^A
$\mathcal{R} \cup \mathcal{P}$	ind.	$\mathcal{R} \cup \mathcal{P}$	ind.	$\prod_{r=1}^{N} \int_0^\infty f_{Z_r^D}(z_r) F_{Z_r^A}(z_r) dz_r$
			syn.	$\int_0^\infty \prod_{r=1}^{N} \left(1 - F_{Z_r^D}(z)\right) f_{Z^A}(z) dz$
	syn.			$\int_0^\infty f_{Z^D}(z) F_{Z^A}(z) dz$
			ind.	$\int_0^\infty \prod_{r=1}^{N} F_{Z_r^A}(z) f_{Z^D}(z) dz$

columns, respectively. Finally, the attacker's gain γ^A for the given combinations is in the fifth column.

To express the attacker's gain, we use a notion similar to that of the basic `FlipIt` game. We let Z_r^i be the random variable representing the time elapsed since player i's last flip on resource r (we omit the index r and denote it by simply Z^i if the player uses a synchronized strategy). We denote the cumulative distribution and density functions of Z_r^i by $F_{Z_r^i}(z)$ and $f_{Z_r^i}(z)$. These functions can easily be computed from the generating distribution of any non-arithmetic renewal strategy (see the extended online version [10]).

It is noteworthy that, when both players use the synchronized combination, the game is equivalent to the basic `FlipIt` game (with $c^i = \sum_r c_r^i$): each player uses only one single-resource (i.e., basic `FlipIt`) strategy, and the state of all resources is the same as they are always flipped together. Consequently, the formula for the attacker's gain is identical to the corresponding formula in [3].

Table 3 shows the attacker's asymptotic gain for various combinations of exponential and periodic strategies. We selected these single-resource strategies because they are known to be optimal in some respect (see Section 2). The table is similar to Table 2, except that the synchronized defender against independent attacker case is omitted to keep the table simple (it can be found in the extended version of this paper [10]) and because it is not a good strategy for either of the players.

The table shows that the independent combination is generally better than the synchronized one for the defender, as her flip rates are added together in the former. This can be explained by the nature of the AND control model: since the defender only needs to control at least one resource, her best strategy is to flip one resource at a time. This forces the attacker to frequently flip all resources back as she cannot know which resources were flipped by the defender (since the exponential process is memoryless).

The formulas also suggest that the attacker should choose the synchronized combination over the independent one. When both players use exponential single-resource strategies, the attacker's gain decays exponentially as the number of resources increases ($\sim k^{-N}$) if she uses the independent combination, but only

Table 3. Asymptotic Gain for Various Combinations of Exponential and Periodic Strategies

Defender single-resource strategy	comb.	Attacker single-resource strategy	comb.	Attacker's gain γ^A
$\mathcal{E}$	ind.	$\mathcal{E}$	ind.	$\prod_{r=1}^{N} \frac{\alpha_r^A}{\alpha_r^A + \alpha_r^D}$
			syn.	$\frac{\alpha^A}{\alpha^A + \sum_{r=1}^{N} \alpha_r^D}$
	syn.			$\frac{\alpha^A}{\alpha^A + \alpha^D}$
	ind.	$\mathcal{P}$	ind.	$\prod_{r=1}^{N} \frac{\alpha_r^A}{\alpha_r^D} \left(1 - e^{-\frac{\alpha_r^D}{\alpha_r^A}}\right)$
			syn.	$\frac{\alpha^A}{\sum_{r=1}^{N} \alpha_r^D} \left(1 - e^{-\frac{\sum_{r=1}^{N} \alpha_r^D}{\alpha^A}}\right)$
	syn.			$\frac{\alpha^A}{\alpha^D} \left(1 - e^{-\frac{\alpha^D}{\alpha^A}}\right)$

according to a power law ($\sim N^{-k}$) if she uses the synchronized one (given that flip rates stay the same). When the attacker uses the periodic single-resource strategy, the relationship between the number of resources and the attacker's gain is more complicated, but similar.

Figure 3 shows the attacker's asymptotic gain as a function of the number of resources for various combinations of exponential and periodic strategies. The plotted pairs of combinations are the following: both players use independent strategies (solid line —), the attacker uses synchronized while the defender uses independent strategy (dashed line - -), and both players use synchronized strategies (dotted line ····). The flip rates are assumed to be uniform, i.e., $\alpha^A = \alpha_r^A = \alpha^D = \alpha_r^D = 1,\ r = 1, \ldots, N$.

The figure shows that, for the given single-resource classes and parameters, the synchronized combination strongly dominates the independent one for the attacker. Again, this can be explained by the nature of the AND control model: since the attacker needs to control all resources, it makes sense to flip them all together. Otherwise, the probability that all resources become controlled by the attacker is very low. However, by using the synchronized combination, the attacker loses the freedom of choosing the flipping rate for each resource independently. Thus, when the heterogeneity of the attacker's flipping costs is very high, the independent combination may outperform the synchronized one.

The figure also supports our finding that the independent combination strongly dominates the synchronized one for the defender. Since a player has complete

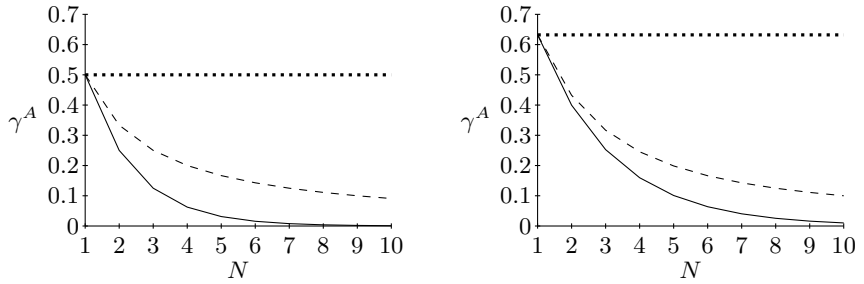

(a) Both players use exponential strategies.

(b) Defender uses exponential, attacker uses periodic strategy.

Fig. 3. The attacker's asymptotic gain as a function of the number of resources for various combinations of exponential and periodic strategies. Plotted pairs of combination are: both players use independent strategies (solid line), attacker uses synchronized strategies while defender uses independent strategies (dashed line), and both players use synchronized strategies (dotted line). In this figure, the flip rates are assumed to be uniform, i.e., $\alpha^A = \alpha_r^A = \alpha^D = \alpha_r^D = 1,\ r = 1, \ldots, N$.

freedom in choosing her flip rates in the independent combination, this combination is better for the defender even for very heterogeneous flipping costs.

Finally, by comparing Subfigures 3a and 3b, we conclude that the periodic strategy dominates the exponential strategy as the attacker's gain is higher when she chooses the former.

5 The Markov Strategy Class

In the previous section, we studied how single-resource strategies can be combined into multi-resource strategies. However, such combinations represent only a tiny fraction of the actual multi-resource strategy space as there are an infinite number of multi-resource strategies that cannot be represented by such simple combinations. For example, a defender might choose to flip one resource periodically, then wait for a time interval chosen according to an exponential distribution, and then flip another resource. To model such complex multi-resource strategies, in this section, we introduce the *Markov* strategy class.

For the clarity of presentation, we derive results for two resources, yet the approach is applicable for any number of resources. Furthermore, as opposed to the basic model, we are going to use discrete time in this section. Note that the discrete time model can be very realistic as players typically do not flip their resources at arbitrary times. Examples are the change of passwords, cryptographic keys or the application of software updates. We denote the duration of a time step by Δ. Finally, we define the time-dependent age functions as follows. The random variables representing the number of time steps elapsed since the last

flip of resource r by the attacker and the defender at time k are denoted by $Z_r^A(k)$ and $Z_r^D(k)$, respectively.

In the case of two resources, the attacker can perform one of the following actions in a given time slot:

- she does not flip any of the resources,
- she flips one of the resources,
- or she flips both resources.

If the decision which action to choose depends only on the times elapsed since the previous flips of the resources, then $\{(Z_1^A(k),\ Z_2^A(k)),\ k = 0, 1, \dots\}$ defines a Markov process. In this case, the behavior of the attacker can be characterized by the following joint distributions corresponding to the events that can happen in two consecutive time steps:

$$\begin{aligned}
p_{i,j}^{(0)} &= \Pr\left[Z_1^A(k) = i,\ Z_2^A(k) = j,\ Z_1^A(k+1) = i+1,\ Z_2^A(k+1) = j+1\right] \\
p_{i,j}^{(1)} &= \Pr\left[Z_1^A(k) = i,\ Z_2^A(k) = j,\ Z_1^A(k+1) = 0,\ Z_2^A(k+1) = j+1\right] \\
p_{i,j}^{(2)} &= \Pr\left[Z_1^A(k) = i,\ Z_2^A(k) = j,\ Z_1^A(k+1) = i+1,\ Z_2^A(k+1) = 0\right] \\
p_{i,j}^{(1,2)} &= \Pr\left[Z_1^A(k) = i,\ Z_2^A(k) = j,\ Z_1^A(k+1) = 0,\ Z_2^A(k+1) = 0\right]\ ,
\end{aligned}$$

where $p_{i,j}^{(0)}$ is the probability that nothing is flipped in the next time step, $p_{i,j}^{(1)}$ ($p_{i,j}^{(2)}$) is the probability that only resource 1 (or 2) is flipped, while $p_{i,j}^{(1,2)}$ is the probability of both resources being flipped in the next time step.

We denote by M_p the Markov strategy generated by a Markov process with event probabilities $p = \{p_{i,j}^{(0)}, p_{i,j}^{(1)}, p_{i,j}^{(2)}, p_{i,j}^{(1,2)}$ for $i, j = 0, 1, \dots\}$, and by $\mathcal{M}$ the class of all Markov strategies. That is,

$$\mathcal{M} = \{M_p \mid p \text{ is a set of event probabilities}\}\ . \tag{8}$$

5.1 Linear Programming Solution

With these definitions and notations, we can define a linear program to determine the optimal probabilities $p_{i,j}^{(\bullet)}$. However, since linear programming problems can only be solved with a finite number of variables (in the general case), we have to restrict the game to a finite time horizon. The last time step we take into consideration is denoted by T.

The attacker wants to maximize her benefit β^A, which is composed of the asymptotic gain and the cost of the flips against both resources as

$$\begin{aligned}
\beta^A = \max_p \Bigg\{ & \underbrace{\sum_{i=0}^{T}\sum_{j=0}^{T} q_{i,j} \Pr\left[Z_1^D > i, Z_2^D > j\right]}_{\gamma^A} \\
& - \underbrace{c_1^A \left(\sum_{i=0}^{T}\sum_{j=0}^{T} p_{i,j}^{(1)} + p_{i,j}^{(1,2)}\right)\frac{1}{\Delta}}_{\alpha_1^A} - \underbrace{c_2^A \left(\sum_{i=0}^{T}\sum_{j=0}^{T} p_{i,j}^{(2)} + p_{i,j}^{(1,2)}\right)\frac{1}{\Delta}}_{\alpha_2^A} \Bigg\}\ ,
\end{aligned} \tag{9}$$

where $q_{i,j}$ is the probability that the number of time steps since the attacker's last flips of resource 1 and 2 are i and j, respectively. This probability can be expressed easily as $q_{i,j} = p^{(0)}_{i,j} + p^{(1)}_{i,j} + p^{(2)}_{i,j} + p^{(1,2)}_{i,j}$; thus, the objective function given by (9) defines a linear relation with respect to $p^{(\bullet)}_{i,j}$.

As variables $p^{(\bullet)}_{i,j}$ must be valid probabilities, we need to apply the inequality constraints $p^{(0)}_{i,j} \geq 0, p^{(1)}_{i,j} \geq 0, p^{(2)}_{i,j} \geq 0, p^{(12)}_{i,j} \geq 0$; and we also need to ensure that the probabilities sum up to 1, that is, $\sum_{i=0}^{T} \sum_{j=0}^{T} p^{(0)}_{i,j} + p^{(1)}_{i,j} + p^{(2)}_{i,j} + p^{(1,2)}_{i,j} = 1$.

Further equality constraints are required to define the possible state transitions, yielding

$$q_{i,j} = p^{(0)}_{i-1,j-1} \quad \text{for } i > 0, j > 0, \qquad q_{0,0} = \sum_{i=0}^{T} \sum_{j=0}^{T} p^{(1,2)}_{i,j} \ ,$$

$$q_{0,j} = \sum_{i=0}^{T} p^{(1)}_{i,j-1} \quad \text{for } j > 0, \qquad q_{i,0} = \sum_{j=0}^{T} p^{(2)}_{i-1,j} \quad \text{for } i > 0, \tag{10}$$

with $q_{i,j}$ given above.

Finally, we require that a resource is always flipped in the next time step if its age has reached the maximum age T:

$$p^{(0)}_{i,j} = 0 \ \text{ for } i = T \text{ or } j = T, \quad p^{(1)}_{i,j} = 0 \ \text{ for } j = T, \quad p^{(2)}_{i,j} = 0 \ \text{ for } i = T. \tag{11}$$

5.2 Results

The linear program defined above answers several questions regarding the `Flip-Them` game, including the following:

- What is the attacker's optimal strategy against a given defender strategy?
- What are the optimal flip rates maximizing the defender's benefit if the attacker always plays an optimal strategy?
- What is the Nash equilibrium of this game?

Solving the optimization problem using a linear programming based approach poses some challenges. In particular, the length of the time horizon T is limited by the capabilities of the linear program solver. For our examples, we used the built-in solver of MATLAB with $T = 30$ (resulting in 900 variables in case of two resources).[4] Note that the number of variables increases polynomially in the length of the time horizon and exponentially in the number of resources. Using custom software, the analysis can be extended to much larger values of T; however, the results we obtained with MATLAB are already revealing and useful.

In the rest of this section, we consider several numerical examples to demonstrate the usefulness of the model. In each of these examples, the attacker is assumed to be non-adaptive (NA), but she is assumed to know the strategy of the

[4] This can model, for example, the key update policy of a company over a duration of 2.5 years assuming that updates are defined by the granularity of a month.

defender (KS). The defender, however, has no information about the attacker. For the definitions and rationale behind these modeling choices, see Section 2.

Optimal attack against a given defender strategy In this example, the defender flips the resources according to independent Poisson processes with parameters $\alpha_1^D = 1$ and $\alpha_2^D = 3$. The joint age function is then $\Pr\left[Z_1^D > i, Z_2^D > j\right] = e^{-\alpha_1^D i\Delta - \alpha_2^D j\Delta}$. The attacker's flip costs are $c_1^A = 0.1$ and $c_2^A = 0.05$. The discrete problem is solved with $T = 30$ and $\Delta = 0.03$.

At this point, we take the opportunity to introduce the conditional state transition probability matrices $\boldsymbol{P^{(0)}}, \boldsymbol{P^{(1)}}, \boldsymbol{P^{(2)}}$ and $\boldsymbol{P^{(1,2)}}$ that help to visualize and understand the strategy of the attacker. The entries of these matrices are

$$[\boldsymbol{P^{(\bullet)}}]_{i,j} = \frac{p_{i,j}^{(\bullet)}}{p_{i,j}^{(0)} + p_{i,j}^{(1)} + p_{i,j}^{(2)} + p_{i,j}^{(1,2)}} \ . \tag{12}$$

To simulate an attack, one has to follow the state of the attacker given by positions i, j in the matrices. In state (i, j), no flips occur with probability $[\boldsymbol{P^{(0)}}]_{i,j}$, and the next state of the attacker is $(i + 1, j + 1)$. With probability $[\boldsymbol{P^{(1)}}]_{i,j}$ (or $[\boldsymbol{P^{(2)}}]_{i,j}$), only resource 1 (or resource 2) is flipped in the next time step, and the next state of the system is $(0, j + 1)$ (or $(i + 1, 0)$). Finally, both resources are flipped in the next time step with probability $[\boldsymbol{P^{(1,2)}}]_{i,j}$, followed by a jump to state $(0, 0)$.

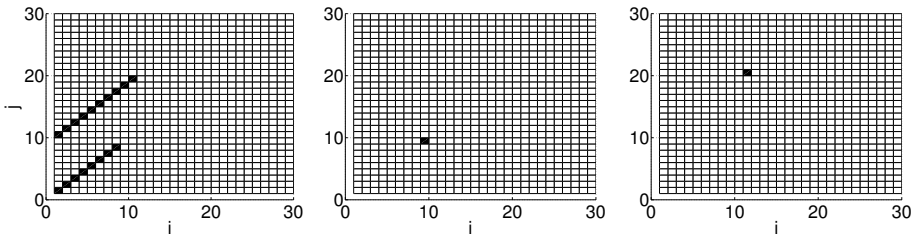

Fig. 4. Optimal attack strategy against two resources flipped according to independent Poisson processes.

By solving the linear program, we obtain the optimal strategy of the attacker, represented by the matrices depicted in Figure 4. In this particular example, the entries of all four matrices are all either 0 (represented by white squares) or 1 (black squares). Matrix $\boldsymbol{P^{(1)}}$ is not depicted as it has only 0 entries. By following the attacker's strategy in the above described manner, we have that she first waits 9 time steps (black squares on the diagonal of $[\boldsymbol{P^{(0)}}]$), then flips one resource (black square in $\boldsymbol{P^{(2)}}(9, 9)$), waits another 10 time steps, and finally flips both resources (black square in $\boldsymbol{P^{(1,2)}}(20, 11)$).

Thus, based on the matrices, a "periodic" attack can be identified with a period of $\delta = 20$. The resources are not flipped in a synchronized manner. Resource 2 is flipped at the 9th time step from the beginning of the period, while both resources are flipped at the end of the period.

If the defender flips both resources according to independent periodic strategies, the joint age process is given by $\Pr\left[Z_1^D > i, Z_2^D > j\right] = (1 - \alpha_1^D i\Delta)(1 - \alpha_2^D j\Delta)$, if $i\Delta < 1/\alpha_1^D,\ j\Delta < 1/\alpha_2^D$, and $\Pr\left[Z_1^D > i, Z_2^D > j\right] = 0$ otherwise. When keeping all parameters the same as before, the optimal strategy of the attacker is more complex in this case (see Figure 5). The period of her strategy is $\delta = 22$ now, and she flips solely resource 2 at time steps 6 and 13, while she flips both resources at time step 22, which also marks the end of her period.

It is noteworthy that the attacker's benefit is 0.265 in the Poisson case, but only 0.047 in the periodic case, which means that the periodic defense is less economical to attack (given, of course, that the attacker has no knowledge on the last move of the defender, thus it is of type NA).

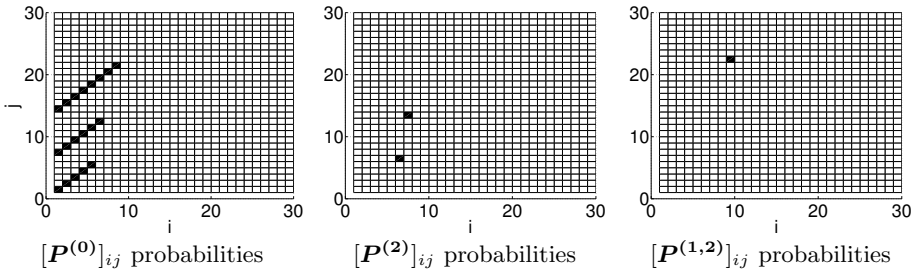

Fig. 5. Optimal attack strategy against two resources flipped according to independent periodic strategies.

Defender's optimal flip rates The linear program can also be used to find the defender's optimal flip rates given that the attacker always uses her best-response strategy. Notice that we do not calculate a Nash equilibrium here, thus the defender does not have to take the strategy of the attacker into consideration.

First, consider the case when the defender flips her resources according to independent Poisson processes. Assume that the attacker's flipping costs are $c_1^A = 0.1$ and $c_2^A = 0.2$. We solved the linear program with various combinations of α_1^D and α_2^D, and with two different settings for the parameters c_1^D and c_2^D. The results are shown in Figure 6. As the benefit of the attacker is the subject of optimization in the linear program, the corresponding plot is obviously smooth, and gives higher values for lower flip rates of the defender. The corresponding gain rates (which are not plotted due to the lack of space), however, are not smooth. As the defender's benefit is directly related to the attacker's gain rate, the plots of the defender's benefit are not smooth either. The maximum benefit for the defender is 0.222, obtained at $\alpha_1^D = 0.8, \alpha_2^D = 0.7$.

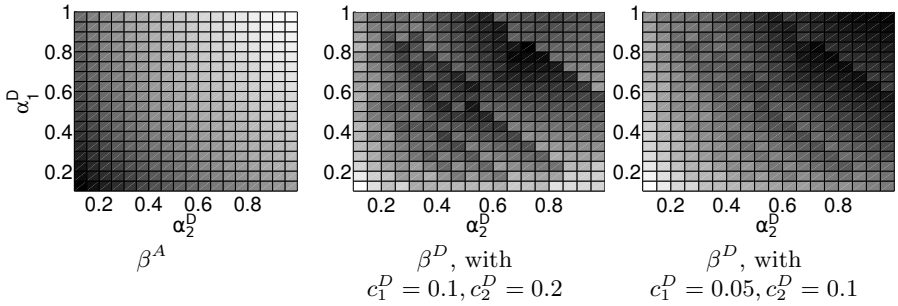

Fig. 6. Benefits of the attacker (β^A) and defender (β^D) for various flip rates of the defender (Poisson case). Darker shades of gray indicate higher benefit.

If the defender flips the resources according to independent periodic strategies, higher flip rates are required to maximize her benefit. The corresponding results are depicted in Figure 7: the optimal flip rates are $\alpha_1^D = 0.9, \alpha_2^D = 1.2$, and her benefit $\beta^D = 0.61595$ is higher compared to the Poisson case. Observe that the attacker always drops out for higher flip rates, which is indicated by the white area on the plot of her benefit and also by the sharp line appearing on the plots of the defender's benefit.

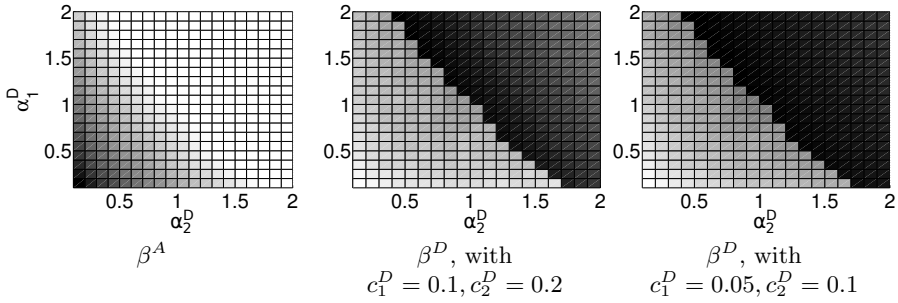

Fig. 7. Benefits of the attacker (β^A) and defender (β^D) for various flip rates of the defender (periodic case). Darker shades of indicate higher benefit.

Optimal flip for a fixed budget In this example, we assume that the defender has a fixed budget, and we are looking for the flip rates maximizing her benefit. By a fixed budget, we mean that the defender spends a fixed amount B on average on flipping her resources, thus $B = c_1^D \alpha_1^D + c_2^D \alpha_2^D$ is fixed, while the ratio of the flip rates $R = \alpha_1^D / \alpha_2^D$ is subject of optimization. Notice that B and R determine the flip rates uniquely as

$$\alpha_1^D = \frac{RB}{c_1^D R + c_2^D}, \qquad \alpha_2^D = \frac{B}{c_1^D R + c_2^D} \ . \tag{13}$$

The flip costs of the attacker and the defender are set to $c_1^A = 0.1, c_2^A = 0.05$ and $c_1^D = c_2^D = 0.001$, the total cost is $B = 0.004$, and we apply a finer discretization in this example with $T = 90$ and $\Delta = 0.01$.

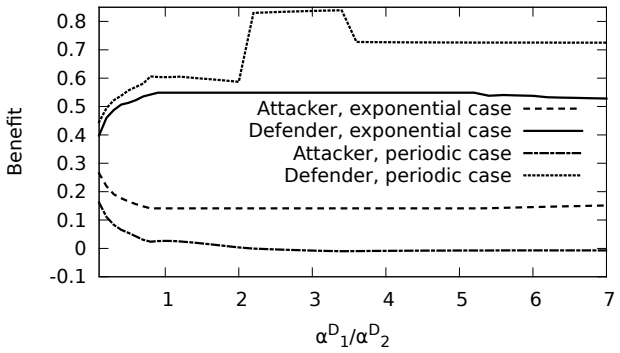

Fig. 8. Benefits of the attacker and the defender as functions of the ratio between the flip rates for the resources.

Figure 8 depicts the players' benefits assuming that the attacker always flips according to her best-response Markov strategy. The optimal ratio R (from the defenders point of view) is 3.4 when she flips her resources periodically, and it is between 0.9 and 5.2 when she uses exponential strategies. By looking at the results closer, we find that, when the defender chooses an optimal ratio, the attacker uses a synchronized periodic attack against the resources in both cases.

Nash equilibrium The proposed linear program can be applied to calculate the optimal strategies of both the defender and the attacker. We can thus use a simple iterative algorithm to find a Nash equilibrium of the game. This algorithm starts with assigning a random strategy to the defender, followed by the alternating optimizations of the attacker's and the defender's strategies. In practice, however, we found that this algorithm does not converge in the vast majority of the cases, but it starts oscillating after a number of iterations, suggesting that no Nash equilibrium exists.

6 Concluding Remarks

Extending the `FlipIt` game to multiple resources requires modeling the players' goals as functions of the compromised resources. We selected the two most intuitive choices, namely the AND and OR control models, to represent the gains derived from controlling the resources. From the attacker's viewpoint, the AND control model represents the case when all resources need to be compromised

to get access to the system. This is similar to the *total effort* model of security interdependence in the state-of-the-art [9,20]. The OR control model represents the case when the compromise of a single resource suffices to get access. This second choice relates to the *weakest link* model of interdependence [9,20].

We proposed two major classes of multi-resource strategies: combinations of single-resource strategies (independent processes and synchronized processes) and the Markov strategy class. Based on our result, we can formulate a set of recommendations for the defender. These recommendations can be readily used in practice where the assumptions of the `FlipThem` game apply, for example, when defining the key update strategy for a security infrastructure.

- In the AND control model, we found that the defender should use independent flipping strategies. In practice, this means that cryptographic keys should not be updated at the same time, but rather independently.
- On the other hand, in the OR control model, the defender should use synchronized flipping strategies. In practice, this means updating cryptographic keys synchronously. However, the defender needs to pay attention to the cost of updating keys in the OR control model. If these costs are very heterogeneous, the key update processes should remain synchronized, but with different update rates across the keys.
- If the attacker is non-adaptive, then the periodic defender strategy is a good choice according to our numerical results.[5] Periodic strategies have multiple advantageous properties such as higher benefits for the defender, robustness to optimization errors and ease of implementation in practice. However, periodic strategies perform poorly against an LM attacker [3]. Thus, the defender needs to carefully assess the potential information available to an attacker when choosing her strategy.
- Surprisingly, the defender's benefit is not a smooth or monotonic function of her flip rates, which makes optimization difficult in practice. Our numerical results imply that this observation holds for any combination of the periodic and the exponential strategy classes. The major reason behind this non-monotonous property is that, as the defender's flip rate reaches a threshold, the attacker drops out of the game. In realistic cases, the defender's flipping cost is much lower than the attacker's flipping cost, which causes the attacker to drop out.

Acknowledgment This work is supported in part by the National Science Foundation (CNS-1238959).

References

1. Bowers, K., van Dijk, M., Griffin, R., Juels, A., Oprea, A., Rivest, R., Triandopoulos, N.: Defending against the unknown enemy: Applying FlipIt to system security. In: Proceedings of the 3rd Conference on the Decision and Game Theory for Security (GameSec). pp. 248–263 (2012)

[5] This complies with the results of the basic `FlipIt` game for a single resource in [3].

2. cnet.com: Comodo hack may reshape browser security. http://news.cnet.com/8301-31921_3-20050255-281.html (Apr 4 2011)
3. van Dijk, M., Juels, A., Oprea, A., Rivest, R.L.: FlipIt: The game of "stealthy takeover". Cryptology ePrint Archive, Report 2012/103 (2012)
4. Falliere, N., Murchu, L.O., Chien, E.: W32.Stuxnet Dossier. http://www.symantec.com/connect/blogs/w32stuxnet-dossier (February 2011)
5. Finkle, J., Shalal-Esa, A.: Hackers breached U.S. defense contractors. http://www.reuters.com/article/2011/05/27/us-usa-defense-hackers-idUSTRE74Q6VY20110527 (May 27 2011)
6. 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) (2014)
7. Kaspersky Lab: Flame...the latest cyber-attack. http://www.kaspersky.com/flame (May 2012)
8. Kaspersky Lab: The MiniDuke Mystery: PDF 0-day Government Spy Assembler 0x29A Micro Backdoor. http://www.securelist.com/en/blog/208194129/The_MiniDuke_Mystery_PDF_0_day_Government_Spy_Assembler_0x29A_Micro_Backdoor (February 2013)
9. Laszka, A., Felegyhazi, M., Buttyán, L.: A survey of interdependent security games. Tech. Rep. CRYSYS-TR-2012-11-15, CrySyS Lab, Budapest University of Technology and Economics (Nov 2012)
10. Laszka, A., Horvath, G., Felegyhazi, M., Buttyán, L.: FlipThem: Modeling targeted attacks with FlipIt for multiple resources (extended version). http://www.crysys.hu/%7Elaszka/papers/laszka2014flipthem.pdf
11. Laszka, A., Johnson, B., Grossklags, J.: Mitigating covert compromises: A game-theoretic model of targeted and non-targeted covert attacks. In: Proceedings of the 9th Conference on Web and Internet Economics (WINE). pp. 319–332 (2013)
12. Laszka, A., Johnson, B., Grossklags, J.: Mitigation of targeted and non-targeted covert attacks as a timing game. In: Proceedings of the 4th Conference on Decision and Game Theory for Security (GameSec). pp. 175–191 (2013)
13. Mandiant: APT1: Exposing one of China's cyber espionage units. http://www.mandiant.com/apt1 (Feb 18 2013)
14. Menn, J.: Key Internet operator VeriSign hit by hackers. http://www.reuters.com/article/2012/02/02/us-hacking-verisign-idUSTRE8110Z820120202 (Feb 2 2012)
15. Nochenson, A., Grossklags, J.: A behavioral investigation of the FlipIt game. In: Proceedings of the 12th Workshop on the Economics of Information Security (WEIS) (2013)
16. Pham, V., Cid, C.: Are we compromised? Modelling security assessment games. In: Proceedings of the 3rd Conference on the Decision and Game Theory for Security (GameSec). pp. 234–247 (2012)
17. Reitter, D., Grossklags, J., Nochenson, A.: Risk-seeking in a continuous game of timing. In: Proceedings of the 13th International Conference on Cognitive Modeling (ICCM). pp. 397–403 (2013)
18. Rivner, U.: Anatomy of an attack. http://blogs.rsa.com/anatomy-of-an-attack/ (Apr 2011)
19. Symantec Security Response: W32.Duqu: The Precursor to the Next Stuxnet. http://www.symantec.com/connect/w32_duqu_precursor_next_stuxnet (October 18 2011)
20. Varian, H.: System reliability and free riding. In: Economics of Information Security, pp. 1–15. Springer (2004)