It's Time to Migrate! A Game-Theoretic Framework for Protecting a Multi-tenant Cloud against Collocation Attacks

Ahmed H. Anwar and George Atia
Department of Electrical and Computer Engineering
University of Central Florida
Orlando, FL 32816
a.h.anwar@knights.ucf.edu, george.atia@ucf.edu

Mina Guirguis
Department of Computer Science
Texas State University
San Marcos, TX 78666
msg@txstate.edu

***Abstract*—We present a novel game-theoretic framework for the Virtual Machine (VM) migration timing problem. In a multi-tenant cloud, a number of VMs are collocated on the same physical machine. This increases the risk of a malicious VM performing side-channel attacks and leaking sensitive information. To this end, this paper develops and analyzes a game-theoretic framework for the timing problem in which the cloud provider decides *when* to migrate a VM to a different physical machine to reduce the risk of being compromised by a collocated malicious VM. The adversary decides the rate at which she launches new VMs to collocate with the victim VMs. Our formulation captures a data leakage model in which the cost incurred by the cloud provider depends on the duration of collocation as well as the overhead in migration. We establish sufficient conditions for the existence of Nash equilibria for general cost functions, as well as for specific instantiations, and characterize the best response for both players. Our theoretical findings are corroborated with extensive numerical results in various settings.**

***Keywords*-Game theory; Cloud migration; Cloud security;**

I. Introduction and related work

One of the main characteristics of the Cloud that allows scalable and cost-effective operation is multi-tenancy. Multi-tenancy is achieved through virtualization to enable cloud providers to host multiple virtual machines (VMs) on the same physical machine while providing isolation between them. Recent attacks, however, have been shown to bypass such isolation [1]. A malicious VM collocating on the same physical machine with a victim VM can seek unauthorized access to sensitive and private data and/or intellectual property, or can render some of its computational functionality unusable.

This has prompted cloud providers to develop various strategies for VM placement, migration and reconfiguration to mitigate some of these attacks. Moving target defense (MTD) strategies aim to dynamically shift the attack surface, making it more difficult for attackers to launch potent attacks [2]. When developing an MTD strategy, two main questions generally arise: *which* targets should be moved and *when* should they be moved? The answer to these questions is highly-dependent on the context of the problem and the nature of the attack. For example, if an attacker contemplates inferring the underlying topology of the cloud, then the connectivity between machines is the target that should be changed over time. In a different setting, if the attacker is interested in cracking the system credentials that protect the users' databases, then the keys are the target that should be constantly reconfigured (i.e., moved). In this paper, we consider collocation attacks whereby an attacker can leak sensitive data from a targeted victim by running a VM on the same physical node (e.g., through launching a side-channel attack). Thus, for securing such system, VMs should be periodically migrated (i.e., moved to a different physical machine). This paper is primarily focused on the second question, that is, when to move the identified targets.

In the MTD literature, this question is usually referred to as the timing problem of the MTD strategy. In this paper, we study this question using a game-theoretic framework seeking an understanding of the interplay of the actions of both the cloud provider (i.e., the defender) and the adversary, [3], [4]. In our formulation, the adversary seeks to prolong the collocation time with the victim VMs to maximize information leakage. Since the adversary has no guarantees to be successfully collocated on the same physical machine with the victim – since different cloud providers implement different placement algorithms according to different criteria that the attacker has no control over – her best-effort would be to increase the number of VMs to launch (which is a cost metric we capture). The adversary can then check after being placed whether she had a successful collocation or not [5]. The cloud provider, on the other hand, seeks to migrate VMs between physical machines to minimize the collocation times between VMs. VM live migration, while efficient, is not free [6] and thus the question so as to when to migrate is crucial in order to mitigate the collocation attack threats while not burdening the system with a large overhead that may not be justified.

Cross-VM side-channel attacks and their impact have been the subject of various recent studies (e.g., [5], [7], [8], [9], [10], [11], [12]). The authors in [13] showed that by controlling the placement process, a defense mechanism can

mitigate the effect of cross-VM attacks through reducing the co-run probability between users. The approach, however, is only effective in the case of time-sensitive attacks and when the number of assigned virtual CPUs is substantial. The use of game theory has largely focused on the VM allocation problem in the presence of adversaries [6], [14], [15], [16], [17], [18], [19], [20]. A common assumption in such formulations is that the adversary is known. This assumption does not hold in practice. Additionally, existing formulations do not consider the timing problem, which is a critical one for the defender wishing to migrate VMs for security.

Contributions: While VM migration strategies have been proposed as defense mechanisms against collocation attacks in various studies, such work focused on the VM assignment problem (mapping VMs to physical nodes) as a single player scheduling problem. In this paper, however, we consider the timing problem of the MTD as a game between the attacker and the cloud provider. Our work contributes to the theory of timing games [21], which is largely unexplored in cloud computing settings. We leverage the results of the leakage model in the FlipIt game considered previously in [22], [23], [24], [25], [26], [27] to develop a novel formulation to study the VM collocation problem in an extended FlipIt game-theoretic framework. To the best of our knowledge, this is the first work to investigate the following aspects of the timing games.

- We provide a new game-theoretic formulation for the VM collocation timing problem.
- Unlike [14], [15], [17], we do not assume the defender has prior knowledge of the exact location of the attacker, thereby allowing for realistic threat and defense models. The defender has to periodically migrate the VMs to protect against malicious collocating users.
- We analytically characterize the Nash equilibrium (NE) for the studied game model and derive sufficient existence conditions.
- We provide extensive numerical experiments to support our theoretical findings and compare our proposed defense policies against other defense policies. In our numerical evaluation, we consider several reward functions to reflect the severity of the attack and different degrees of information leakage.

This paper is organized as follows. In Section II, we provide the system model and game formulation. In Section III, we provide theoretical analysis and establish existence conditions of NE for the formulated game. Our numerical results are presented in Section IV and we conclude the paper in Section V.

II. System Model

A. The cloud

We model the cloud as a set of physical machines whereby each machine can host a number of VMs from different users. The cloud provider uses a placement strategy to initially assign VMs to physical machines. The details of the placement strategy do not affect our analysis and we assume that the adversary (or any user) has no control over it. We assume the adversary is interested in targeting a set of victim VMs by collocating with them on the same physical machine. We study the interaction between the cloud provider (defender) and the adversary through a game-theoretic framework in which the rewards are time-dependent. In particular, the defender's strategy is to choose the time to re-assign VMs to different machines to defend against collocation attacks. The adversary, on the other hand, chooses an attack rate to launch more VMs to increase her collocation duration to maximize information leakage from her victims as described in Fig. 1. We define the game next.

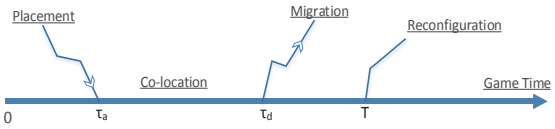

Figure 1. System model illustration

B. The game

A game is defined as a tuple $\Gamma(\mathcal{P}, \mathcal{A}, \mathcal{U})$, where $\mathcal{P}$ is the set of players (here, $\mathcal{P} = \{1, 2\}$, denoting the defender (player 1) and the adversary (player 2)), $\mathcal{A} = \mathcal{A}_d \times \mathcal{A}_a$ is the action space for the defender and adversary, and $\mathcal{U} = \{u_d, u_a\}$ is the reward function, $\mathcal{U} : \mathcal{A} \rightarrow \mathbb{R}^2$.

1) Defender's action space: Since we are investigating the timing factor of the game, the defender is assumed to control the re-allocation period. Let $\tau_d \in \mathcal{A}_d$ denote the time at which the defender migrates a running VM to a new physical node, such that $\mathcal{A}_d = [\tau_{\min}, T]$, where T is a system parameter at which the credentials are reset and $\tau_{\min}$ the smallest reconfiguration time. Since we assume a leakage model, at time T when the system credentials are reset, the attacker can no longer benefit from the current side-channel attack. In other words, the whole game will be reset every T. The defender seeks to optimize the value of τ_d to minimize chances for information leakage and avoid loading the system with unnecessary migrations. Thus, the defender's goal is to optimize the tradeoff between security and stability.

2) Attacker's action space: In our model, we consider a realistic threat model in which the attacker does not know the placement engine algorithms, hence only tries to increase her co-residency chances via increasing the number of requests submitted to the cloud provider. Let $\lambda_a \in \mathcal{A}_a$ denote the rate of requests (rate of attack) submitted to the cloud, where $\mathcal{A}_a = [\lambda_{\min}, \lambda_{\max}]$ is an interval of non-negative attack rates. The game starts at time $t = 0$, and τ_a denotes the actual time at which the attacker successfully collocates with her targeted victim. Hence, $\tau_a > 0$ is a non-negative random

variable with a probability density function (pdf) $f_a(.;\lambda_a)$ parametrized by λ_a. Since the attacker pays a cost for each submitted job, she needs to optimize over the attack rate λ_a. Hence, the attacker's tradeoff can be summarized as follows. When λ_a is very small, it is less probable for the attacker to successfully co-reside with her victim and in turn leak any information before it is migrated. When λ_a is very large, the attacker increases her chances of successful collocation at the expense of a higher attack cost.

Therefore, the pdf f_a should be such that $f_a(\tau_a;\lambda_{a_1})$ yields a higher probability of early collocation than $f_a(\tau_a;\lambda_{a_2})$, when $\lambda_{a_1} > \lambda_{a_2}$, i.e., $F_a(t;\lambda_{a_1}) \geq F_a(t;\lambda_{a_2})$, where $F_a(t;\lambda_a)$ denotes the Cumulative Distribution Function (CDF) of the collocation time τ_a.

If $\lambda_{\min} = 0$, then the attacker can choose to back off, i.e., not attack. In such case, $F_a(T;0) = 0$ since the probability of collocation is 0. Next, we define the players' reward (payoff) functions. We assume a nonzero-sum two-person game.

3) Attacker's reward: Since the game ends at $\tau_d \leq T$ then repeated, we consider the reward per unit time. Once the attacker is successfully placed on the same node where the victim VM resides, she immediately starts accumulating rewards by leaking information. Let $G(\tau_d,\tau_a)$ denote the reward accumulated by the attacker such that $G(\tau_d,\tau_a) = G(\tau_d-\tau_a)$ for $\tau_a < \tau_d$ and 0 otherwise. Naturally, we assume that G is monotonically non-decreasing in the collocation duration $\tau_d - \tau_a$. The attacker incurs a cost C_a for launching the attack. Hence, the total cost is scaled by the rate of attack λ_a. Therefore, the attacker's payoff is

$$\tilde{u}_a(\tau_d,\tau_a,\lambda_a) = \frac{1}{\tau_d}\left[G(\tau_d,\tau_a)\cdot \mathbf{1}_{\{\tau_a<\tau_d\}} - \lambda_a C_a\right] , \quad (1)$$

where $\mathbf{1}_{\{\}}$ is an indicator function, and the tilde notation signifies the payoff for a given realization of τ_a. Hence, the expected payoff is

$$u_a(\tau_d,\lambda_a) = \int_0^\infty \tilde{u}_a(\tau_d,\tau_a,\lambda_a) f_a(\tau_a;\lambda_a)\, d\tau_a . \quad (2)$$

4) Defender's reward: The defender, on the other hand, incurs a loss due to collocation of a victim VM with the attacker equal in magnitude to the gain of the attacker. In addition, the defender pays a cost per migration, which increases the system overhead and overloads the placement engine. The cost of migration is denoted C_d. Accordingly, the defender's payoff can be written as

$$\tilde{u}_d(\tau_d,\tau_a,\lambda_a) = \frac{1}{\tau_d}\left[-G(\tau_d,\tau_a)\cdot \mathbf{1}_{\{\tau_a<\tau_d\}} - C_d\right] . \quad (3)$$

Averaging over τ_a, the expected payoff for the defender can be calculated as

$$u_d(\tau_d,\lambda_a) = \int_0^\infty \tilde{u}_d(\tau_d,\tau_a,\lambda_a) f_a(\tau_a;\lambda_a) d\tau_a . \quad (4)$$

III. Theoretical Analysis

In this section, we provide sufficient conditions for the existence of NE for the formulated game. Existence of a NE depends on the properties of the payoff functions. First, we derive existence conditions for a general accumulated reward function $G(\tau_d,\tau_a)$ and pdf of the collocation time $f_a(\tau_a;\lambda_a)$, then analyze conditions for a special instance of such functions. We also characterize the best response curves for both players and derive conditions for NE strategies if they exist. Next, we state the main theoretical results and discuss their significance in Section III-B. For proofs of the stated results, we refer the reader to an extended version of this work [28].

A. General reward functions

The following theorem establishes sufficient conditions for the formulated game to admit a pure strategy NE for the general payoff formulation described in (2) and (4).

Theorem 1. *The 2-person nonzero-sum game defined in Section II-B with the payoff functions in (2) and (4) admits a NE in pure strategy if $f_a(\tau_a;\lambda_a)$ is continuous and strictly concave in $\lambda_a \in \mathcal{A}_a$, $\frac{G(\tau_d-\tau_a)}{\tau_d}$ is convex in $\tau_d \in \mathcal{A}_d$, and G is continuous in $\tau_d \in \mathcal{A}_d$.*

The proof of Theorem 1 rests upon establishing sufficient conditions for strict concavity of the payoff functions, specifically ensuring that u_d is strictly concave in τ_d for every $\lambda_a \in \mathcal{A}_a$ and that u_a is strictly concave in λ_a for every $\tau_d \in \mathcal{A}_d$. The following proposition explicitly identifies an important NE in which the attacker backs off and the defender stops migration.

Proposition 2. *For the game defined in Section II-B with $\lambda_{\min} = 0$, there exists an equilibrium in which the attacker backs off (i.e., does not attack) and the defender does not migrate if the reward function G satisfies*

$$\mathbb{E}_{\lambda_a}\left[G(T-\tau_a)\right] \leq \lambda_a C_a, \quad (5)$$

for every $\lambda_a \in \mathcal{A}_a$, where $\mathbb{E}_{\lambda_a}[.]$ denotes the expectation w.r.t. the measure induced by $f(.;\lambda_a)$.

The following theorem characterizes the best response for both players.

Theorem 3. *For the 2-person nonzero-sum game defined in Section II-B, if the attacker's payoff function in (2) is strictly concave in λ_a, then the attacker's best response λ_a^* to any defense strategy can be described as*

- $\lambda_a^* = \lambda_{\max}$, *if* $\frac{\partial u_a}{\partial \lambda_a} > 0,\ \forall\ \lambda_a \in \mathcal{A}_a$
- $\lambda_a^* = \lambda_{\min}$, *if* $\frac{\partial u_a}{\partial \lambda_a} < 0,\ \forall\ \lambda_a \in \mathcal{A}_a$
- $\lambda_a^* \in \left\{\lambda_a \mid \int_0^{\tau_d} G(\tau_d,\tau_a)\frac{\partial f_a(\tau_a;\lambda_a)}{\partial \lambda_a} d\tau_a = C_a\right\}$, *if* $\frac{\partial u_a}{\partial \lambda_a} = 0$, *for any* $\lambda_a \in \mathcal{A}_a$.

Also, if the defender's payoff function in (4) is strictly concave in τ_d, then the best response τ_d^* can be described as

- $\tau_d^* = T$, *if* $\frac{\partial u_d}{\partial \tau_d} > 0,\ \forall\ \tau_d \in \mathcal{A}_d$
- $\tau_d^* = \tau_{\min}$, *if* $\frac{\partial u_d}{\partial \tau_d} < 0,\ \forall\ \tau_d \in \mathcal{A}_d$
- $\tau_d^* \in \left\{ \tau_d \mid \int_0^{\tau_d} \left(\tau_d \frac{\partial G}{\partial \tau_d} - G \right) f_a(\tau_a; \lambda_a) d\tau_a = C_d \right\}$, *if* $\frac{\partial u_d}{\partial \tau_d} = 0$, *for any* $\tau_d \in \mathcal{A}_d$.

B. Special reward functions

In Section III-A, we provided conditions for the existence of an equilibrium for generic reward functions. The conditions imposed were the strict concavity of f_a in addition to the non-negativity, monotonicity and stationarity of G (stationarity in that the accumulated reward depends on the collocation and migration times only through their difference, i.e., the duration of collocation). In this section, we study existence conditions for equilibrium and characterize the best response sets of both players for specific choices of the reward function G and the collocation pdf $f_a(\tau_a; \lambda_a)$ as special cases of interest. In particular, we provide an analysis of the formulated timing game for the case where $G(t)$ increases linearly in the collocation duration t, i.e.,

$$G(\tau_d, \tau_a) = \begin{cases} \alpha\left(\tau_d - \tau_a\right), & \tau_a \leq \tau_d \leq T \\ 0, & \text{otherwise} \end{cases} \tag{6}$$

for some constant $\alpha > 0$. Without loss of generality, we always consider $\alpha = 1$. The case $\alpha \neq 1$ is equivalent to $\alpha = 1$ with the migration cost C_d replaced by $\frac{C_d}{\alpha}$. In Section IV-C, we provide numerical results on the best response for other (non-linear) functions including, for example, when $G(t)$ is quadratic in t.

Motivated by the interpretation of $\lambda_a \in \mathcal{A}_a$ as the rate of attack launched by the adversary, in our numerical evaluation we consider an exponential pdf for the collocation time, i.e.,

$$f_a(\tau_a; \lambda_a) = \lambda_a e^{-\lambda_a \tau_a}\ , \quad \tau_a \geq 0\,. \tag{7}$$

Next, we derive sufficient conditions for the existence of a NE for the choice of functions in (6) and (7).

Theorem 4. *Consider the 2-person nonzero-sum game defined in Section II-B with $G(t)$ and $f_a(\tau_a; \lambda_a)$ defined in (6) and (7). If*

$$1 - \lambda_a C_d < \left(1 + \lambda_a \tau_d + \frac{\lambda_a^2 \tau_d^2}{2}\right) e^{-\lambda_a \tau_d}, \forall (\tau_d, \lambda_a) \in \mathcal{A}\,,$$

then the game admits a pure strategy NE.

The best response for both players for the choice of functions in (6) and (7) can be derived following the same proof of the general characterization in Theorem 3 as in [28] and is omitted here for brevity. The following theorem characterizes bounds on both the attack cost C_a and the migration cost C_d beyond which the players' best response strategies are on the boundaries of their action intervals.

Theorem 5. *For the two person nonzero-sum game defined in Section II-B with the reward function in (6) and the exponentially distributed collocation time τ_a in (7), if*

$$C_a > \frac{1 - (1 + \lambda_{\max}\tau_d) e^{-\lambda_{\max}\tau_d}}{\lambda_{\min}^2}\,,$$

then the attacker's best response to the action τ_d of the defender is $\lambda_a^*(\tau_d) = \lambda_{\min}$. Also, if

$$C_d > \frac{1 - (1 + \lambda_a T) e^{-\lambda_a T}}{\lambda_a}\,,$$

then the defender's best response to the action λ_a of the attacker is to stop migrations, i.e, $\tau_d^*(\lambda_a) = T$.

IV. Numerical Analysis

In this section, we provide numerical analysis of the studied game model. We consider the payoff functions for both players using $G(t)$ and $f_a(\tau_a; \lambda_a)$ defined in (6) and (7). We study the behavior of the payoff functions for both players, then we investigate the effect of the migration cost C_d and the attack cost C_a on the reward functions, the players' best responses, and the existence of a NE. We also investigate different scaling regimes for the reward function.

Figure 2a shows the NE existence region for $C_d = 0.3$. Per Theorem 4, at $T = 1.5$ and $\lambda_{\max} = 5$ as marked with the highlighted rectangle, the game admits a Nash equilibrium in pure strategies. The figure also illustrates the best response curves along with the game action space when $C_d = 0.3$ and $C_a = 0.5$. In the figure, the horizontally dashed region is the region of concavity of u_a in λ_a for all $\tau_d \in \mathcal{A}_d$. Similarly, the vertically dashed region designates the region in which u_d is concave in τ_d for every $\lambda_a \in \mathcal{A}_a$. Any game defined in the region of intersection characterized in Theorem 4 admits a NE in pure strategies. In this setting, the NE is unique – shown as the unique intersection point of the best response curves for both players at $\tau_d^* = 1.27$ and $\lambda_a^* = 0.61$.

A. Payoff functions

Figure 2b shows the payoff function of the defender u_d versus the migration time τ_d for $C_d = 0.3$ and $T = 4$. The figure highlights the tradeoff faced by the defender as he seeks to optimize τ_d to secure the system through VM migration while avoiding a large migration overhead. Evidently, the optimal migration time τ_d^* depends on the attacker's strategy λ_a. The tradeoff agrees with our intuition based on the studied game model. Specifically, a very small τ_d – signifying a high VM migration rate – is associated with a high migration cost that dominates the payoff function u_d. On the other hand, a larger τ_d implies that the VMs dwell for a longer period of time on the same physical node giving the attacker more room to collocate and leak data from her targeted VM. In Figure 2b, we compare the defender's reward at different attack rates λ_a. When the attack is less aggressive, the defender is able to maximize his

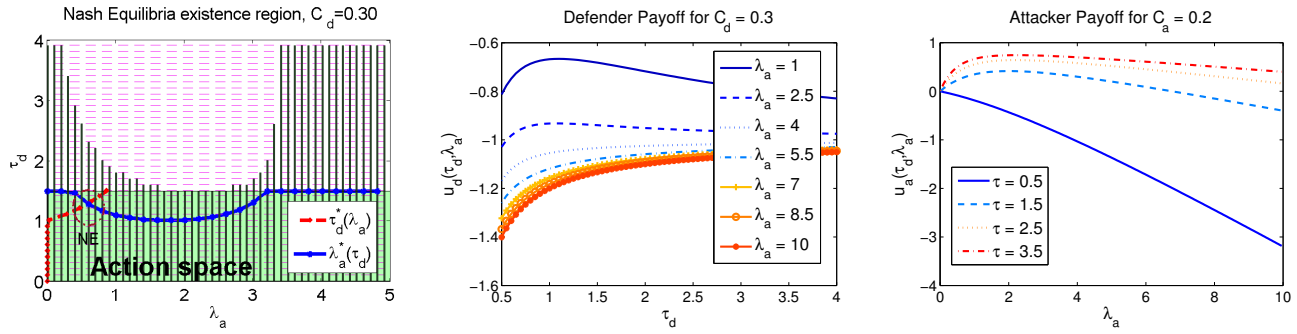

(a) (b) (c)

Figure 2. (a) Verifying NE existence in the region characterized in Theorem 4; (b) Defender's reward versus migration time τ_d; (c) Attacker's reward versus attack rate λ_a.

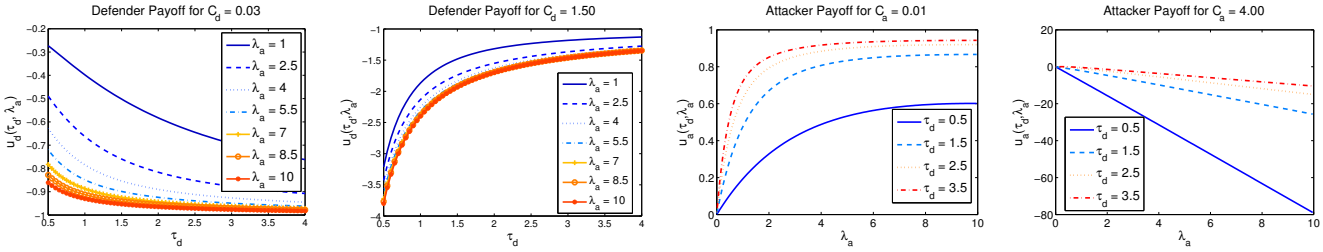

(a) (b) (c) (d)

Figure 3. Defender's reward versus migration time τ_d for (a) $C_d = 0.03$ and (b) $C_d = 1.5$, and attacker's reward versus attack rate λ_a for (c) $C_a = 0.01$ and (d) $C_a = 4$.

payoff by reducing the migration time τ_d at the expense of higher migration cost. Therefore, when λ_a increases from 1 to 2.5, the optimal τ_d reduces from 1.25 to 0.85 resulting in a higher migration frequency. However, when the attacker is very aggressive, the defender is better off avoiding the migration cost by increasing τ_d to T.

In Figure 2c, we plot the attacker's expected payoff u_a versus the attack rate λ_a for different defense actions τ_d for an attack cost $C_a = 0.2$. As shown, the optimal attack rate depends on the defender's action. As the attack rate increases, the cost of attack increases and becomes the dominating term in the payoff function. Moreover, as the defender reduces his time to migrate τ_d, the attacker's reward decreases. This is due to the fact that when the migration rate is higher, there is a shorter time window for the attacker to successfully collocate with her victim. Contrariwise, when the migration rate is not too high (i.e., τ_d is fairly large), the attacker can maximize her reward by increasing the attack rate λ_a. However, if the defender is migrating the VMs at a very high rate, i.e, τ_d is very small, the attacker's best response is to attack at the minimum possible rate or completely back-off since the attack is useless.

B. Cost effect and monotonicity

To show the effect of the migration and attack costs C_d and C_a, we plot the players' payoff functions for different values of the cost. In Figure 3a, we plot the defender's payoff versus τ_d for different attack strategies for a fairly small migration cost $C_d = 0.03$. At this small migration cost, the defender's best response is to always migrate at the highest permissible rate, i.e, $\tau_d^* = \tau_{\min}$ regardless of the attack rate λ_a. Hence, the leakage loss term dominates the defender's payoff function u_d. On the other hand, when the migration cost is high as shown in Figure 3b where $C_d = 1.5$, the defender's best response is $\tau_d^* = T$ to reduce the associated migration cost, a fact which was established analytically in Theorem 5. Similarly, the effect of the attack cost C_a is exhibited in Figure 3c and Figure 3d. At a very small attack cost, $C_a = 0.01$, the attacker's best strategy is to attack aggressively with $\lambda_{\max}$ to maximize the chances of successful collocation regardless of the defender's action as shown in Figure 3c. In Figure 3d, the attack cost is high as $C_a = 4$, so the cost of the attack term dominates the payoff function. Therefore, the best action for the attacker is $\lambda_{\min}$ regardless of the action of the defender, which follows our result in Theorem 5.

C. Best response curves

In this section, we study the best response curves to provide more insight into the optimal action of a player as function of the action of the opponent. Figure 4a shows the defender's best response curve τ_d^* as function of λ_a

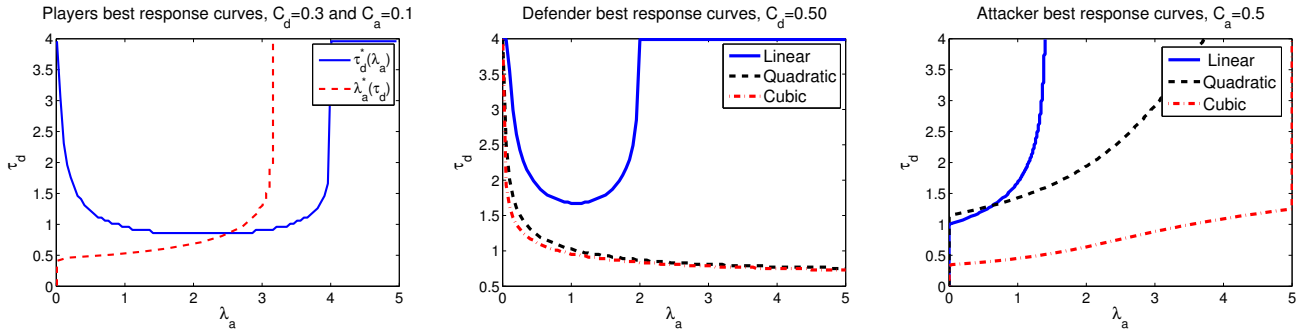

(a) (b) (c)

Figure 4. Players best response curves.

and the attacker's best response curve λ_a^* as function of the defender's action τ_d. In this scenario, we set $T = 4$, $\lambda_{\max} = 5$, $C_d = 0.3$, and $C_a = 0.1$. In Figure 4a, the intersection point of the two response curves is the unique Nash equilibrium. The point(s) of equilibria depend on the values of C_a and C_d. The best response curves also underscore the tradeoff for each player. For example, at equilibrium the defender migrates with $\tau_d = 0.86$ while the attacker uses rate $\lambda_a = 2.51$ for the attack. Clearly, at low attack rate, VM migration at a very small migration rate, i.e, larger τ_d, is more favorable. As the attack rate increases, the defender is urged to migrate the VMs at faster rate, wherefore τ_d^* decreases as λ_a increases, but only until a certain point where faster migration becomes futile. Indeed, when the attack rate is overwhelming, it is more rewarding for the defender to use a large τ_d to alleviate high migration costs. On the attacker's side, a similar tradeoff is observed. The attacker attacks the system at the minimum rate $\lambda_{\min}$ as long as the VM stays on the same physical node for a duration $\tau_d < 0.4$ since it is very hard to collocate when migration is taking place at such high rates. If the defender increases the time before migrating, i.e. $\tau_d > 0.4$, the attacker is enticed to attack the system at higher rates to increase information leakage. However, the maximum attack rate the attacker will select is $\lambda_a = 3.16$, which is strictly smaller than $\lambda_{\max}$, since the resultant attack cost yields a smaller overall payoff.

D. Higher order reward regimes

In order to shed light on the importance of the data leakage model, we study other scaling regimes for the reward function. In particular, we consider the scenario where the reward function $G(\tau_d, \tau_a)$ scales quadratically or cubically with the collocation duration. In Figure 4b, we plot the defender's best response curves for linear, quadratic, and cubic reward functions. Intuitively, higher order reward functions are more disposed to dominate the payoff functions than for the linear scaling. It is obvious that in the linear regime the defender is facing the tradeoff discussed earlier in Section IV-C. However, for higher order reward regimes, the reward term dominates the payoff over the entire range of attack rates in this setting. Therefore, the defender is consistently urged to increase the migration rate as the attacker increases her attack rates.

In Figure 4c, the attacker's best response curves are plotted for the different reward functions. In the linear regime, the attacker's best response rate is non-vanishing and increasing in τ_d for $\tau_d > 1$, but saturates at $\lambda_a = 1.4$ as soon as the cost of the attack starts to dominate the attacker's payoff. For both the quadratic and cubic regimes, the higher reward from data leakage entices the attacker to attack at higher rates as τ_d increases.

V. Conclusion

We developed a moving target defense framework for the virtual machines migration timing problem. Live migration of virtual machines between different physical nodes is studied in a game-theoretic framework to defend multi-tenant clouds against side channel attacks launched by malicious users who are co-residing on the same physical node. We established sufficient conditions for existence of Nash equilibrium for the proposed game model, as well as best response strategies. We also verified our theoretical results numerically for different settings of the game. The theoretical and numerical analyses provided characterize the performance of the migration defense approach against collocation attacks.

Acknowledgment

This work was supported in part by NSF awards CCF-1552497 and CNS-1149397.

References

[1] K. Ren, C. Wang, and Q. Wang, "Security challenges for the public cloud," *IEEE Internet Computing*, vol. 16, no. 1, pp. 69–73, 2012.

[2] Moving target defense. [Online]. Available: https://www.dhs.gov/science-and-technology/csd-mtd

[3] A. H. Anwar, G. Atia, and M. Guirguis, "Game theoretic defense approach to wireless networks against stealthy decoy attacks," in *54th Annual Allerton Conference on Communication, Control, and Computing*, 2016, pp. 816–821.

[4] ——, "Dynamic game-theoretic defense approach against stealthy jamming attacks in wireless networks," in *Communication, Control, and Computing (Allerton), 2017 55th Annual Allerton Conference on*. IEEE, 2017, pp. 252–258.

[5] Y. Yarom and K. Falkner, "FLUSH+RELOAD: A high resolution, low noise, L3 cache side-channel attack." in *USENIX Security Symposium*, 2014, pp. 719–732.

[6] W. Voorsluys, J. Broberg, S. Venugopal, and R. Buyya, "Cost of virtual machine live migration in clouds: A performance evaluation." *CloudCom*, vol. 9, pp. 254–265, 2009.

[7] T. Ristenpart, E. Tromer, H. Shacham, and S. Savage, "Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds," in *Proceedings of the 16th ACM conference on Computer and communications security*, 2009, pp. 199–212.

[8] K. Suzaki, K. Iijima, T. Yagi, and C. Artho, "Memory deduplication as a threat to the guest OS," in *Proceedings of the Fourth European Workshop on System Security*. ACM, 2011.

[9] R. Owens and W. Wang, "Non-interactive OS fingerprinting through memory de-duplication technique in virtual machines," in *IEEE 30th International Performance Computing and Communications Conference (IPCCC)*, 2011, pp. 1–8.

[10] Y. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart, "Cross-VM side channels and their use to extract private keys," in *Proceedings of the ACM conference on Computer and communications security*, 2012, pp. 305–316.

[11] F. Liu, Y. Yarom, Q. Ge, G. Heiser, and R. B. Lee, "Last-level cache side-channel attacks are practical," in *IEEE Symposium on Security and Privacy (SP)*, 2015, pp. 605–622.

[12] G. Irazoqui, M. S. Inci, T. Eisenbarth, and B. Sunar, "Wait a minute! A fast, Cross-VM attack on AES," in *International Workshop on Recent Advances in Intrusion Detection*. Springer, 2014, pp. 299–319.

[13] W. Qi, J. Wang, H. Hovhannisyan, K. Lu, J. Wang, and J. Zhu, "A generic mitigation framework against Cross-VM covert channels," in *IEEE 25th International Conference on Computer Communication and Networks (ICCCN)*, 2016, pp. 1–10.

[14] Y. Han, T. Alpcan, J. Chan, and C. Leckie, "Security games for virtual machine allocation in cloud computing," in *International Conference on Decision and Game Theory for Security*. Springer, 2013, pp. 99–118.

[15] C. A. Kamhoua, L. Kwiat, K. A. Kwiat, J. S. Park, M. Zhao, and M. Rodriguez, "Game theoretic modeling of security and interdependency in a public cloud," in *IEEE 7th International Conference on Cloud Computing (CLOUD)*, 2014, pp. 514–521.

[16] L. Kwiat, C. A. Kamhoua, K. A. Kwiat, J. Tang, and A. Martin, "Security-aware virtual machine allocation in the cloud: A game theoretic approach," in *IEEE 8th International Conference on Cloud Computing (CLOUD)*, 2015, pp. 556–563.

[17] C. Kamhoua, A. Martin, D. K. Tosh, K. A. Kwiat, C. Heitzenrater, and S. Sengupta, "Cyber-threats information sharing in cloud computing: A game theoretic approach," in *IEEE 2nd International Conference on Cyber Security and Cloud Computing (CSCloud)*, 2015, pp. 382–389.

[18] B. Di Martino and A. Esposito, "Automatic dynamic data structures recognition to support the migration of applications to the cloud," *International Journal of Grid and High Performance Computing (IJGHPC)*, vol. 7, no. 3, pp. 1–22, 2015.

[19] C.-H. Hsu, S.-J. Peng, T.-Y. Chan, K. Slagter, and Y.-C. Chung, "An adaptive pre-copy strategy for virtual machine live migration," in *International Conference on Internet of Vehicles*. Springer, 2014, pp. 396–406.

[20] C. Dhule and U. Shrawankar, "Performance analysis for pareto-optimal green consolidation based on virtual machines live migration," *International Journal of Grid and High Performance Computing (IJGHPC)*, vol. 9, no. 4, pp. 36–56, 2017.

[21] D. Blackwell, "The noisy duel, one bullet each," arbitrary accuracy. Technical report, The RAND Corporation, D-442, Tech. Rep., 1949.

[22] K. D. Bowers, M. Van Dijk, R. Griffin, A. Juels, A. Oprea, R. L. Rivest, and N. Triandopoulos, "Defending against the Unknown Enemy: Applying FlipIt to System Security," in *GameSec*. Springer, 2012, pp. 248–263.

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

[24] P. Hu, H. Li, H. Fu, D. Cansever, and P. Mohapatra, "Dynamic defense strategy against advanced persistent threat with insiders," in *IEEE Conference on Computer Communications (INFOCOM)*, 2015, pp. 747–755.

[25] M. Zhang, Z. Zheng, and N. B. Shroff, "Stealthy attacks and observable defenses: A game theoretic model under strict resource constraints," in *IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, 2014, pp. 813–817.

[26] J. Pawlick, S. Farhang, and Q. Zhu, "Flip the cloud: cyber-physical signaling games in the presence of advanced persistent threats," in *International Conference on Decision and Game Theory for Security*. Springer, 2015, pp. 289–308.

[27] S. Farhang and J. Grossklags, "FlipLeakage: a game-theoretic approach to protect against stealthy attackers in the presence of information leakage," in *International Conference on Decision and Game Theory for Security*. Springer, 2016, pp. 195–214.

[28] A. H. Anwar, G. Atia, and M. Guirguis, "It's time to migrate! a game-theoretic framework for protecting a multi-tenant cloud against collocation attacks," 2018. [Online]. Available: https://arxiv.org/