A Game-Theoretic Framework for the Virtual Machines Migration Timing Problem

Ahmed H. Anwar[1], George Atia[1], and Mina Guirguis[2]

[1]Department of Electrical and Computer Engineering, University of Central Florida, Orlando, FL 32816

[2]Department of Computer Science, Texas State University, San Marcos, TX 78666
a.h.anwar@knights.ucf.edu, george.atia@ucf.edu, msg@txstate.edu

Abstract—In a multi-tenant cloud, a number of Virtual Machines (VMs) are collocated on the same physical machine to optimize performance, power consumption and maximize profit. This, however, increases the risk of a malicious VM performing side-channel attacks and leaking sensitive information from neighboring VMs. As such, this paper develops and analyzes a game-theoretic framework for the VM migration 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 with malicious VMs. It also captures costs incurred by the adversary in launching new VMs and by the defender in migrating VMs. 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. Furthermore, we extend our model to characterize its impact on the attacker's payoff when the cloud utilizes intrusion detection systems that detect side-channel attacks. Our theoretical findings are corroborated with extensive numerical results in various settings as well as a proof-of-concept implementation in a realistic cloud setting.

Index Terms—Cloud security, game theory, VM migration.

1 Introduction

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 effective attacks [2]. When developing an MTD strategy, two main questions generally arise: *which* targets should be moved? and *when* should they be moved? The answers to these questions largely depend on the context of the problem and the nature of the attack. For example, if an attacker contemplates to infer the underlying topology of the cloud, then the target is the machine connectivity that should then be adapted over time. However, if the attacker seeks to crack system credentials that protect the users' databases, then the target are the keys that should be constantly reconfigured (i.e., moved). In this paper, we consider collocation attacks whereby an attacker can access 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 systems, VMs should be periodically migrated, i.e., moved to different physical machines. While much work focused on the scheduling and placement aspect of VM migration, the timing problem is largely understudied. This motivates the work in this paper, which 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 in a game-theoretic framework seeking an understanding of the interplay of the strategies of the cloud provider (i.e., the defender) and the adversary. In our formulation, the adversary seeks to prolong the collocation time with the victim VMs to maximize the amount of information she can access. Since the adversary has no guarantees for being successfully collocated on the same node with the victim (different cloud providers implement different placement algorithms according to different criteria that the attacker has no control over), her best-effort is to increase the number of VMs to launch (which is a cost metric we capture). After the adversary is placed on given physical machines, she can check whether she had a successful collocation or not [3]. The cloud provider, on the other hand, migrates VMs between physical machines to minimize the collocation times between VMs. VM live migration, while efficient at not significantly disrupting the tasks running on a VM in the event of migration, is not free [4]. In practice, the number of cache pages read by an adversary from shared memory pages is proportional to the duration of a side-channel attack. It also depends on

the technique used to access the last level cache (LLC) (e.g, PRIME+PROBE and FLUSH+RELOAD attacks) as shown in [5]. To read the cache, an attacker would need to adjust the time of the PROBE phase, which in turn affects the error rate of the attack covert channel. Thus, the question as to when to migrate is crucial. At the same time, the defender controls the migration time in order to mitigate the collocation attack threats while not burdening the system with significant overhead, e.g., due to VM downtime and undue memory usage.

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 [6], [7], which is largely unexplored in cloud computing settings. We leverage the results of the leakage model in the FlipIt game considered previously in [8], [9], [10], [11], [12], [13] 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 timing games.

- We provide a new game-theoretic formulation for the VM collocation timing problem.
- Unlike [14], [15], [16], 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 migrate the VMs at the right time(s) to defend against malicious collocating users.
- We analytically characterize the Nash equilibrium (NE) for the studied game model and derive sufficient existence conditions.
- We study the behavior of the adversary when the defender adopts an intrusion detection system (IDS). In this case, the adversary not only takes attack actions, but also decides when to stop her attack to reduce the risk of being detected.
- We provide extensive numerical experiments to support our theoretical findings. In our numerical evaluation, we consider several reward functions to reflect the degree of the attack and the severity of the data breach. As a proof of concept, we also implement the migration defense approach on a realistic cloud setup using the Xen hypervisor.

This paper is organized as follows. In Section 3, we present the system model and the game formulation. In Section 4, we provide a theoretical analysis and establish existence conditions of NE for the formulated game. Section 5 studies an extension of the game model in the presence of an IDS. Our numerical results are presented in Section 6 and we conclude the paper in Section 7.

2 Related Work

This work is at the intersection of two areas focused on securing cloud computing: Cross-VM side-channel attacks and mitigation, and game-theoretic modeling and techniques in cloud security. In this section, we put our work in context within these two areas.

2.1 Cross-VM side channel attacks and mitigation strategies

Cloud security has received considerable attention recently [1], [17]. Various studies have investigated the impact of cross-VM side-channel attacks [18], [19], [20], [3], [21], [5], [22]. Users' cryptographic keys have been shown to be vulnerable to exfiltration attacks when adversaries perform Prime+Probe attacks on the square-and-multiply implementation of GnuPG [21]. The authors in [3], [22], [5] have shown that some side-channel attacks can extract cryptographic keys by exploiting the last-level shared caches of the memory. Other attacks have identified pages that a VM shares with its collocated neighboring VMs revealing information about the victim's applications [19] and OS [20].

To combat cross-VM side-channel attacks, various approaches have been proposed at the hypervisor [23], [24], [21], [25], [26]), the guest OS [27], the hardware level [28], [29], and the application layer [30]. These techniques, however, suffer from two fundamental limitations. First, they cannot be generalized to different types of side-channel attacks [31]. Second, they require major changes to the hypervisor, OS, hardware, and applications [32]. VM live migration, on the other hand, has been proposed as an effective mechanism to combat side-channel attacks [4], [33]. The authors in [34] provided a detection mechanism known as *CloudRadar* that works as a real-time side-channel attack detector based on monitoring hardware performance counters. The authors in [35] proposed another detection system that can differentiate between friendly and other malicious activities of neighboring tenants. The authors in [36] 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 large. Motivated by the MTD concept, the authors in [37] presented a migration engine in which VMs are migrated to balance the load between different nodes in the cloud. Although MTD is a well-known defense methodology, the authors in [38] demonstrated that in certain scenarios the migrated VMs can be tracked by adversaries. Hence, they proposed a stealthy approach to migrate VMs that can hide them on the network. In [39], the authors study an MTD migration strategy against an attacker that seeks to collocate with VMs of high rewards by solving a multi-armed bandit problem.

2.2 Cloud security using game-theoretic techniques

The use of game theory has largely focused on the VM allocation problem in the presence of adversaries [14], [15], [40], [16], [41]. A common assumption in such formulations is that the adversary is known, which may not hold in practice. Additionally, existing formulations do not consider the timing question for the VM migration problem, which is a critical one for the cloud provider wishing to migrate VMs for security. A more practical leakage model was considered

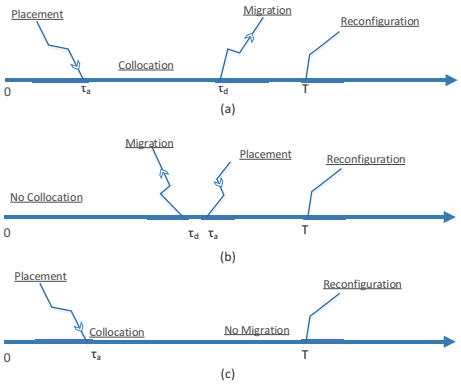

Fig. 1: System model illustration for different placement events.

in [42], [43], based on the FlipIt game model. FlipIt is a two-player game in which a defender and an attacker compete over the control of a given resource, which can only be held by one player at a time. A flip is an action performed by a player to gain control of the resource. The goal is to hold the resource for the longest duration possible with the least number of flips (i.e., flips are costly). Over time, the resource generates rewards for the player holding the resource. The state of the resource is obscured from each player until they "flip". Several variants of the FlipIt game model were considered to study different security situations [8], [9], [44], [45], [10], [11], [12], [13]. In [9], the authors studied different strategies for each player and calculated dominant strategies and Nash equilibria. In [44], the game model was studied under the assumption that the players know the state of the resource before taking actions. In [45], [10] the game was extended to the case of a system where insiders can work in favor of external adversaries. The authors in [11] considered the game with both players having limited budgets. Pawlick *et al*. investigated the game model with characteristics of signaling games [12]. In [13], Farhang *et al*. studied a variant of the FlipIt game with an associated data leakage model in which the defender can partially eliminate the foothold of the attacker. The attacker exploits the system vulnerabilities that appear based on a periodic process. The authors assume that the attacker's strategy is fixed since she always starts to attack right after the defender takes his action. This, however, requires the attacker to fully observe the defender's strategy which we do not assume here.

In this work, we consider a significantly different and a realistic threat model that captures data leakage due to cross-VM side-channel attacks and develop defense strategies for identifying the best time(s) to migrate VMs. We do this through a game-theoretic framework in which the attacker only controls the attack rate and does not fully observe the defender's strategy. In addition, we assume that the attacker controls the probability of a successful attack by choosing the attack rate as opposed to the time to launch the attack.

3 System Model

3.1 The cloud

We model the cloud as a set of physical machines and 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 machines. 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 chances for prolonged collocation with her victims. Fig. 1 illustrates three possible placement scenarios for the game. In plot (a), the attacker's VM is successfully collocated at time τ_a with her target VM on the same hypervisor before the target VM is migrated to another node at time τ_d. This scenario represents a successful collocation event, which results in information leakage. In plot (b), the target VM is migrated before the malicious VM is placed on the hypervisor, hence the collocation event does not occur. Finally, the plot in (c) illustrates a no-migration policy, where the collocation duration is maximized. We define the game next.

3.2 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.
- $\mathcal{U} = \{u_d, u_a\}$ is the reward function, $\mathcal{U} : \mathcal{A} \to \mathbb{R}^2$.

3.2.1 *Defender's action space*

Since we are investigating the timing factor, the cloud provider (referred to as the system defender) is assumed to control the re-allocation period. Let $\tau_d \in \mathcal{A}_d$ denote the time instant 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 $\tau_{\min}$ is 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 side-channel attack. Therefore, 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. In particular, a smaller τ_d ensures the system is more secure since the co-residency times between any two VMs will be small. However, the system's overhead increases due to frequent migration of the VMs between the physical nodes. The overhead of VM live migration has been investigated in [4], [46], and in general depends on the VM workload. The work in [46] has shown that the main

factors affecting the VM migration overhead are the VM memory size and the network speed. On the other hand, a larger τ_d leads to a more stable system. However, the co-residency times between VMs on the same node will be large making the system more susceptible to a data breach through collocation attacks.

3.2.2 Attacker's action space

Here, we assume that the attacker does not know the system placement 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 is assumed to start at time $t = 0$, and let τ_a denote 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 steal any information before VMs are 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 $\lambda_{a_1} > \lambda_{a_2}$. Mathematically, this requirement is expressed through the following assumption.

Assumption 1. *$F_a(t; \lambda_{a_1}) \geq F_a(t; \lambda_{a_2})$ for $\lambda_{a_1} \geq \lambda_{a_2}$, where $F_a(t; \lambda_a) := \Pr(\tau_a \leq t)$ denotes the cumulative distribution function (CDF) of the collocation time.*

If $\lambda_{\min} = 0$, then the attacker can choose to back off (i.e., not attack). In such case, $f_a(\tau_a; 0)$ is a degenerate deterministic distribution such that $F_a(T; 0) = 0$ since the probability of collocation is 0.

We focus only on the timing factor of the problem, and the mapping of VMs to physical nodes is carried out through the placement engine. The separation of the placement and timing strategies allows for layered functionality highly desirable in practice. In particular, the developed timing policies can be implemented on any existing platform without modifying the existing placement engine. This is especially true since allocation decisions are typically developed around widely differing load balancing and power reduction objectives, and other operational constraints. Next, we define the players' reward (payoff) functions in a nonzero-sum two-person game.

3.2.3 Attacker's reward

Once the attacker's VM is successfully placed on the same node where the victim VM resides, she immediately starts accumulating rewards by reading out data from the target VM. The amount of information leakage depends in practice on the duration of collocation as shown in [5]. Let $G(\tau_d, \tau_a)$ denote the reward accumulated by the attacker capturing the relation between the collocation duration and the amount of the data leaked.

Assumption 2. *$G(\tau_d, \tau_a)$ is a stationary function and monotonically non-decreasing in the collocation duration $t = \tau_d - \tau_a$. Therefore, $G(\tau_d, \tau_a) = G(\tau_d - \tau_a, 0) = G(t)$, where $G(t)$ is an abbreviated notation indexed by one variable.*

Stationarity signifies that the attacker's accumulated reward depends on the collocation and migration times only through their difference, i.e., the duration of collocation. The accumulated reward is assumed to be zero if $\tau_a \geq \tau_d$. The attacker incurs a cost C_a for launching an attack. Hence, the total cost is scaled by the rate of attack λ_a. Therefore, the attacker's expected payoff is given by

$$u_a(\tau_d, \lambda_a) = \int_0^{\tau_d} G(\tau_d, \tau_a) f_a(\tau_a, \lambda_a)\, d\tau_a - C_a \lambda_a. \quad (1)$$

3.2.4 Defender's reward

The defender, on the other hand, incurs a loss due to the 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 denoted by C_d. This cost captures the migration overhead, which stems from the VM downtime, performance degradation of the running applications (e.g., due to successive iterations of memory pre-copying [47]), and the amount of memory and cache usage. Accordingly, the defender's expected payoff can be written as

$$u_d(\tau_d, \lambda_a) = -\int_0^{\tau_d} G(\tau_d, \tau_a) f_a(\tau_a, \lambda_a)\, d\tau_a - \frac{C_d}{\tau_d}. \quad (2)$$

4 Theoretical Analysis

A NE characterizes a solution for non-cooperative games in which no player can gain by deviating from his own equilibrium strategy while the other players' strategies are fixed [48], [49]. In this section, we establish sufficient conditions for the existence of a NE for the formulated game model in Theorems 3 and 8. We characterize the players' best responses (c.f. Definition 1) in Theorem 5 and Lemma 9. By definition, a NE secures a minimum reward for the cloud admin since the defender's reward cannot decrease if only the adversary deviates from her best response. Since the rewards for both players depend on the cost parameters C_a and C_d (c.f. Section 3), the role of these parameters is also analyzed in Lemmas 6, 7 and Theorems 10 and 11.

Existence of 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 $f_a(\tau_a; \lambda_a)$ of collocation time, then we provide analysis for a special instantiation of the payoff functions. We also characterize the best response curves for both players and derive existence conditions for the corresponding NE strategies. First, we restate a general theorem from [49] that provides sufficient conditions for N-person nonzero-sum games to admit a pure strategy NE.

Theorem 1. *[49] For each player i in the set $\mathcal{N}$ of N players, let the action space U_i of player i be a closed, bounded and convex subset of a finite-dimensional Euclidean space, and the cost functional $J_i : U_1 \times \cdots \times U_N \to \mathbb{R}$ be jointly continuous in all its arguments and strictly convex in $u_i \in U_i$, for every $u_j \in U_j, j \in \mathcal{N}, j \neq i$. Then, the associated N-person nonzero-sum game admits a Nash equilibrium in pure strategy.*

4.1 General reward functions

For the general payoff formulation described in equations (1) and (2), the following lemma proved in Appendix A establishes sufficient conditions for the concavity of the payoff functions.

Lemma 2. *For the 2-person nonzero-sum game defined in Section 3.2 with payoff functions defined in equations (1) and (2) under Assumptions 1 and 2, if* $\mathbb{E}\left[G(\tau_d,\tau_a)\mathbf{1}_{\{\tau_a<\tau_d\}}\right]$ *is strictly concave in* $\lambda_a \in \mathcal{A}_a$ *for any* τ_d*, then* $u_a(\tau_d,\lambda_a)$ *is strictly concave in* λ_a *for any* $\tau_d \in \mathcal{A}_d$*, and if* $\mathbb{E}\left[G(\tau_d,\tau_a)\mathbf{1}_{\{\tau_a<\tau_d\}}\right]$ *is convex in* $\tau_d \in \mathcal{A}_d$*, then* $u_d(\tau_d,\lambda_a)$ *is strictly concave in* τ_d *for any* $\lambda_a \in \mathcal{A}_a$*, where* $\mathbf{1}_{\{.\}}$ *is an indicator function.*

Therefore, we can readily state sufficient conditions for our game to admit a pure strategy NE.

Theorem 3. *The 2-person nonzero-sum game defined in Section 3.2 under Assumptions 1 and 2 with the payoff functions in (1) and (2) admits a NE in pure strategy if* $\mathbb{E}\left[G(\tau_d,\tau_a)\mathbf{1}_{\{\tau_a<\tau_d\}}\right]$ *is continuous and strictly concave in* $\lambda_a \in \mathcal{A}_a$*, and* $\mathbb{E}\left[G(\tau_d,\tau_a)\mathbf{1}_{\{\tau_a<\tau_d\}}\right]$ *is convex and* G *is continuous in* $\tau_d \in \mathcal{A}_d$*.*

The proof of Theorem 3 follows directly from Lemma 2, which establishes strict concavity of the payoff functions under the conditions in the statement of the theorem, and Theorem 1 from [49].

Proposition 4. *For the game defined in Section 3.2 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* $G(t)$ *satisfies*

$$\mathbb{E}_{\lambda_a}\left[G(T-\tau_a)\right] \leq \lambda_a C_a, \tag{3}$$

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

Proof. If the attacker backs off, i.e., chooses $\lambda_a = \lambda_{\min} = 0$, then the defender's payoff in (2) becomes

$$u_d(\tau_d, 0) = \frac{-C_d}{\tau_d},$$

which attains its maximum at $\tau_d = T$ for any $C_d > 0$. Hence, the defender's best response is to not migrate over the game interval. Also, if condition (3) in the statement of Proposition 4 is satisfied, then the attacker's best response to the defender's action $\tau_d = T$ is $\lambda_a = 0$. To see that note that if

$$\mathbb{E}_{\lambda_a}\left[G(T,\tau_a)\right] = \int_0^\infty G(T,\tau_a) f_a(\tau_a;\lambda_a) d\tau_a \leq \lambda_a C_a\ ,$$

then

$$\int_0^{\tau_d} G(\tau_d,\tau_a) f_a(\tau_a;\lambda_a) d\tau_a \leq \lambda_a C_a$$

since $G(t)$ is monotonically non-decreasing in t per Assumption 2. Recalling the attacker's payoff function in (1), the attacker's decision to back off is at least as good as launching an attack at an alternative non-vanishing rate since the cost of the attack upper bounds the leakage reward for any $\lambda_a \neq 0$.

□

Definition 1. *In an N-person nonzero sum game, let* $u_i(a_1,\ldots,a_i,\ldots,a_N)$ *be the reward function of player* i*. For each player* $i \in \{1,\ldots,N\}$*, assume that the maximum reward of* u_i *with respect to* $a_i \in \mathcal{A}_i$ *can be attained for any players' action profile* $a_{-i} \in \mathcal{A}_{-i}$*, where* $a_{-i} := \{a_1,\ldots,a_{i-1},a_{i+1},\ldots,a_N\}$ *and* $\mathcal{A}_{-i} \equiv \mathcal{A}_1 \times \ldots \times \mathcal{A}_{i-1} \times \mathcal{A}_{i+1} \times \ldots \times \mathcal{A}_N$*. Then, the set* $R_i(a_{-i}) \subset \mathcal{A}_i$ *defined by*

$$R_i(a_{-i}) = \{\zeta \in \mathcal{A}_i : u_i(\zeta, a_{-i}) \geq u_i(a_i, a_{-i}),\ \forall a_i \in \mathcal{A}_i\}\,,$$

is called the optimal (or best) response of player i*. If* R_i *is a singleton for every* $a_{-i} \in \mathcal{A}_{-i}$*, then it is called the reaction curve [49].*

Accordingly, it follows from the definition of a NE (in that no player can gain by a unilateral change of strategy if the strategies of the other players remain unchanged) that the intersection points of the best responses are NE. In the following theorem, we characterize the best response for both players.

Theorem 5. *For the 2-person nonzero-sum game defined in Section 3.2, if the attacker's payoff function in (1) 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* $\frac{\partial u_a}{\partial \lambda_a} < 0,\ \forall\ \lambda_a \in \mathcal{A}_a$
- $\lambda_a^* \in \left\{\lambda_a \mid \frac{\partial}{\partial \lambda_a}\mathbb{E}_{\lambda_a}[G(\tau_d,\tau_a)\mathbf{1}_{\{\tau_a<\tau_d\}}] = C_a\right\}$*, 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 (2) 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* $\frac{\partial u_d}{\partial \tau_d} < 0,\ \forall\ \tau_d \in \mathcal{A}_d$
- $\tau_d^* \in \left\{\tau_d \mid \tau_d^2 \frac{\partial}{\partial \tau_d}\mathbb{E}_{\lambda_a}[G(\tau_d,\tau_a)\mathbf{1}_{\{\tau_a<\tau_d\}}] = C_d\right\}$*, if* $\frac{\partial u_d}{\partial \tau_d} = 0$*, for any* $\tau_d \in \mathcal{A}_d$*.*

Proof. Given the concavity of the payoff function u_a in $\lambda_a \in \mathcal{A}_a$, the derivative $\frac{\partial u_a}{\partial \lambda_a}$ is monotone. Hence, there exist three possibilities for the behavior of u_a. If $\frac{\partial u_a}{\partial \lambda_a} > 0$, then u_a is strictly increasing in λ_a for all $\lambda_a \in \mathcal{A}_a$, thus the payoff is maximized by $\lambda_a^* = \lambda_{\max}$. If $\frac{\partial u_a}{\partial \lambda_a} < 0,\ \forall \lambda_a \in \mathcal{A}_a$, then u_a is strictly decreasing in λ_a for all $\lambda_a \in \mathcal{A}_a$, thus the payoff is maximum at $\lambda_a^* = \lambda_{\min}$. Otherwise, u_a attains its maximum when $\frac{\partial u_a}{\partial \lambda_a} = 0$, hence the best response λ_a^* belongs to the set $\Lambda_a = \left\{\lambda_a \mid \int_0^{\tau_d} \frac{\partial f_a}{\partial \lambda_a} G(\tau_d,\tau_a) d\tau_a = C_a\right\}$ at which $\frac{\partial u_a}{\partial \lambda_a} = 0$. The second part of Theorem 5 which characterizes the defender's best response can be proven similarly. □

Next, we study the effect of the attack cost C_a and the moving cost C_d and state bounds on the costs beyond which no player is interested in the game. When the cost C_a exceeds a certain threshold, the cost of the attack dominates the attacker's tradeoff, i.e., the attacker is better off backing off over attempting to access the victim's information. Similarly, if C_d is too high, the defender incurs a cost for migration that exceeds any benefit he would get at any migration rate.

In the following lemma, we derive a lower bound on the attack cost C_a beyond which the attacker is always better

off attacking with the minimum rate $\lambda_{\min}$. If $\lambda_{\min} = 0$, then the attacker will back off.

Lemma 6. *For the two person nonzero-sum game Γ defined in Section 3.2, if $\mathbb{E}\left[G(\tau_d, \tau_a)\mathbf{1}_{\{\tau_a < \tau_d\}}\right]$ is strictly concave in $\lambda_a \in \mathcal{A}_a$, and $C_a > \frac{\partial}{\partial \lambda_a}\mathbb{E}_{\lambda_a}[G(\tau_d, \tau_a)\mathbf{1}_{\{\tau_a < \tau_d\}}] \mid_{\lambda_a = \lambda_{\min}}$, then the attacker's best response to any defense strategy τ_d is to attack at the minimum permissible rate $\lambda_{\min}$.*

Proof. We argue that under the condition stated in the lemma, the attacker's payoff is monotonically decreasing in λ_a. Hence, $\lambda_a^* = \lambda_{\min}$ is the attacker's best response to any τ_d. To show that $\lambda_{\min}$ is the unique best response, assume for contradiction there exists $\lambda^* = \lambda_1 \neq \lambda_{\min}$. If $C_a > \int_0^{\tau_d} G f_a'(\lambda_a)\, d\tau_a \mid_{\lambda_a = \lambda_{\min}}$, where $f_a'(\lambda_a) = \frac{\partial f_a}{\partial \lambda_a}$, then u_a is monotonically decreasing, therefore $u_a(\lambda_{\min}) > u_a(\lambda_1)$ since $\lambda_1 > \lambda_{\min}$. Hence, λ_1 is not in the best response set. □

Similarly, the following lemma establishes a lower bound on the migration cost C_d of the defender, beyond which it is more advantageous not to migrate before the system reconfiguration cycle T.

Lemma 7. *For the two person nonzero-sum game Γ defined in Section 3.2, if $\mathbb{E}_{\lambda_a}\left[G(\tau_d, \tau_a)\mathbf{1}_{\{\tau_a < \tau_d\}}\right]$ is strictly convex in λ_a and G is continuous in $\tau_d \in \mathcal{A}_d$, and $C_d > T^2 \frac{d}{d\tau_d}\mathbb{E}_{\lambda_a}\left[G(\tau_d, \tau_a)\mathbf{1}_{\{\tau_a < \tau_d\}}\right] \mid_{\tau_d = T}$, then the action of not migrating any VM before T is the defender's unique best response regardless of the attacker's strategy λ_a, where $\mathbb{E}_{\lambda_a}[.]$ is the expectation with respect to $f_a(\tau_a; \lambda_a)$.*

Proof. By an argument similar to the proof of Lemma 6, under the condition in the statement of the lemma, the defender's payoff is monotonically increasing in τ_d. Hence, $T \in R_1(\lambda_a)$ for any λ_a. Establishing the uniqueness of T as a best response action follows the same argument used in the proof of Lemma 6. □

4.2 Special instantiation analysis

In Section 4.1, 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 $f_a(\tau_a; \lambda_a)$. Since the amount of information leakage depends in practice on the duration of collocation, here we provide an analysis for the case where $G(t)$ increases linearly in the collocation duration t. Hence, we analyze the formulated timing game for the following choice of G,

$$G(\tau_d, \tau_a) = \begin{cases} \alpha\left(\tau_d - \tau_a\right), & \tau_a \leq \tau_d \leq T \\ 0, & \text{otherwise.} \end{cases} \tag{4}$$

In Section 6.4, we provide numerical results on the best response for other (non-linear) functions, including when G scales sublinearly and quadratically in t. Without loss of generality, we always consider $\alpha = 1$. The case $\alpha \neq 1$ corresponds to the case $\alpha = 1$ with the migration cost C_d replaced by $\frac{C_d}{\alpha}$.

In our numerical evaluation we consider an exponential pdf f_a for the collocation time, i.e.,

$$f_a(\tau_a; \lambda_a) = \lambda_a e^{-\lambda_a \tau_a}, \quad \tau_a \geq 0. \tag{5}$$

This choice of $f_a(.; \lambda_a)$ is motivated by the interpretation of λ_a as the rate of attacks launched by the adversary. In other words, the attacker controls the rate of the submitted requests to the cloud server. Her requests are served within the queue of the placement engine and hence assigned to different physical machines according to a Poisson arrival process justifying the exponential arrival time.

Next, we derive sufficient conditions for the existence of a NE for the choice of functions in (4) and (5).

Theorem 8. *The 2-person nonzero-sum game defined in Section 3.2 with $G(t)$ and $f_a(\tau_a; \lambda_a)$ defined in (4) and (5) admits a pure strategy NE.*

The proof of Theorem 8 provided in the appendix rests upon proving the strict concavity of u_a and u_d, which translates into existence of a NE in pure strategy from [49, Theorem 1]. To characterize NE for both players, we start off by characterizing the best response set for each player in the following lemma whose proof follows the same argument used in the proof of Theorem 5.

Lemma 9. *For the 2-person game defined in Section 3.2 with the reward function $G(t)$ and the probability density function $f_a(\tau_a; \lambda_a)$ defined in (4) and (5), the attacker's best response pure strategy is characterized as*

- $\lambda_a^* = \lambda_{\max}$, *if* $1 - \lambda_a \tau_d e^{-\lambda_a \tau_d} - e^{-\lambda_a \tau_d} - C_a \lambda_a^2 > 0$
- $\lambda_a^* = \lambda_{\min}$, *if* $1 - \lambda_a \tau_d e^{-\lambda_a \tau_d} - e^{-\lambda_a \tau_d} - C_a \lambda_a^2 < 0$
- $\lambda_a^* = \left\{\lambda_a \mid 1 - \lambda_a \tau_d e^{-\lambda_a \tau_d} - e^{-\lambda_a \tau_d} = C_a \lambda_a^2\right\}$, *otherwise,*

for any action τ_d by the defender.

The best response strategy for the defender is characterized as

- $\tau_d^* = T$, *if* $C_d - \tau_d^2\left(1 - e^{-\lambda_a \tau_d}\right) > 0$
- $\tau_d^* = \tau_{\min}$, *if* $C_d - \tau_d^2\left(1 - e^{-\lambda_a \tau_d}\right) < 0$
- $\tau_d^* = \left\{\tau_d \mid \tau_d^2\left(1 - e^{-\lambda_a \tau_d}\right) = C_d\right\}$, *otherwise,*

for any action λ_a by the attacker.

Remark 1. *The optimal NE strategies can be obtained analytically using the players' best response curves derived in Lemma 9 (c.f. Section 6). In situations where an analytical solution is intractable, they can be obtained numerically as the equilibrium solution of convex optimization problems associated with each player. Specifically, for $\lambda_a \in \mathcal{A}_a, \lambda_a^* = \arg\max u_a$ and for $\tau_d \in \mathcal{A}_d$, $\tau_d^* = \arg\max u_d$. Since both objective functions were shown to be convex and differentiable, each player is guaranteed to converge to his optimal action for every action by the opponent. The equilibrium (τ_d^*, λ_a^*) satisfies the Lagrangian equations corresponding to both problems [50] and can be obtained by solving both problems simultaneously using standard techniques such as Newton's method with convergence guarantees [51], [52].*

The following two theorems whose proof is provided in Appendix B establish 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 10. *For the two person nonzero-sum game defined in Section 3.2 with the reward function in (4) and the exponentially distributed collocation time τ_a in (5), 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 $\lambda_a^*(\tau_d) = \lambda_{\min}$.

Theorem 11. *For the two person nonzero-sum game defined in Section 3.2 with the reward function in (4) and the exponentially distributed collocation time τ_a in (5), if*

$$C_d > T^2(1-e^{-\lambda_a T}),$$

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

5 Generalization: Game Model with IDS

In the aforementioned model, the attacker's goal is to be collocated with her victim as soon as possible before the victim is migrated. Evidently, upon collocation with her victim, the attacker will choose to reside there until τ_d since no detection mechanism is in place to urge her to evade. In this section, we extend the existing system model and consider the case in which the cloud data center is equipped with an IDS. The IDS monitors suspicious activities and captures malicious behavior of any user after a sufficient period of time δ, which is a random variable with distribution $y(\delta)$, $\delta \in [0,T]$. For useful detection, $\delta < \tau_d$. Hence, the attacker may need to stop her collocation attacks before being detected. This introduces another control variable s to be optimized by the attacker, namely how long she should continue to carry on the attack after successful collocation. The distribution $y(\delta)$ accounts for the entire range of priors between the extreme of an uninformative prior (a uniform distribution) in which the players do not have useful information about the time-to-detection δ, and the extreme of a fully degenerate distribution in which the players know δ exactly. For the latter case, the attacker will surely choose to stop after a duration δ from the onset of successful collocation, i.e., right before $\tau_a + \delta$.

Next, we modify the attacker's payoff function u_a in order to account for the probability of detection. In the event of detection, the attacker incurs a cost D (since this user will be black-listed), but her gain is in the amount of data read out until detection. Therefore, we redefine the attacker's expected reward by averaging over both the time-to-detection δ and the collocation time τ_a,

$$\begin{aligned} u_a(\tau_d,\lambda_a,s) &= \int_{\delta=s}^{T} G(s)y(\delta)d\delta \int_{\tau_a=0}^{\tau_d-s} f_a(\tau_a;\lambda_a)d\tau_a \\ &+ \int_{\delta=0}^{s} (G(\delta)-D)\left(\int_0^{\tau_d-\delta} f_a(\tau_a;\lambda_a)d\tau_a\right) y(\delta)d\delta \\ &+ \int_{\delta=0}^{s}\left(\int_{\tau_d-\delta}^{\tau_d} G(\tau_d-\tau_a)f_a(\tau_a;\lambda_a)d\tau_a\right) y(\delta)d\delta \\ &+ \int_{\delta=s}^{T}\left(\int_{\tau_d-s}^{\tau_d} G(\tau_d-\tau_a)f_a(\tau_a;\lambda_a)d\tau_a\right) y(\delta)d\delta - C_a\lambda_a. \end{aligned} \tag{6}$$

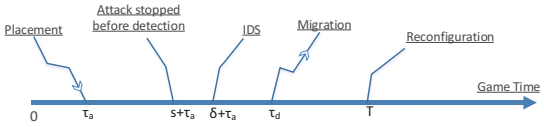

Fig. 2: Attacker evades IDS by early stopping of malicious activity.

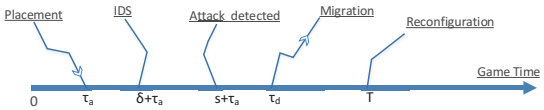

Fig. 3: Attacker detected by the IDS.

The first term in (6) accounts for the attacker's expected payoff in the event of no detection as the attacker stopped malicious activities before the IDS alarm, i.e., $s < \delta$, as illustrated in Fig. 2. The second term represents the event of detection, hence collocation ends at $\tau_a + \delta$, i.e., after a collocation duration δ as $\delta < s$, as shown in Fig. 3. Therefore, the attacker incurs a detection loss D. The third and fourth terms account for the event of no detection but due to the migration mechanism. In other words, the attacker is not identified because $\tau_d - \tau_a < \min(\delta, s)$. The last term accounts for the cost of launching the attack.

Similarly, we redefine the defender's expected payoff function,

$$\begin{aligned} u_d(\tau_d,\lambda_a,s) &= -\int_{\delta=s}^{T} G(s)y(\delta)d\delta \int_{\tau_a=0}^{\tau_d-s} f_a(\tau_a;\lambda_a)d\tau_a \\ &- \int_{\delta=0}^{s} (G(\delta)-D)\left(\int_0^{\tau_d-\delta} f_a(\tau_a;\lambda_a)d\tau_a\right) y(\delta)d\delta \\ &- \int_{\delta=0}^{s}\left(\int_{\tau_d-\delta}^{\tau_d} G(\tau_d-\tau_a)f_a(\tau_a;\lambda_a)d\tau_a\right) y(\delta)d\delta \\ &- \int_{\delta=s}^{T}\left(\int_{\tau_d-s}^{\tau_d} G(\tau_d-\tau_a)f_a(\tau_a;\lambda_a)d\tau_a\right) y(\delta)d\delta - \frac{C_d}{\tau_d}. \end{aligned} \tag{7}$$

6 Numerical Analysis

In this section, we provide numerical analysis of the proposed game model. To characterize the payoff functions for both players, we need to specify $G(t)$ and $f_a(\tau_a;\lambda_a)$. For the linear reward function $G(t)$ and the exponential density function $f_a(\tau_a;\lambda_a)$ described in (4) and (5), the reward functions can be readily expressed as

$$u_a(\tau_d,\lambda_a) = \frac{\lambda_a\tau_d + e^{-\lambda_a\tau_d} - C_a\lambda_a^2 - 1}{\lambda_a}, \tag{8}$$

$$u_d(\tau_d,\lambda_a) = \frac{1-\lambda_a\tau_d - e^{-\lambda_a\tau_d}}{\lambda_a} - \frac{C_d}{\tau_d}, \tag{9}$$

for $\tau_d \in \mathcal{A}_d, \lambda_a \in \mathcal{A}_a$. In the following analysis, we study the behavior of the payoff functions for both players. We illustrate the reward of the defender as a function of the migration time τ_d for a range of attack rates λ_a. For the attacker, we plot her reward as a function of λ_a for different τ_d. Afterwards, we investigate the effect of the migration cost C_d and the attack cost C_a on the reward functions

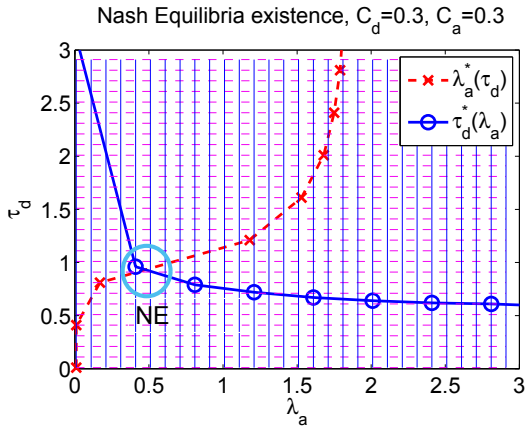

Fig. 4: For the shown action space, $\mathcal{A} = \mathcal{A}_d \times \mathcal{A}_a$, the game admits a NE in pure strategies.

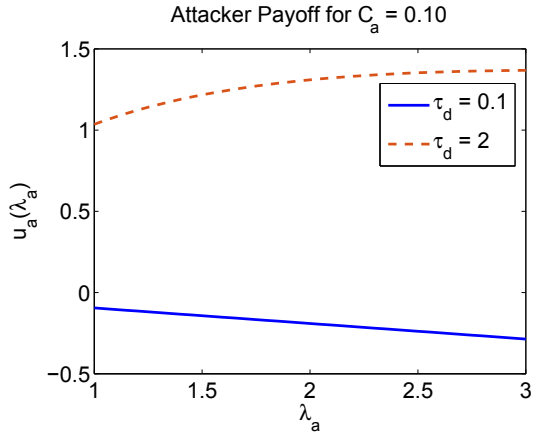

Fig. 5: At $\tau_d = 0.1$, the attacker payoff is monotonically decreasing, but not for $\tau_d = 2$, in agreement with the bound on C_a in Theorem 10.

and the players' best response curves. We also demonstrate existence of NE when the game satisfies the concavity conditions. Finally, we generalize our analysis to investigate different scaling regimes of the reward function, including sublinear and superlinear regimes.

We start our numerical analysis by reflecting on the theoretical analysis in Section 4.2. In Fig. 4, we plot the NE existence region that satisfies strict concavity of both u_a and u_d. Per Theorem 8, for G and f_a as defined in (4) and (5), the game played over the illustrated action space admits a NE in pure strategies circled in Fig. 4. The figure illustrates the best response curves along with the game action space at $C_d = 0.3$ and $C_a = 0.3$. Fig. 4 verifies our analytical results of NE existence. An equilibrium point lies at the intersection of the two best response curves for both players. By definition, this point is a NE at which each player makes the best decision taking into account the opponent's best fixed strategy. In this setting, the NE is unique – the unique intersection point of the best response curves for both players at $\tau_d^* = 0.93$ and $\lambda_a^* = 0.48$. Theorems 10 and 11 established lower bounds on C_a and C_d beyond which u_a and u_d are monotone. Fig. 5 shows the attacker's reward function at different migration rates, verifying Theorem 10. We numerically verify the monotonicity of u_a for $\tau_d = 0.1$ and $\tau_d = 2$. Let $\lambda_{\min} = 1$ and $\lambda_{\max} = 3$, hence according to Theorem 10, u_a is monotonically decreasing when $C_a > 0.04$ when $\tau_a = 0.1$. However, at $\tau_d = 2$, the attack cost $C_a > 0.98$ ensures that u_a is monotonically decreasing in λ_a. In Fig. 5 where $C_a = 0.1$, it is shown that the corresponding u_a is monotonically decreasing for all $\lambda_a \in [1, 3]$ for $\tau_d = 0.1$. At $\tau_d = 2$ when the condition on C_a is not satisfied, the payoff u_a is not monotonically decreasing. Next, we study and discuss the effect of different system parameters on the players' payoff and best response in comparison to other defense and attack policies.

6.1 Payoff functions

Fig. 6a shows the payoff function of the defender u_d versus the migration time τ_d for $C_d = 0.3$, $\tau_{\min} = 0.1$, and $T = 3$. The figure highlights the tradeoff of the defender as he seeks to optimize τ_d to secure the system through VM migration while avoiding a large overhead. Evidently, the optimal migration time τ_d^* depends on the attacker's strategy λ_a. The tradeoff shown in Fig. 6a agrees with our intuition based on the 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, with a larger τ_d, the VMs dwell for a longer duration on the same physical node leaving more room for the attacker to collocate and steal data from her target VM. In Fig. 6a, we compare the defender's reward at different attack rates λ_a. Clearly, when the attack is more aggressive, the defender is able to maximize his 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 0.8 to 0.6.

In Fig. 6b, 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 eventually becomes the dominant 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 τ_d is small (a higher migration rate), 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. For example, when τ_d is reduced from $\tau_d = 3.5$ to $\tau_d = 1.5$, the optimal attack rate that maximizes the payoff u_a decreases from $\lambda_a^* = 2.21$ to $\lambda_a^* = 2$. 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. To better understand the effect of the migration (attack) cost on the optimal migration (attack) rate for the defender (attacker), in the following two subsections we study the behavior of the payoff functions at different values of the cost. We also study the behavior of the best response curves to gain more insight into the tradeoffs associated with this game.

6.2 Cost effect and monotonicity

To show the effect of the migration and attack costs C_d and C_a, we plot the players' reward functions for different

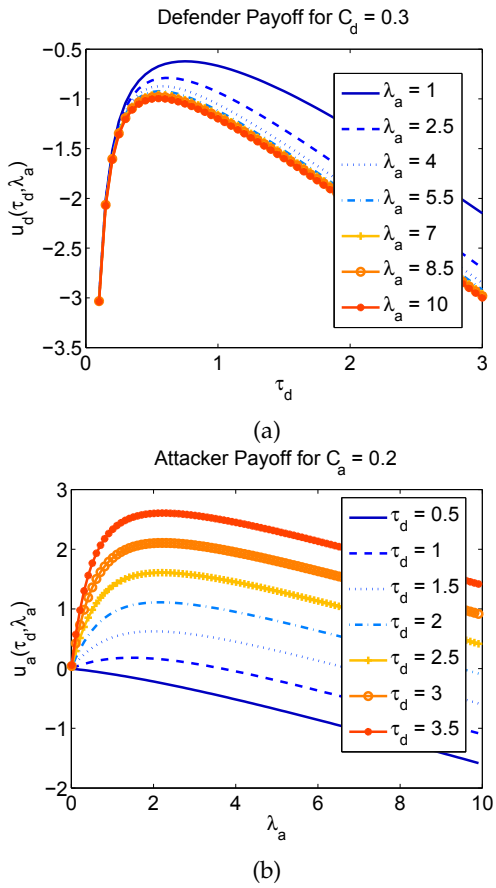

Fig. 6: (a) Defender's reward versus migration time τ_d; (b) Attacker's reward versus attack rate λ_a.

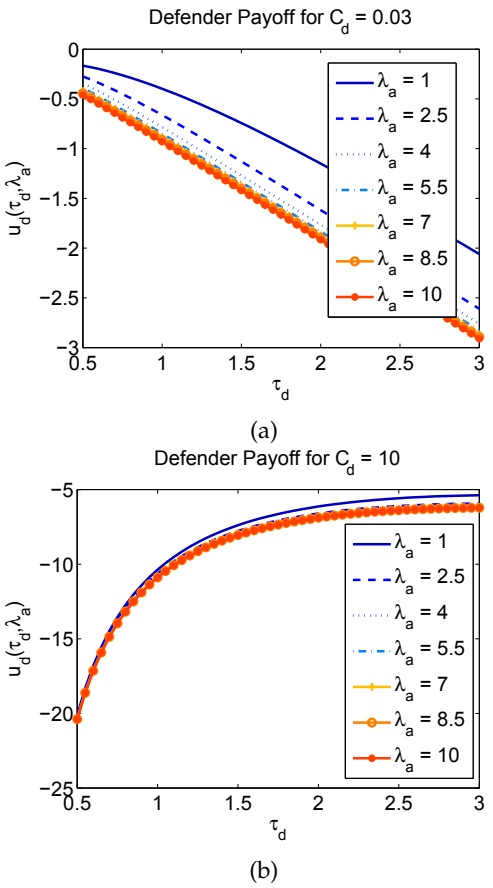

Fig. 7: Defender's reward versus migration time τ_d for (a) $C_d = 0.03$, and (b) $C_d = 10$.

values of the cost. In Fig. 7a, 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 at this small migration cost. Indeed, referring to (9), u_d is monotonically decreasing in τ_d when $C_d \to 0$. On the other hand, when the migration cost is too high as shown in Fig. 7b where $C_d = 10$, the defender's best response is $\tau_d^* = T$ to reduce the associated migration cost. We remark that the reward function is monotonically increasing in τ_d for such high migration cost, a fact which was established analytically in Theorem 11.

Similarly, the effect of the attack cost C_a can be shown in Fig. 8. At a very small attack cost, $C_a = 0.01$, as shown in Fig. 8b, the attacker's best attack strategy is to attack aggressively at $\lambda_{\max}$ to maximize the chances of successful collocation regardless of the defender's action. Recalling the attacker's payoff function in (8), u_a is monotonically increasing in λ_a when $C_a \to 0$. In case of a high attack cost, the behavior of the payoff function is reversed as shown in Fig. 8b where $C_a = 6$. In this case, 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. This behavior is confirmed by the analysis in Theorem 10.

6.3 Best response curves

In this section, we study the best response curves for both players based on Definition 1 to provide more insight into the optimal action of a player as function of the action of the opponent. The solid blue line in Fig. 9a shows the defender's best response curve τ_d^* as function of λ_a. The attacker's best response curve λ_a^* as function of the defender's action τ_d is shown in dashed red line. In this scenario, we set $T = 3$, $\lambda_{\max} = 3$, $C_d = 0.3$, and $C_a = 0.1$. In Fig. 9a, the intersection point of the two response curves is the unique NE. The point(s) of equilibria depend on the values of C_a and C_d as detailed next. The best response curves also underscore the tradeoff for each player. For example, at equilibrium the defender migrates with $\tau_d = 0.66$ while the attacker uses rate $\lambda_a = 1.8$ 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. 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 the amount of data leaked out to the attacker as long as the defender is reducing the

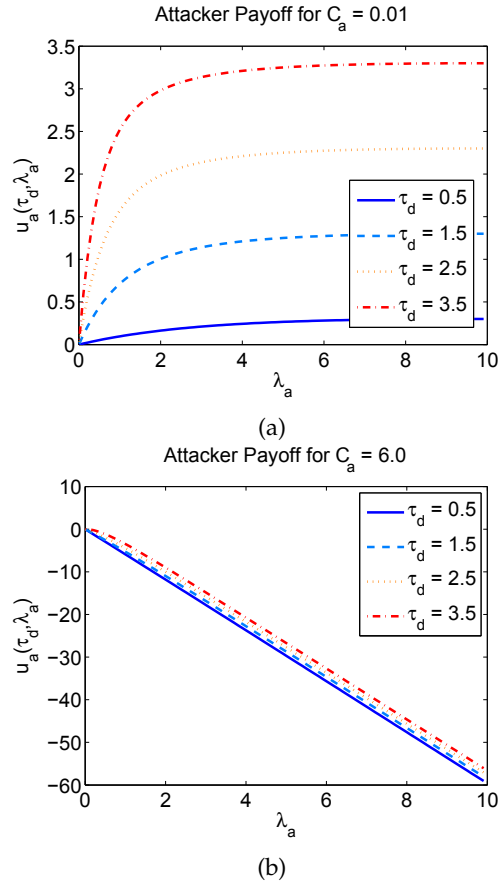

Fig. 8: Attacker's reward versus attack rate λ_a for (a) $C_a = 0.01$, and (b) $C_a = 6$.

migration rate. The best response curves also demonstrate the monotonicity of the payoff functions with respect to C_a and C_d as explained earlier in Section 6.2. To show this, Fig. 9b, 9c and 9d illustrate the best response curves at extreme cost values. In particular, in Fig. 9b, both C_d and C_a are set to zero. It is obvious that the defender is migrating with the highest permissible frequency such that, $\tau_d^* = \tau_{\min}$ for any attack rate. In response, the attacker's best action is $\lambda_a = \lambda_{\max}$ regardless of the defender's action. Hence, when the costs of migration and attack are zero, both players do not face any tradeoffs and the game is zero-sum. Fig. 9c shows another extreme scenario where only the defender faces a very high cost for migration. His best response is $\tau_d^* = T = 3$, which corresponds to the lowest migration rate possible. In Fig. 9d, the attack cost $C_a = 6$ while the defender incurs zero cost for migration. Hence, the defender adopts the highest migration rate at $\tau_d^* = \tau_{\min}$ against any attack rate. In response, it is more rewarding for the attacker to attack at $\lambda_{\min}$ for any τ_d.

6.4 Different reward scaling regimes

In the numerical analysis above, we considered the reward function G to be linearly increasing in the collocation duration. However, the reward function need not be linear. In this section, we study other scaling regimes. In particular, we consider the scenario where the reward $G(\tau_d, \tau_a)$ scales sub-linearly, quadratically or cubically with the collocation

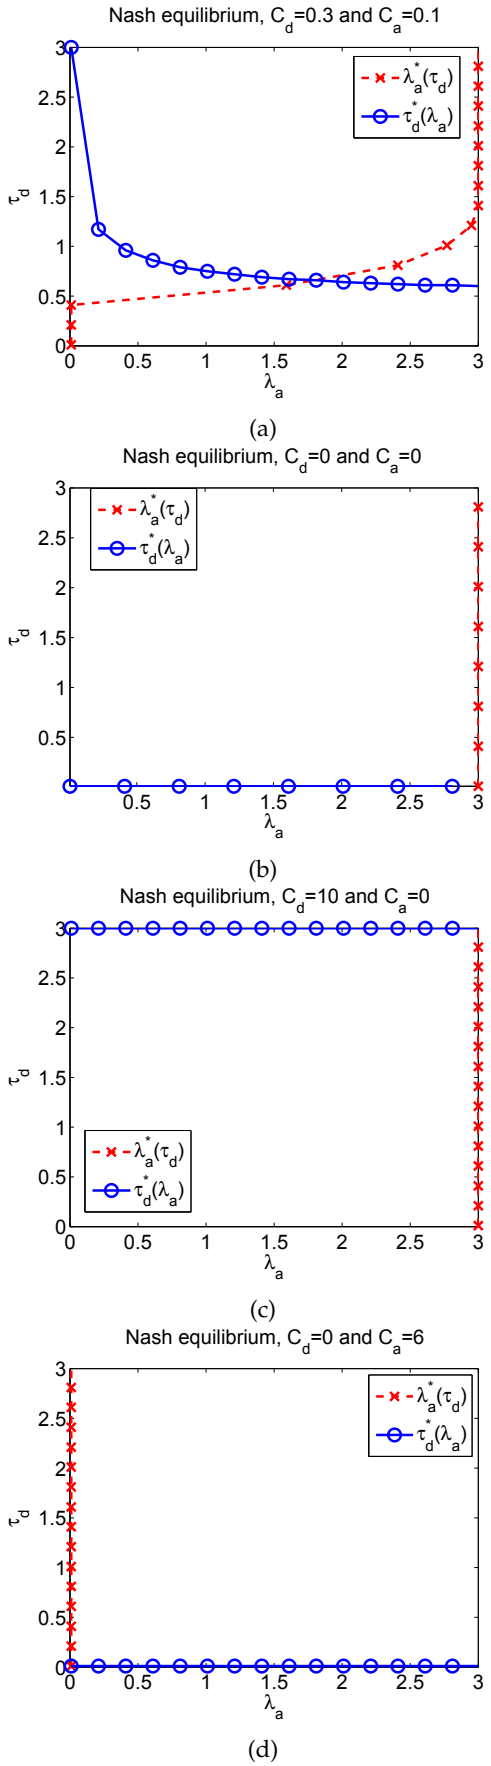

Fig. 9: Players best response curves for different cost values.

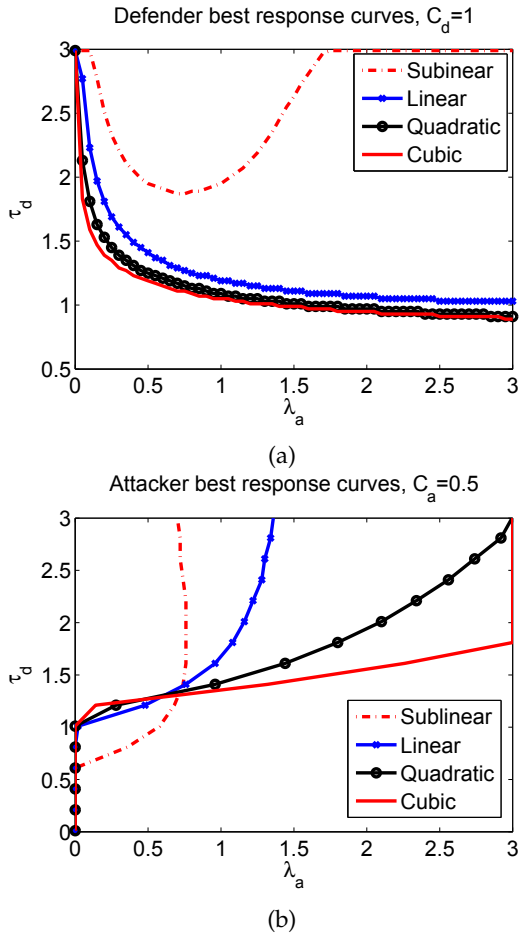

Fig. 10: Defender's (a) and Attacker's (b) best response curves for different reward scaling regimes.

duration. In other words, $G(\tau_d, \tau_a) = \max\{0, (\tau_d - \tau_a)^n\}$, where $n = \{0.2, 1, 2, 3\}$. The scaling of the reward function relates to the power of the attack as well as the vulnerability of the system under attack. The more powerful the attacker, the more data she is able to read out for a given duration of collocation. The power of the attack also depends on the side-channel technique used. The authors in [5] have shown that the amount of data leaked practically depends on the technique used to access the last level cache (LLC) (e.g., PRIME+PROBE and FLUSH+RELOAD attacks [18], [21]). For example, to read the cache, an attacker would need to adjust the time of the PROBE phase, which in turn affects the error rate of the attack covert channel.

In Fig. 10a, we plot the defender's best response curves for sub-linear, linear, quadratic, and cubic reward functions. Intuitively, higher order reward functions are more disposed to dominate the payoff functions than for the linear and sub-linear scaling. In Fig. 10a, the migration cost is set to $C_d = 1$. For the sub-linear regime, the defender's payoff function is more dominated by the migration cost. As the attack rate λ_a increases, the defender is urged to migrate the VMs at a faster rate (wherefore τ_d^* decreases), 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. For the linear regimes, the defender is facing exactly the same tradeoff discussed earlier in Section 6.3. Similarly, for the higher order reward regimes, the leakage term dominates the payoff over the entire range of attack rates. Therefore, the defender is consistently urged to increase the migration rate as the attacker increases her attack rates. With the quadratic and cubic reward functions the defender's best response is shown to exhibit a similar behavior, but conceivably the cubic reward incentivizes a faster increase in the rate of migration.

In Fig. 10b, the attacker's best response curves are plotted for different reward functions. The higher the order of the reward regime, the more is the attacker enticed to attack. For the sub-linear regime, the attacker attacks at rates higher than $\lambda_{\min}$ when $\tau_d > 0.6$. However, in the linear regime, the attacker's best response rate is non-vanishing and increasing in τ_d for $\tau_d > 1$, reaches $\lambda_a = 1.4$ as soon as the cost of the attack starts to dominate the attacker's payoff. Evidently, the higher the order of the reward, the more is the attacker willing to attack at higher rates. As shown in Fig. 10b, the cubic regime is extremely rewarding to the attacker, and as a result the attacker affords to attack at the maximum permissible rate as the reward term dominates her payoff function.

6.5 Game simulation and implementation

In this section, we compare the payoff of both players playing NE strategies to the payoffs of other defense and attack strategies. As per our theoretical analysis in Section 4, the players' optimal (NE) policies depend on the values of the associated costs C_d and C_a. Table 1 presents the results of a simulation of the game for the linear reward regime in which $G(\tau_d, \tau_a) = (\tau_d - \tau_a)^+ = \max(\tau_d - \tau_a, 0)$ at different values of C_d and C_a. For the numerical results, the maximum collocation time is set to $T = 3$ and the maximum attack rate is $\lambda_{\max} = 3$. The results underscore that a rational attacker would adapt the attack rate to the attack cost to avoid incurring high cost and/or launching useless attacks. For example, when $C_a = 8$, the payoff corresponding to the NE converges to that of the No Attack strategy, which is substantially higher than the payoff of an aggressive attacker of -23.68 due to a substantial attack cost. Similarly, the defender should not resort to very high frequency migration (equivalently, small τ_d), unless the migration cost is fairly low. For example, the results in the table show that the payoff of the defender adopting the NE policy tends to that of the No defense policy as C_d increases. The last column designated as worst case (for the defender) corresponds to the scenario where the attacker is attacking at the highest rate while the defender does not adopt any migration policy. The loss of the defender for not migrating compared to the NE strategy is more pronounced when the NE point has $\tau_d^* < T$, i.e., when the defender is in a position to defend the system through VM migrations.

Real system implementation: To demonstrate the effectiveness of the proposed approach, we implemented the migration defense approach on a proof-of-concept cloud setup using the Xen hypervisor [53], which allows us to run many instances of an operating system on a single machine. The setup is composed of five physical nodes in addition to the orchestrating controller node. The specifications of

the nodes are: Dell Inc. PowerEdge 1900 Intel(R) Xeon(R), 4-core CPU E5335 (2.00GHz) and 8GB Ram running Ubuntu 16.04.4 LTS with Xen 4.6. Each hypervisor initially runs 20 VMs. The number of VMs residing on each hypervisor changes slightly over time due to live migration. Three target VMs are uniformly distributed over the five hypervisors. We validate our results by comparing the performance of the proposed defense approach to a no-defense approach and to a random migration defense policy with a uniform distribution over the interval $[\tau_{\min}, T]$. In Fig. 13, we plot the collocation duration per target VM (left y-axis) and the average number of collocation events (right y-axis). As shown, the no-migration approach results in collocation events of longer duration. The proposed defense approach can reduce both the durations and number of occurrences of collocation events. VM migration defense policies are shown to reduce the duration of collocation events by half. In Fig. 14, we evaluate the defender's reward at different migration cost values. The proposed defense is shown to outperform the random migration policy, which does not adapt to the migration cost C_d as it chooses a random τ_d.

6.6 Extended model with IDS

In Fig 11, we compare the attacker's payoff with and without an IDS in place based on the analysis in Section 5. In this experiment, we set $D = 0.2$, $C_a = 0.2$ for different stop times s. It is clear that the IDS drastically reduces the attacker's reward. In addition, while the IDS is in place the attacker can increase her expected reward by shortening the attack duration. Hence, the defender's reward increases since the amount of data leaked out is reduced.

Fig. 12 illustrates the attacker's payoff as function of the attack stopping time s for different values of the detection cost D. Obviously, the best time to stop the attack depends on the detection cost D. As D increases, the attacker's payoff decreases and the optimal stopping time s (corresponding to the highest payoff) is shown to decrease. At a certain point, the attacker is forced to stop as soon as she collocates to evade a high penalty if detected.

7 Conclusion

In this paper, we developed a MTD framework for the VM migration timing problem. Live migration of VMs between different physical nodes is studied in a game-theoretic framework to defend multi-tenant clouds against side channel attacks launched by malicious users co-residing on the same physical node. We characterized best strategies for the players and established NE existence conditions. We also considered an extended system model in which the cloud is equipped with an IDS. The IDS is a reactive defense approach which, combined with our proactive VM migration defense approach, enhances the cloud security against side channel attacks. 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. We also demonstrated the proposed migration defense on a cloud network implemented in a Xen-based cluster. A large scale implementation in a public cloud

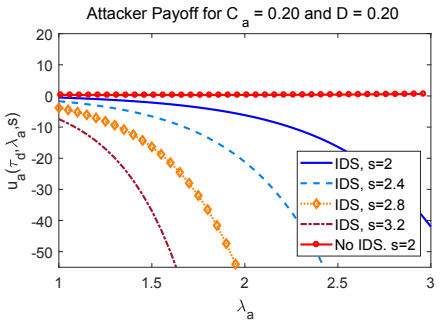

Fig. 11: Attacker's payoff with and without IDS.

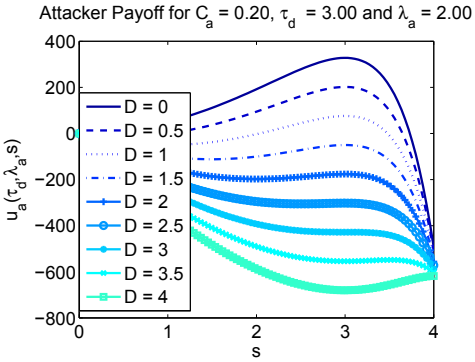

Fig. 12: Attacker's reward vs stopping time s at different costs of detection.

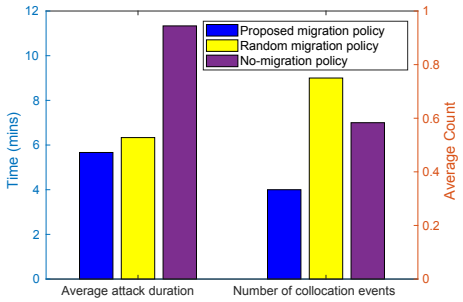

Fig. 13: Comparison between the proposed migration approach and no-migration.

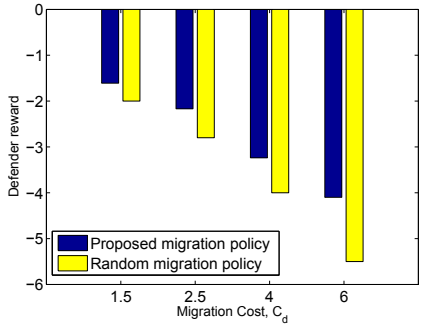

Fig. 14: Comparing the defender's reward with the proposed and the random migration policies at several migration cost values.

TABLE 1
Players' Payoff For Several Attack and Defense Strategies.

Cost		NE		No Defense		No Attack		Aggressive Attack		Worst case	
C_d	C_a	u_d	u_a	u_d	u_a	u_d	u_a	u_d	u_a	u_d	u_a
0	0	-1.49E-04	1.49E-04	-2.6778	2.6778	-5.00E-07	5.00E-07	-1.47E-04	1.47E-04	-2.6722	2.6722
0.1	0	-0.4184	0.1745	-2.711	2.6778	-0.2447	8.39E-04	-0.4164	0.1725	-2.7054	2.6722
0.1	0.1	-0.292	0.0147	-1.9578	1.8396	-0.1937	3.50E-04	-0.447	-0.0414	-2.7054	2.3762
0	0.1	-5.00E-07	-1.00E-03	-0.0448	0.0438	-5.00E-07	-1.00E-03	-1.47E-04	-0.2959	-2.6722	2.3762
0.4	0.2	-0.7561	0.0603	-2.2234	1.8825	-0.4912	0.0014	-0.9998	-0.08	-2.8051	2.0802
0.4	0.4	-0.6141	0.0331	-1.5869	1.254	-0.3864	0.0015	-1.1082	-0.4567	-2.8051	1.4882
0.4	0.6	-0.5024	0.0158	-1.0553	0.7664	-0.3379	0.0013	-1.2121	-0.8944	-2.8051	0.8962
0.8	0.6	-0.9674	0.0732	-1.7374	1.1656	-0.5974	0.0032	-1.6164	-0.7478	-2.938	0.8962
0.8	1	-0.7914	0.0345	-1.1881	0.6624	-0.5098	0.0029	-1.7719	-1.685	-2.938	-0.2878
2	4	-0.9212	0.0145	-0.9206	0.0162	-0.7112	0.0046	-3.3289	-9.1778	-3.3367	-9.1678
10	0	-6	2.6667	-5.999	2.6767	-3.3779	0.0446	-5.9955	2.6622	-5.9945	2.6722
0	8	-5.00E-07	-0.08	-0.0448	-0.0352	-5.00E-07	-0.08	-1.47E-04	-23.6799	-2.6722	-21.0078

environment is subject of future work, along with a detailed performance analysis quantifying actual costs of migration, including downtime, memory and cache usage, and application-specific performance metrics. Other avenues for future research include studying VM allocation dynamics and jointly optimizing timing and allocation policies in stochastic game model formulations.

ACKNOWLEDGEMENT

The authors would like to thank Jason High for his help with the cloud network implementation. This work was supported by NSF CAREER Award CCF-1552497, NSF grant No. CCF-1320547 and NSF CNS award No. 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] 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.
[4] 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.
[5] F. Liu, Y. Yarom, Q. Ge, G. Heiser, and R. B. Lee, "Last-level cache side-channel attacks are practical," in *Security and Privacy (SP), 2015 IEEE Symposium on*. IEEE, 2015, pp. 605–622.
[6] D. Blackwell, "The noisy duel, one bullet each," arbitrary accuracy. Technical report, The RAND Corporation, D-442, Tech. Rep., 1949.
[7] T. Radzik, "Results and problems in games of timing," *Lecture Notes-Monograph Series*, pp. 269–292, 1996.
[8] 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.
[9] 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.
[10] P. Hu, H. Li, H. Fu, D. Cansever, and P. Mohapatra, "Dynamic defense strategy against advanced persistent threat with insiders," in *Computer Communications (INFOCOM), 2015 IEEE Conference on*. IEEE, 2015, pp. 747–755.
[11] M. Zhang, Z. Zheng, and N. B. Shroff, "Stealthy attacks and observable defenses: A game theoretic model under strict resource constraints," in *Signal and Information Processing (GlobalSIP), 2014 IEEE Global Conference on*. IEEE, 2014, pp. 813–817.
[12] 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.
[13] 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.
[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 *Cloud Computing (CLOUD), 2014 IEEE 7th International Conference on*. IEEE, 2014, pp. 514–521.
[16] 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 *Cyber Security and Cloud Computing (CSCloud), 2015 IEEE 2nd International Conference on*. IEEE, 2015, pp. 382–389.
[17] S. Singh and I. Chana, "A survey on resource scheduling in cloud computing: Issues and challenges," *Journal of grid computing*, vol. 14, no. 2, pp. 217–264, 2016.
[18] 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*. ACM, 2009, pp. 199–212.
[19] 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, p. 1.
[20] R. Owens and W. Wang, "Non-interactive os fingerprinting through memory de-duplication technique in virtual machines," in *Performance Computing and Communications Conference (IPCCC), 2011 IEEE 30th International*. IEEE, 2011, pp. 1–8.
[21] 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 2012 ACM conference on Computer and communications security*. ACM, 2012, pp. 305–316.
[22] 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.
[23] B. C. Vattikonda, S. Das, and H. Shacham, "Eliminating fine grained timers in xen," in *Proceedings of the 3rd ACM workshop on Cloud computing security workshop*. ACM, 2011, pp. 41–46.
[24] P. Li, D. Gao, and M. K. Reiter, "Stopwatch: a cloud architecture for timing channel mitigation," *ACM Transactions on Information and System Security (TISSEC)*, vol. 17, no. 2, p. 8, 2014.
[25] H. Raj, R. Nathuji, A. Singh, and P. England, "Resource management for isolation enhanced cloud services," in *Proceedings of the 2009 ACM workshop on Cloud computing security*. ACM, 2009, pp. 77–84.
[26] T. Kim, M. Peinado, and G. Mainar-Ruiz, "Stealthmem: System-level protection against cache-based side channel attacks in the cloud." in *USENIX Security symposium*, 2012, pp. 189–204.
[27] Y. Zhang and M. K. Reiter, "Düppel: Retrofitting commodity operating systems to mitigate cache side channels in the cloud," in *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM, 2013, pp. 827–838.
[28] Z. Wang and R. B. Lee, "A novel cache architecture with enhanced performance and security," in *Microarchitecture, 2008. MICRO-41. 2008 41st IEEE/ACM International Symposium on*. IEEE, 2008, pp. 83–93.
[29] F. Liu and R. B. Lee, "Random fill cache architecture," in *Microarchitecture (MICRO), 2014 47th Annual IEEE/ACM International Symposium on*. IEEE, 2014, pp. 203–215.
[30] E. Pattuk, M. Kantarcioglu, Z. Lin, and H. Ulusoy, "Preventing

cryptographic key leakage in cloud virtual machines." in *USENIX Security Symposium*, 2014, pp. 703–718.

[31] Y. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart, "Cross-tenant side-channel attacks in paas clouds," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2014, pp. 990–1003.

[32] S.-J. Moon, V. Sekar, and M. K. Reiter, "Nomad: Mitigating arbitrary cloud side channels via provider-assisted migration," in *Proceedings of the 22nd acm sigsac conference on computer and communications security*. ACM, 2015, pp. 1595–1606.

[33] V. Shrivastava, P. Zerfos, K.-W. Lee, H. Jamjoom, Y.-H. Liu, and S. Banerjee, "Application-aware virtual machine migration in data centers," in *INFOCOM, 2011 Proceedings IEEE*. IEEE, 2011, pp. 66–70.

[34] T. Zhang, Y. Zhang, and R. B. Lee, "Cloudradar: A real-time side-channel attack detection system in clouds," in *International Symposium on Research in Attacks, Intrusions, and Defenses*. Springer, 2016, pp. 118–140.

[35] Y. Zhang, A. Juels, A. Oprea, and M. K. Reiter, "Homealone: Co-residency detection in the cloud via side-channel analysis," in *Security and Privacy (SP), 2011 IEEE Symposium on*. IEEE, 2011, pp. 313–328.

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

[37] Q. Sun, Q. Shen, C. Li, and Z. Wu, "Selance: Secure load balancing of virtual machines in cloud," in *Trustcom/BigDataSE/I SPA, 2016 IEEE*. IEEE, 2016, pp. 662–669.

[38] S. Achleitner, T. La Porta, P. McDaniel, S. V. Krishnamurthy, A. Poylisher, and C. Serban, "Stealth migration: Hiding virtual machines on the network," in *Infocom. IEEE*, 2017.

[39] T. Penner and M. Guirguis, "Combating the bandits in the cloud: A moving target defense approach," in *Cluster, Cloud and Grid Computing (CCGRID), 2017 17th IEEE/ACM International Symposium on*. IEEE, 2017, pp. 411–420.

[40] 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 *Cloud Computing (CLOUD), 2015 IEEE 8th International Conference on*. IEEE, 2015, pp. 556–563.

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

[42] A. Laszka, B. Johnson, and J. Grossklags, "Mitigation of targeted and non-targeted covert attacks as a timing game," in *International Conference on Decision and Game Theory for Security*. Springer, 2013, pp. 175–191.

[43] B. Johnson, A. Laszka, and J. Grossklags, "Games of timing for security in dynamic environments," in *International Conference on Decision and Game Theory for Security*. Springer, 2015, pp. 57–73.

[44] V. Pham and C. Cid, "Are we compromised? modelling security assessment games," *Decision and Game Theory for Security*, pp. 234–247, 2012.

[45] X. Feng, Z. Zheng, P. Hu, D. Cansever, and P. Mohapatra, "Stealthy attacks meets insider threats: a three-player game model," in *Military Communications Conference, MILCOM 2015-2015 IEEE*. IEEE, 2015, pp. 25–30.

[46] H. Liu, C.-Z. Xu, H. Jin, J. Gong, and X. Liao, "Performance and energy modeling for live migration of virtual machines," in *Proceedings of the 20th international symposium on High performance distributed computing*. ACM, 2011, pp. 171–182.

[47] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, and A. Warfield, "Live migration of virtual machines," in *Proceedings of the 2nd Conference on Symposium on Networked Systems Design & Implementation-Volume 2*. USENIX Association, 2005, pp. 273–286.

[48] H. Von Stackelberg, *Marktform und gleichgewicht*. J. springer, 1934.

[49] T. Başar and G. J. Olsder, *Dynamic noncooperative game theory*. SIAM, 1998.

[50] M. S. Bazaraa, H. D. Sherali, and C. M. Shetty, *Nonlinear programming: theory and algorithms*. John Wiley & Sons, 2013.

[51] E. Altman, K. Avrachenkov, and A. Garnaev, "Fair resource allocation in wireless networks in the presence of a jammer," *Performance Evaluation*, vol. 67, no. 4, pp. 338–349, 2010.

[52] A. H. Anwar, G. Atia, and M. Guirguis, "Adaptive topologies against jamming attacks in wireless networks: A game-theoretic approach," *Journal of Network and Computer Applications*, vol. 121, pp. 44–58, 2018.

[53] D. Chisnall, *The definitive guide to the xen hypervisor*. Pearson Education, 2008.

Ahmed H. Anwar (S'09) received his B.Sc. degree (with highest honors) in electrical engineering from Alexandria University, Alexandria, Egypt, in 2011 and the M.Sc. degree in wireless Information Technology from Nile University, Giza, Egypt, in 2013. He is currently a Ph.D. candidate at the Department of Electrical and Computer Engineering, University of Central Florida, Orlando, FL, USA. His research interests include game theory applications, security games in wireless networks and cloud computing.

George K. Atia (S'01-M'09-SM'18) received the B.Sc. and M.Sc. degrees from Alexandria University, Egypt, in 2000 and 2003, respectively, and the Ph.D. degree from Boston University, MA, in 2009, all in electrical and computer engineering.

He joined the University of Central Florida in Fall 2012, where he is currently an associate professor in the Department of Electrical and Computer Engineering. From Fall 2009 to 2012, he was a postdoctoral research associate at the Coordinated Science Laboratory (CSL) at the University of Illinois at Urbana-Champaign (UIUC). His research interests include statistical signal processing, machine learning, stochastic control, wireless communications, detection and estimation theory, and information theory. Dr. Atia is the recipient of many awards, including the UCF Reach for the Stars Award and the CECS Research Excellence Award in 2018, the Dean's Advisory Board Fellowship and the Inaugural UCF Luminary Award in 2017, the NSF CAREER Award in 2016, and the Charles Millican Faculty Fellowship Award (2015-2017). He organized and chaired the first IEEE GlobalSIP Symposium on Controlled Sensing for Inference in 2013.

Mina S. Guirguis is an Associate Professor of Computer Science at Texas State University, which he joined in 2006. His research is broadly driven by the interplay of security, networks and stochastic control with research contributions in the areas of Cyber-Physical Systems (CPS), Networks and Computing Systems, and Mobile Cloud Computing. His research work has been published in over fifty refereed papers, posters and journals, and one book chapter. Guirguis' research and educational activities are funded from the NSF, DoD, AFOSR, DHS, IEEE, Cisco and Texas State. Guirguis received the NSF CAREER award in 2012. Guirguis has been a visiting faculty researcher at the Air Force Research Laboratory (AFRL) and at the DHS Center for Risk and Economic Analysis of Terrorism Events (CREATE). Guirguis has a wide range of industrial experience at various companies including Fortress Technologies and Microsoft. He has served on various Technical Program Committees for many conferences, on NSF panels and on the Editorial Board for the International Journal on Advances in Networks and Services. Guirguis is serving as an Academic Alliance Member for NCWIT and as a Senior Research Fellow in the LBJ Institute for STEM Education and Research.

Guirguis earned his B.Sc. in Computer Science and Automatic Control at Alexandria University in 1999, his M.A. in Computer Science at Boston University in 2005 and his Ph.D. in Computer Science at Boston University in 2007 under Azer Bestavros and Ibrahim Matta.