

A Game Theoretic Analysis of Intrusion Detection in Access Control Systems

Tansu Alpcan and Tamer Başar

Abstract— We present a game-theoretic analysis of intrusion detection in access control systems. A security game between the attacker and the intrusion detection system is investigated both in finite and continuous-kernel versions, where in the latter case players are associated with specific cost functions. The distributed virtual sensor network based on software agents with imperfect detection capabilities is also captured within the model introduced. This model is then extended to take the dynamic characteristics of the sensor network into account. Properties of the resulting dynamic system and repeated games between the players are discussed both analytically and numerically.

I. INTRODUCTION

The increasing electronic interaction and collaboration between various organizations and economic entities on a global scale results in information management systems, which are today far more complex and sophisticated than their forerunners. Such systems have to protect the integrity and confidentiality of stored information, while enabling individual users to access the allowed data and services. Achieving these goals is only possible with a proper authentication mechanism for correctly identifying users, which is part of an access control mechanism determining what information users are entitled to access. The policy and role based access control (PRBAC) server developed by the Boeing company is a good example for this type of systems. The security of an access control system is of prime importance and is crucial for a successful operation. However, static protective measures are not sufficient to secure a complex networked system. Therefore, access control systems need intrusion detection (ID) as an integral part of their operation. Intrusion detection systems (IDSs) increase security by monitoring the events in the networked system, analyzing them for signs of security problems [1], and alerting the system administrators as appropriate.

Majority of the earlier literature on intrusion detection relies on ad-hoc schemes and experimental work. Hence, there is a need for a quantitative decision and control framework in order to address issues like attack modeling, analysis of detected threats, and decision on response actions. A rich set of tools have been developed within the game theory discipline to address problems where multiple players with different objectives compete and interact with each other on the same system, and they are successfully used in many disciplines including economics, political science,

decision theory, and control. Therefore, game theory is a strong candidate to provide the much needed mathematical framework for analysis, modeling, decision, and control processes for information security and intrusion detection. Such a mathematical abstraction is useful for generalization of problems, combining the existing ad-hoc schemes under a single umbrella, and future research. Consequently, game theory has been recently proposed by several studies for a theoretical analysis of ID [2]–[4].

In [2], application of game theory to the network security area has been discussed with a special focus on information warfare. Furthermore, several matrix games between an attacker and a defending administrator have been formulated, and their equilibrium properties investigated. In [3], the interaction between these players has been modeled as a two-player stochastic game, and the Nash equilibrium or best-response strategies have been calculated using a non-linear program. While the framework considered has been mathematically comprehensive, the approach suffers from drawbacks such as scalability and extensive computations required to find the equilibrium solution. In the study [4], on the other hand, a game theoretic approach for estimating the attacker’s intent, objective, and strategies has been discussed in detail, and a numerical example has been given.

This paper investigates a game theoretic approach for intrusion detection in access control systems by building on and extending the concepts proposed in [5]. Our goal is to establish a quantitative approach with a reasonable degree of abstraction in order to study the underlying principles for development of IDSs as well as the best ID strategies. In the next section we present the underlying mathematical model considered. The security game, the cost function, and the existence of a unique Nash equilibrium are discussed in Section III. In Section IV, we investigate system dynamics and analyze various strategies numerically, which is followed by the concluding remarks of Section V.

II. THE MODEL

Building on the framework introduced in [5], consider a distributed IDS with a network of sensors, $\mathcal{S} := \{s_1, s_2, \dots, s_{max}\}$, which we call as a *virtual sensor network* in order to distinguish it from physical sensor networks. A *virtual sensor* is defined as an autonomous software agent that monitors the system and collects data for detection purposes [6]. These sensors report possible intrusions or anomalies occurring in a subsystem of a large network using a specific technique like signature comparison, pattern detection, statistical analysis, etc. The

Research supported by The Boeing Company.

T. Alpcan and T. Başar are with the Coordinated Science Laboratory, University of Illinois, 1308 West Main Street, Urbana, IL 61801 USA. (alpcan, tbasar)@control.csl.uiuc.edu

system monitored by the IDS can be represented as a set of subsystems, $\mathcal{T} = \{t_1, t_2, \dots, t_{max}\}$, which may be targeted by an attacker. We note that these subsystems could be actual computer programs or parts of the network, as well as abstract processes distributed over multiple hosts. Define $\mathcal{I} = \{I_1, I_2, \dots, I_{max}\}$ to be the set of documented threats and detectable anomalies, which may indicate a possible intrusion, as well as various types of possible attacks. Let us associate in this context the generic term “attack” with two specific attributes: target subsystem, $t \in \mathcal{T}$, and threat or anomaly type, $I \in \mathcal{I}$. We hence define the set of attacks $\mathcal{A} = \{a_1, a_2, \dots, a_{t_{max} \times I_{max}}\}$ as the cross-product of the sets \mathcal{T} and \mathcal{I} , $\mathcal{A} := \mathcal{T} \times \mathcal{I}$.

III. THE NETWORK SECURITY GAME

We model the interaction between the attacker(s) and the IDS as a noncooperative non-zero sum game. In addition to the attacker(s) and the IDS, we introduce the sensor network as a third “fictitious” player similar to the “nature” player in standard game theory [7, p. 57]. The strategy of this player consists of a fixed probability distribution given a specific attack, and it represents the output of the sensor network during that attack. This way, we capture the imperfect conveyance of the attack information to IDS by the sensors.

A. The Security Game in Extensive Form

The finite version of the security game extends the ideas of the game in [5] by modeling the general case of multiple attackers and/or complex attacks. We can explain and illustrate the finite security game through a specific example. For simplicity let us consider a network consisting of a single subsystem and a single detectable threat, i.e. $\mathcal{A} = \{a\}$. We also limit the possible actions of the IDS to “set an alert” or “do nothing”. Thus, the strategy spaces of the attacker(s) and the IDS are $U^I = \{u_1^I, u_2^I\}$ and $U^A = \{u_1^A, u_2^A\}$, where u_2^A corresponds to “no attack”. The strategy space of the sensor network is then $U^S = \{\mathbf{p}_1, \mathbf{p}_2\}$ given $u^A \in U^A$. Here, $\mathbf{p} = [p_1, p_2] \in \mathbb{R}^2$, $\mathbf{p} \geq 0$, $p_1 + p_2 = 1$, where p_1 and p_2 give the likeliness of an attack and of no attack at all, respectively. A representation of this game in extensive form is shown in Figure 1 using the GAMBIT software [8]. The payoff or benefit values for the IDS and the attacker are chosen for illustrative purposes and given by $[(R_1^A, R_1^I), \dots, (R_8^A, R_8^I)]$.

Let us further explain the game in Figure 1 by describing a specific scenario step by step, which corresponds to following a path from left to right in accordance with the order of players’ actions. The lower left branch in the figure labeled A indicates an attack by the attacker(s) to the system. The sensor network labeled as “chance” detecting this attack is represented by the $Sensor_A1$ branch. Finally, given the information from the sensor network, the IDS decides in branch U to take a predefined response action. The outcome of this scenario is quantified by a benefit of -5 to the attacker and $+5$ to the IDS.




Fig. 1. The finite version of the security game example shown in extensive form.

We next investigate the existence of a Nash equilibrium (NE) as in [5]. A Nash equilibrium for a 2-player game is defined as a pair of strategies and the corresponding pair of costs, with the property that no player can benefit by modifying its own strategy while the other player keeps her/his fixed. Hence, NE provides a suitable solution for the analysis of the security game. This particular game does not admit any Nash equilibrium solution in pure strategies, and hence, we extend the analysis by considering mixed strategies of the attacker and the IDS defined as probability distributions on the space of their pure strategies [7, p. 23]. Solving the game on this extended strategy space using GAMBIT, we obtain a unique NE in mixed strategies which also corresponds to the unique solution in behavioral strategies. Figure 1 depicts the game in extensive form and displays the probability values associated with the NE strategies of the players under the branches. Note that, unlike the other two players, the sensor network -chance player- is associated with a predefined probability distribution, which models the imperfect flow of information from the attacker to the IDS. In the NE, the attacker(s) target the system with a probability 0.20. A reason for this low probability is the discouraging effect of the sensor network’s capability of correct detection with probability 2/3. We note that there are two information sets for the IDS, one indicating an attack and one for no alarm. The NE strategy of the IDS given this information by the sensor network is “no response” (NR) with probability 1 if there is no alarm, and a response (U) with probability 0.86 if an alarm is set. We can argue that the IDS in this case has a high degree of trust on the information conveyed by the sensor network. However, it is important to note that the NE strategies of the players are very much dependent on the outcome payoffs of the game [5] as well as the detection probability distribution of the sensor network. Thus, it is crucial for correct analysis that the payoff values in the game reflect the trade-offs of the system at hand. A possible way of achieving this may be to utilize a supervised learning scheme to approximate the actual player payoffs and detection capabilities of the sensors.

Although the finite version of the security game provides

a detailed visualization of the interaction between the players, it has some limitations and disadvantages. One drawback is the scalability. The strategy spaces of the attacker and the IDS become too large for a more comprehensive analysis of a larger system. Another disadvantage is in the choice of the payoff values, which have to be determined separately for each branch of the game tree. This process may become tedious and inaccurate for a large system. In order to address these limitations, we next investigate a continuous-kernel version of the security game which is slightly different from the one above. In this game we adopt the convention that the players are minimizers (of costs) rather than being maximizers (of payoffs).

B. The Cost Functions of the Security Game

We address various security tradeoffs and establish the continuous-kernel security game by associating specific cost functions with the IDS and the attacker. Given the set of attacks with cardinality A_{max} , the strategy space of the attacker is defined as $U^A := \{\mathbf{u}^A \subset \mathbb{R}^{A_{max}} : u_i^A \geq 0, i = 1, \dots, A_{max}\}$. Similarly, the strategy space of the IDS is given by $U^I := \{\mathbf{u}^I \subset \mathbb{R}^{R_{max}} : u_i^I \geq 0, i = 1, \dots, R_{max}\}$, where R_{max} is the cardinality of the set of responses available to the IDS. The actions of the sensor network, on the other hand, belong to the space $U^S := \{\mathbf{u}^S \subset \mathbb{R}^{A_{max} \times A_{max}} : 0 \leq u_i^S \leq 1 \forall i\}$, and can be represented conveniently in matrix form by $\bar{P} := [\bar{p}_{ij}]$, $\bar{P} \in U^S$, $i, j = 1, \dots, A_{max}$. The matrix \bar{P} represents how well the sensor network detects the attacks on the average, and maps the actions of the attacker to the sensor output. Furthermore, we define a simple metric for the detection of each attack, $a \in \mathcal{A}$, monitored by the sensor network

$$dq(i) := \frac{\bar{p}_{ii}}{\sum_{j=1}^{A_{max}} \bar{p}_{ij}}, \quad i = 1, \dots, A_{max}.$$

For notational convenience, let us also define the matrix

$$P := [p_{ij}] = \begin{cases} p_{ij} = -\bar{p}_{ij} & \text{if } i = j \\ p_{ij} = \bar{p}_{ij} & \text{if } i \neq j \end{cases}. \quad (1)$$

We now introduce the cost parameters, which we take to be nonnegative. Let $\mathbf{c}^I := [c_1^I, \dots, c_{A_{max}}^I]$ represent the cost of each attack for the IDS, whereas $\mathbf{c}^A := [c_1^A, \dots, c_{max}^A]$ quantifies the gain of the attacker from the attack, if it is successful. The nonnegative matrix Q with diagonal entries greater than or equal to 1 models the vulnerability of a specific subsystem to attacks. On the other hand, the matrix $\bar{Q} := [\bar{Q}]_{A_{max} \times R_{max}}$ with entries of ones and zeros correlates IDS response actions with the attacks. The vectors $\alpha := [\alpha_1, \dots, \alpha_{R_{max}}]$ and $\beta := [\beta_1, \dots, \beta_{A_{max}}]$ are the cost of the response and the cost of the effort required to carry out an attack for the IDS and the attacker, respectively. The cost of false-alarms and capture as well as the benefit of detection and deception for the IDS and the attacker are associated with the scalar value γ . Consequently, we define

the cost function of the IDS, $J^I(\mathbf{u}^A, \mathbf{u}^I, P)$, and the one of the attacker(s), $J^A(\mathbf{u}^A, \mathbf{u}^I, P)$, as

$$J^I(\mathbf{u}^A, \mathbf{u}^I, P) := \gamma(\mathbf{u}^A)^T P \bar{Q} \mathbf{u}^I + (\mathbf{u}^I)^T \text{diag}(\alpha) \mathbf{u}^I + \mathbf{c}^I (Q \mathbf{u}^A - \bar{Q} \mathbf{u}^I), \quad (2)$$

and

$$J^A(\mathbf{u}^A, \mathbf{u}^I, P) := -\gamma(\mathbf{u}^A)^T P \bar{Q} \mathbf{u}^I + (\mathbf{u}^A)^T \text{diag}(\beta) \mathbf{u}^A + \mathbf{c}^A (\bar{Q} \mathbf{u}^I - Q \mathbf{u}^A), \quad (3)$$

where $(x)^T$ denotes the transpose of the vector or matrix, and $\text{diag}(x)$ is a diagonal matrix with the diagonal entries given by the elements of the vector x .

With these specific structures of the cost functions J^I and J^A , we attempt to capture various aspects of the security game between the attacker and the IDS. The first terms of each cost function, $\gamma(\mathbf{u}^A)^T P \bar{Q} \mathbf{u}^I$ and $-\gamma(\mathbf{u}^A)^T P \bar{Q} \mathbf{u}^I$ represent the cost of false-alarms and benefit of detection for the IDS as well as the cost of capture and benefit of deception for the attacker, respectively. Notice that, this part of the cost is zero sum. The second terms $(\mathbf{u}^I)^T \text{diag}(\alpha) \mathbf{u}^I$ and $(\mathbf{u}^A)^T \text{diag}(\beta) \mathbf{u}^A$ quantify the cost of specific responses and attacks. Depending on the response action, this reflects the cost of the use of resources, possible restrictions on system usage, or sensor reconfigurations for the IDS. On the other hand, it represents for the attacker the cost of resources required by the attack. The last terms $\mathbf{c}^I (Q \mathbf{u}^A - \bar{Q} \mathbf{u}^I)$ and $\mathbf{c}^A (\bar{Q} \mathbf{u}^I - Q \mathbf{u}^A)$ give the actual cost or benefit of a successful attack. False alarms and detection capabilities of the sensor network at a given time are incorporated into the values of the matrix P . In the ideal case of the sensor network functioning perfectly, i.e. no false alarms and 100% detection, the matrix $-P$ is equal to the identity matrix, $Id = \text{diag}([1, \dots, 1])$.

For notational convenience, define the vectors $\theta^I(\mathbf{c}^I, \bar{Q}, \alpha) := [(c^I \bar{Q})_1 / (2\alpha_1), \dots, (c^I \bar{Q})_{R_{max}} / (2\alpha_{R_{max}})]$ and $\theta^A(\mathbf{c}^A, Q, \beta) := [(c^A Q)_1 / (2\beta_1), \dots, (c^A Q)_{A_{max}} / (2\beta_{A_{max}})]$. The reaction functions of the attacker and the IDS are obtained by minimizing the respective strictly convex cost functions (2) and (3). Hence, they are uniquely given by $\mathbf{u}^I(\mathbf{u}^A, P) = [u_1^I, \dots, u_{R_{max}}^I]^T$ and $\mathbf{u}^A(\mathbf{u}^I, P) = [u_1^A, \dots, u_{max}^A]^T$, respectively, where

$$\mathbf{u}^I(\mathbf{u}^A, P) = [\theta^I - \gamma[\text{diag}(2\alpha)]^{-1} \bar{Q}^T P^T \mathbf{u}^A]^+ \quad (4)$$

and

$$\mathbf{u}^A(\mathbf{u}^I, P) = [\theta^A + \gamma[\text{diag}(2\beta)]^{-1} P \bar{Q} \mathbf{u}^I]^+. \quad (5)$$

The function denoted by $[x]^+$ maps all negative values of x to zero. It is desirable for the IDS that the sensor grid is configured such that all possible threats are covered. It is also natural to assume a worst-case scenario where for each attack (type) targeting a subsystem there exists at least one attacker who finds it beneficial for him to attack. Hence, we expect in many practical cases $u_i^A > 0 \forall i$ or $u_j^I > 0 \forall j$.

C. Existence and Uniqueness of a Nash Equilibrium

The Nash Equilibrium (NE) which has been widely utilized in noncooperative game theory is also a useful concept for the analysis of the continuous-kernel security game. Within the context of the security game defined in Section III-B, a pair of strategies $(\mathbf{u}^{I*}, \mathbf{u}^{A*})$ of the IDS and the attacker is in NE if it satisfies $\mathbf{u}^{I*} = \arg \min_{\mathbf{u}^I} J^I(\mathbf{u}^{A*}, \mathbf{u}^I, P)$ and $\mathbf{u}^{A*} = \arg \min_{\mathbf{u}^A} J^A(\mathbf{u}^{I*}, \mathbf{u}^A, P)$.

Theorem III.1. *There exists a unique NE in the security game defined in Section III-B. Furthermore, if*

$$\gamma < \min \left(\frac{\min_i \theta^I}{\left[\max_i (\text{diag}(2\alpha))^{-1} \bar{Q}^T P^T \theta^A \right]^+}, \frac{\min_i \theta^A}{\left[\max_i (\text{diag}(2\beta))^{-1} (-P) \bar{Q} \theta^I \right]^+} \right), \quad (6)$$

then the NE is an inner solution, $\mathbf{u}^{I*} > 0$ and $\mathbf{u}^{A*} > 0$, and is given by

$$\mathbf{u}^{A*} = (Id + Z)^{-1} \cdot [\theta^A + \gamma [\text{diag}(2\beta)]^{-1} P \bar{Q} \theta^I] \quad (7)$$

and

$$\mathbf{u}^{I*} = (Id + \bar{Z})^{-1} \cdot [\theta^I - \gamma [\text{diag}(2\alpha)]^{-1} \bar{Q}^T P^T \theta^A], \quad (8)$$

where $Z := \gamma^2 [\text{diag}(2\beta)]^{-1} P \bar{Q} [\text{diag}(2\alpha)]^{-1} \bar{Q}^T P^T$, $\bar{Z} := \gamma^2 [\text{diag}(2\alpha)]^{-1} \bar{Q}^T P^T [\text{diag}(2\beta)]^{-1} P \bar{Q}$, and Id is the identity matrix.

Proof. The existence of a NE in the game follows from the facts that the objective functions are strictly convex, they grow unbounded as $|u| \rightarrow \infty$, and the constraint set is convex [7, p. 174]. We next establish a unique strictly positive (equivalently inner) NE under the given sufficient condition. Let $\bar{\nabla}$ be the pseudo-gradient operator, defined through its application on the cost vector $J := [J^I \ J^A]$, as

$$\bar{\nabla} J := \left[\nabla_{u_1^I}^T J^I \dots \nabla_{u_{R_{max}}^I}^T J^I \quad \nabla_{u_1^A}^T J^A \dots \nabla_{u_{A_{max}}^A}^T J^A \right]^T, \quad (9)$$

and define $g(\mathbf{u}) := \bar{\nabla} J$ where $\mathbf{u} := [\mathbf{u}^I \ \mathbf{u}^A]$. Let $G(\mathbf{u})$ be the Jacobian of $g(\mathbf{u})$ with respect to \mathbf{u} . Define the symmetric matrix $\mathcal{G}(\mathbf{u}) := \frac{1}{2}(G(\mathbf{u}) + G(\mathbf{u})^T)$. It immediately follows that $\mathcal{G}(\mathbf{u}) = \text{diag}([\alpha \ \beta])$, which is positive definite. Thus, due to the positive definiteness of the Hessian-like matrix $\mathcal{G}(\mathbf{u})$, the game admits a unique NE solution [9]. Note that this result does not use the condition (6) on γ , which however comes into picture if we further look for an inner solution as discussed below.

We now obtain an analytical description of the inner NE solution. Let us substitute for \mathbf{u}^I in (5) the expression in (4). Hence, we obtain a fixed-point equation $\mathbf{u}^{A*} = \mathbf{u}^A(\mathbf{u}^I(\mathbf{u}^{A*}, P), P)$, given by

$$\mathbf{u}^{A*} = \theta^A + [\text{diag}(\frac{2\beta}{\gamma})]^{-1} P \bar{Q} \theta^I - \text{diag}(\frac{\gamma}{2\beta}) P \bar{Q} [\text{diag}(\frac{2\alpha}{\gamma})]^{-1} \bar{Q}^T P^T \mathbf{u}^{A*}. \quad (10)$$

Solving for \mathbf{u}^{A*} yields (7) where the inverse exists because Z is nonnegative definite. The equilibrium solution \mathbf{u}^{I*} in (8), on the other hand, can be derived by simply substituting for \mathbf{u}^{A*} from (7) into (4). It is then straightforward to show that if (6) holds then $\mathbf{u}^{I*} > 0$, and hence the NE is strictly positive. Moreover, there cannot be a boundary solution in this case due to the uniqueness of the NE. As a result, the game admits a unique inner NE under (6). \square

IV. THE SYSTEM DYNAMICS AND REPEATED GAMES

We consider a discrete-time system model in order to capture dynamic nature of the system and take into account the interactions between players over a time period. Dynamics such as varying detection capability and (re)configuration of the sensor network given the strategies of the attacker and the IDS are quantified through the entries of the \bar{P} matrix.

Let us define n as the time variable, and κ , δ , and ε as (small) scalar positive parameters. We define the random matrix $W := [w_{ij}]$, $i = 1, \dots, A_{max}$, $j = 1, \dots, R_{max}$ where w_{ij} 's are independent uniformly distributed on the interval $[-1, 1]$. Hence, W models the transients and imperfect nature of the sensor grid. Similarly, define ω as a scalar random variable uniformly distributed on the interval $[-1, 1]$, and independent of w_{ij} 's. Let us also define an upper bound, $dt_{max} < 1$, and a lower bound $dt_{min} > 0$ on the elements of \bar{P} . In doing so we can model the cases where sensors have a limited detection capability. A possible dynamic equation for \bar{P} is then given by

$$\bar{P}(n+1) = \left[\bar{P}(n) + 2\delta(\omega + \kappa)(\text{diag}(\text{diag}(\mathbf{u}^A) \bar{Q} \mathbf{u}^I) - \delta \text{col}(\text{diag}(\mathbf{u}^A) \bar{Q} \mathbf{u}^I)) + \varepsilon W(n) \right]^N, \quad (11)$$

where $\text{col}(x)$ is an $A_{max} \times A_{max}$ matrix with repeating x vectors constituting the columns, and the normalization function $[x]^N$ maps entries of x onto the interval $[dt_{min}, dt_{max}]$. With \bar{P} generated by (11), $P(n)$ can then be obtained directly from (1). The dynamics in (11) represent a somewhat optimistic point of view, as it models a situation where a past attack and follow-up response result in better detection capabilities for the sensor network. A justification for this is the efficient reconfiguration of the sensors or direct intervention by the system administrator.

We consider repeated games as a simple and suitable dynamic model where the attacker and the IDS make instantaneous myopic optimizations given the state of the system (performance of the sensor network). Consequently, the set of equations characterizing the dynamic game consist of (11) and

$$\mathbf{u}^I(n+1) = \left[\frac{\mathbf{c}^I \bar{Q}}{2\alpha} - [\text{diag}(\frac{2\alpha}{\gamma})]^{-1} (P(n))^T \mathbf{u}^A(n) \right]^+ \\ \mathbf{u}^A(n+1) = \left[\frac{\mathbf{c}^A Q}{2\beta} + [\text{diag}(\frac{2\beta}{\gamma})]^{-1} P(n) \mathbf{u}^I(n) \right]^+, \quad (12)$$

where P is related to \bar{P} through (1). Existence of a unique NE for a fixed P (\bar{P}) has already been established in Theorem III.1. Consequently, we investigate the convergence and stability properties of the system (12). Let us define

$$Idl := [idl_{ij}] = \begin{cases} dt_{max} & \text{if } i = j \\ dt_{min} & \text{if } i \neq j \end{cases},$$

which sets a limit on the best-case scenario in terms of detection capabilities of the sensor network. It immediately follows from (11) that

$$\begin{aligned} |\bar{p}_{ij}(n+1) - Idl_{ij}| &< |\bar{p}_{ij}(n) - Idl_{ij}| + \varepsilon |w_{ij}(n)| \\ &\quad + \delta \xi |\omega(n)| \\ &< |\bar{p}_{ij}(n) - Idl_{ij}| + \varepsilon + \delta \xi, \end{aligned}$$

where

$$\xi := \max_{i,j,n} \left| \begin{aligned} &2\text{diag}(\text{diag}(\mathbf{u}^A(n))\bar{Q}\mathbf{u}^I(n)) \\ &- \text{col}(\text{diag}(\mathbf{u}^A)\bar{Q}\mathbf{u}^I) \end{aligned} \right|_{i,j}.$$

Hence, if $\varepsilon = 0$ and $\omega(n) = 0 \forall n$ then as $n \rightarrow \infty$ $\bar{P}(n)$ clearly converges to the Idl matrix. Furthermore, for small fixed $\delta, \varepsilon > 0$, and starting from any feasible initial point, $\bar{P}(0) \in U^S$, $E[\bar{P}(n)]$ converges asymptotically to the region

$$Reg(\varepsilon) := \{[\bar{p}_{ij}] \in U^S : dt_{max} - (\varepsilon + \delta \xi) \leq \bar{p}_{ii} \leq dt_{max} \text{ and } 0 < dt_{min} \leq \bar{p}_{ij} \leq dt_{min} + \varepsilon + \delta \xi \forall i \neq j\}.$$

A. Dynamic Strategies and Numerical Analysis

We analyze some simple strategies available to the attacker and the IDS within the dynamic model (12) in order to gain further insight into IDS and attacker behaviors. Let us first consider strategies with fixed actions over a finite time period. Assume that the attacker starts an attack at a given time with action $\mathbf{u}^A(n)$ and sustains it over a fixed time period N such that $\mathbf{u}^A(n) = \mathbf{u}^A(n+1) = \dots = \mathbf{u}^A(n+N)$. Then, given $\mathbf{u}^A(t)$, $t = n, n+1, \dots, n+N$, the response of the IDS will be according to (4), as the IDS cannot know the fact that the attacker has chosen a fixed strategy. From (12) we immediately conclude that $J^A(n)$ will increase with n and $J^I(n)$ will decrease with n . In other words, it is suboptimal for the attacker to have a fixed action strategy. Similarly, deploying a fixed response strategy is suboptimal for the IDS as it gives the attacker an opportunity to exploit the weaknesses of the sensor network [10]. Another problem with choosing a fixed action strategy for both players is determining what this strategy should be. As we will soon demonstrate, any deviation from NE response results in higher costs for the player. Therefore, it is beneficial for both the IDS and the attacker to frequently update their strategies as part of a multi-step optimization process.

The conclusions of the discussion above can be illustrated through numerical analysis. We choose a simple scenario with three specific attacks monitored by the sensor network. For comparison purposes cost parameters are chosen to be

the same for both the attacker and the IDS: $\mathbf{c}^I = \mathbf{c}^A = [50, 50, 50]$, $\alpha = \beta = [10, 10, 10]$, $\gamma = 10$, $\varepsilon = 0.01$, $\kappa = 0.1$, and $\delta = 0.001$. The responses of the IDS are also limited to three, and $\bar{Q} = Q = Id$ for simplicity. In addition, $\bar{p}_{ij} \in [0.3, 0.7]$. We first simulate the system described in (12). The costs and actions of the attacker and the IDS as well as detection quality of sensors are shown in Figure 2.




Fig. 2. Dynamics of the system (12).

In the next simulation, we fix the IDS response as $u^I = [5, 5, 5]$, which is roughly equal to the NE solution for the static game. From Figure 3, we observe that the attacker can exploit this by limiting his/her attack to a short time period. Furthermore, temporary degradations in sensor detection quality are utilized by the attacker to decrease his own cost, while they drive the IDS cost higher. We also investigate the cases when IDS response is chosen as $u^I = [8, 8, 8]$ and $u^I = [2, 2, 2]$ respectively. We observe in the former scenario that the IDS can increase the cost of the attacker if it accepts a significant cost for itself also. On the other hand, the latter case where the IDS does not take sufficient precautions proves to be very costly for it. Clearly, both of these suboptimal fixed actions result in higher costs for the IDS while benefiting the attacker.




Fig. 3. Simulation of the system (12) with the IDS's actions fixed as $u^I = [5, 5, 5]$.

In addition, we analyze the case where the attacker deploys a fixed strategy $u^A = [5, 5, 5]$ while the IDS adjusts its actions according to (12). For clarity of presentation we

choose $\kappa = 0.3$. It is observed that the cost of the attacker increases as the quality of sensor detection improves over time. Thus, it is a better strategy for the attacker to deploy short high intensity attacks intermittently over a time period.

Next, we investigate what happens if the attacker discovers an inherent vulnerability in the system monitored by the IDS. In order to capture this scenario within our model we increase Q to $\text{diag}([2, 1, 1])$ after a fixed time point. As shown in Figure 4, the cost of the IDS increases significantly after the discovery of the vulnerability by the attacker. On the other hand, increased attack intensity on the first subsystem results in a stronger IDS response. We note the increased variation in the detection quality of the specific attack, which is due to the random imperfections in sensor reconfiguration mechanism.




Fig. 4. Simulation of the system (12) when Q is modified as $Q = \text{diag}([2, 1, 1])$ after a time point.

Finally, the inherent assumption that both the attacker and the IDS have perfect knowledge on the performance of the sensor network is relaxed. Figure 5a depicts the NE costs of both players and the difference between the two costs under the assumption that IDS estimates (from left to right) a perfectly functioning sensor network ($\bar{P} = Id$) to the worst-case ($\bar{P} = \text{Ones} - Id$), where Ones is the matrix of ones. The counterpart of this for the attacker is also depicted in Figure 5b. Clearly, a correct estimation of \bar{P} decreases the difference between the costs, which is beneficial to the player. We also observe that assuming a perfect detection the IDS can increase both its and the attacker's cost, and hence, discouraging an attack at his own expense. Likewise, Figure 5b shows that incentive to attack varies inversely proportionally to how the attacker perceives the success rate of the sensor network. As expected, IDS having a good sensor network discourages the attacker.

V. CONCLUSIONS

We have presented a game theoretic approach to intrusion detection in access control systems. Modeling the interaction between the attacker(s) and the IDS as both




Fig. 5. The NE costs of both players under the assumption that the IDS (the attacker) estimates -from left to right- a perfectly functioning sensor network ($\bar{P} = Id$) to the worst-case ($\bar{P} = \text{Ones} - Id$).

finite and continuous-kernel noncooperative security games, we have established a quantitative mathematical framework which provides insight into and addresses a wide range of resource allocation problems in intrusion detection. The imperfect flow of information from the attacker to the IDS through a virtual sensor network is also captured within this framework. Existence of a unique Nash equilibrium and best-response strategies for players under specific cost functions are investigated. In addition, the interaction between the players over a time period is analyzed using repeated games and a specific dynamic model for the sensor network. Finally, some basic strategies for the IDS and the attacker are discussed through several numerical studies.

REFERENCES

- [1] R. Bace and P. Mell, "Intrusion detection systems," NIST Special Publication on Intrusion Detection Systems, <http://www.snort.org/docs/nist-ids.pdf>.
- [2] D. A. Burke, "Towards a game theoretic model of information warfare," Master's thesis, Air Force Institute of Technology, Air University, November 1999.
- [3] K.-W. Lye and J. Wing, "Game strategies in network security," in *Foundations of Computer Security Workshop in FLoC'02*, Copenhagen, Denmark, July 2002.
- [4] P. Liu and W. Zang, "Incentive-based modeling and inference of attacker intent, objectives, and strategies," in *Proc. of the 10th ACM Computer and Communications Security Conference (CCS'03)*, Washington, DC, October 2003, pp. 179–189.
- [5] T. Alpcan and T. Başar, "A game theoretic approach to decision and analysis in network intrusion detection," in *Proc. of the 42nd IEEE Conference on Decision and Control*, Maui, HI, December 2003, pp. 2595–2600.
- [6] D. Zamboni, "Using internal sensors for computer intrusion detection," Ph.D. dissertation, Purdue University, August 2001.
- [7] T. Başar and G. J. Olsder, *Dynamic Noncooperative Game Theory*, 2nd ed. Philadelphia, PA: SIAM, 1999.
- [8] Gambit, "Gambit game theory analysis software and tools," <http://www.hss.caltech.edu/gambit>, 2002.
- [9] J. B. Rosen, "Existence and uniqueness of equilibrium points for concave n-person games," *Econometrica*, vol. 33, pp. 520–534, July 1965.
- [10] K. M. Tan, K. S. Killourhy, and R. A. Maxion, "Undermining an anomaly-based intrusion detection system using common exploits," in *Recent Advances in Intrusion Detection : 5th International Symposium, RAID 2002, Zurich, Switzerland, October 16-18, 2002. Proceedings*, ser. Lecture Notes in Computer Science. Springer-Verlag Heidelberg, January 2002, vol. 2516 / 2003, pp. 54 – 73.