

A Game Theoretic Investigation of Deception in Network Security

Thomas E. Carroll

Department of Computer Science,
Wayne State University
5143 Cass Avenue,
Detroit, MI 48202 USA.
Email: tec@cs.wayne.edu

Daniel Grosu

Department of Computer Science,
Wayne State University
5143 Cass Avenue,
Detroit, MI 48202 USA.
Email: dgrosu@cs.wayne.edu

Abstract—We perform a game theoretic investigation of the effects of deception on the interactions between an attacker and a defender of a computer network. The defender can employ camouflage by either disguising a normal system as a honeypot, or by disguising a honeypot as a normal system. We model the interactions between defender and attacker using a signaling game, a non-cooperative two player dynamic game of incomplete information. For this model, we determine which strategies admit perfect Bayesian equilibria. These equilibria are refined Nash equilibria in which neither the defender nor the attacker will unilaterally choose to deviate from their strategies. We discuss the benefits of employing deceptive equilibrium strategies in the defense of a computer network.

I. INTRODUCTION

Defenders can employ deception to increase the effectiveness of their protections. Deception impedes attacks by increasing their costs. Attackers require additional resources (e.g., time, money) to comprehend the situation and to tailor their attacks to the specifics that the situation demands. Deception has a long history of effective use within the military and is now being deployed for the protection of information systems [1], [2].

It is common practice nowadays for a defender to deploy honeypots within her network. A *honeypot* is a computer system that is a trap to detect unauthorized accesses [3]. Unlike normal systems, honeypots produce a rich source of information detailing the attack methods used when attackers attempt to compromise them. As a consequence, attackers have reasons to avoid attacking them [4].

Attackers may be able to determine if a system is a honeypot by considering clues such as slow I/O and other time delays, unusual system calls, temptingly obvious file names (e.g., “passwords”), and the addition of data in memory [5]. To complicate the job of detecting honeypots, Rowe *et al.* [2] have proposed the use of “fake honeypots”, which are normal systems that have been disguised to appear as honeypots. This is a form of deception in which objects are *camouflaged* or *masked* to appear as some other object [6]. Once fake honeypots are implemented, attackers will avoid compromising them, thinking that the systems are actually honeypots. As this defensive technique becomes common knowledge, attackers

must expend additional resources to determine whether a system is a true honeypot or not.

In this work, we perform a game theoretical investigation of deception in network security. The scenario we examine is as follows. A defender deploys honeypots in her network. She can choose to employ camouflage which means either disguising normal systems as honeypots or honeypots as normal systems. The attacker observes the system without being able to detect its real type. If the system is camouflaged, she observes the disguise; otherwise, she observes the system’s actual type. Then, the attacker must determine whether or not to proceed compromising the system. We model the defender-attacker interaction as a *signaling game*, a dynamic game of incomplete information. A *dynamic game* is a game in which players take turns choosing their actions. In the scenario under study, the defender moves first by choosing whether or not to disguise the system, after which, the attacker decides whether to compromise the system. The *incomplete information* arises from the attacker’s uncertainty of system type. We determine and characterize the *perfect Bayesian equilibria* of the game. At an equilibrium, the defender and the attacker do not have incentives to unilaterally deviate by changing their strategies. We show that camouflage is an equilibrium strategy for the defender. Finally, we discuss the benefits of these deceptive equilibrium actions in defending a network.

Related work. Recently, honeypots have become one of the main tools used to collect information about attacks. A honeypot is a system that is specifically setup to trap unauthorized accesses [3]. Unlike a normal system, a honeypot has rich logging facilities that record all activities and permit detailed analysis. Since its sole purpose is engaging hackers, any activity directed at it is by definition unauthorized. HoneyNet is a volunteer project that has deployed honeynets, networks of honeypots, around the world. One goal of the HoneyNet project is to collect data from their honeynets and analyze and correlate activities to achieve early warning and prediction [7].

Besides honeypots, deception techniques have been proposed for defending information systems. Cohen [8] provides a comprehensive discussion of deception as a means to protect information systems. Among his many conclusions is that de-




Fig. 1. A signaling game.

ception has a positive effect for the defenders, and conversely, a negative effect for the attackers. Cohen and Koike [1] showed how deception can control the path of an attack. Rowe *et al.* [2] showed how fake honeypots, normal systems disguised as honeypots, decreased the amount of attacks a network witnesses. Several tools for evaluating honeypot deceptions were proposed in [9].

Game theory has been used for studying various security related problems. Grossklags *et al.* [10] provided a game-theoretic analysis of security investment decision-making. Garg and Grosu [11] develop a model for deception in honeynets using a dynamic game in extensive form. Patcha and Park [12] used a signaling game to study intrusion detection in ad hoc wireless networks. To the best of our knowledge, this is the first work that models deception in computer network security as a signaling game.

II. SIGNALING GAMES

The analysis performed in this paper is based on *game theory*, a subfield of economics that focuses on analyzing the interactions among decision makers. In our setting we have two decision makers, the defender and the attacker. The defender employs deception which masks the type of her systems (*i.e.*, normal system or honeypot). Due to *incomplete information*, the attacker is uncertain upon initial inspection if the system she is attacking is a normal system that is beneficial to compromise, or is a honeypot that is harmful.

The defender and the attacker play a *dynamic game*: the defender plays first followed by the attacker. Most importantly, the attacker observes some information about the defender's action which she uses to optimize her choice of action. Dynamic games are usually specified in *extensive form*. This form represents the game as a tree. Each *decision node* represents a state in which a player must choose an action. Each leaf gives a payoff for a particular sequence of choices. We model the interaction between the defender and attacker as a *signaling game* [13], a non-cooperative two player dynamic game of incomplete information. An extensive form representation of a basic signaling game is given in Fig. 1.

A game of incomplete information can be transformed into

a game of imperfect information by adding a hypothetical player called Nature (denoted by C here) and conditioning the payoffs on the Nature's unknown moves. The Nature player moves first by randomly choosing the type of Player 1 from an *a priori* probability distribution over all Player 1's types. This distribution is known by all players. In our example, Nature assigns type T with probability $P \in [0, 1]$ and type B with probability $1 - P$. The *type* is private to Player 1 and cannot be directly observed by Player 2. Once Player 1 learns her type, she decides what signal or message to send to Player 2. The *signal* provides indirect information to Player 2 about the type of Player 1. In our example, Player 1 can send either signal t or b . Player 2 observes the signal and then responds with an action, either L or R .

The set of decision nodes is partitioned into information sets. An *information set* is a set of one or more decision nodes of a player that determines the possible moves conditioned on what the player has observed so far. Decision nodes belonging to the same information set are indicated by a dotted line connecting them. Player 1 has two information sets, one where Nature assigns T and the other where Nature assigns B . Player 2 also has two information sets, one where she receives signal t and the other where she receives signal b .

The game in our example has eight *outcomes*. One example of outcome is Nature assigning T to Player 1, Player 1 sending b , and Player 2 responding with L . Each outcome results in a payoff. A *payoff* measures the benefit that each player earns if that outcome is reached. Payoffs corresponding to outcome i are represented as tuples (a_i, b_i) , the first component being Player 1's payoff and the second being Player 2's payoff.

A *strategy* is a plan that accounts for all contingencies and consists of one action for each information set. Continuing with the above example, one of Player 1's strategies is to send t if she is of type B , and to send b if she is of type T . For Player 2, one strategy is to respond with L if she receives t , and to respond with R if she receives b . Since Player 1 has two information sets and two signals, she has $2^2 = 4$ strategies. Similarly, Player 2 has $2^2 = 4$ strategies. The players are self-interested and welfare-maximizing, thus, they choose strategies that maximize their payoff. A *strategy profile* consists of a tuple of strategies, one strategy for each player. When investigating non-cooperative games, we are interested in strategy profiles that results in equilibria.

Definition 1 (Nash equilibrium): A *Nash equilibrium* is a strategy profile in which each player cannot improve her payoff by unilaterally changing her strategies.

A Nash equilibrium results in a steady state in which each player chooses not to deviate as doing so reduces her payoff. It was proved by Nash that a game with a finite set of players, each having a finite set of actions, has at least one equilibrium [14]. In a signaling game, an information set can be on the equilibrium path or off the equilibrium path. An information set is *on the equilibrium path* if it is reached with positive probability when the equilibrium strategies are played, and it is *off the equilibrium path* if it is reached with zero probability. Certain Nash equilibria result

in unlikely plays for off the equilibrium path information sets. The *perfect Bayesian equilibrium* (PBE) concept refines the solution concept by excluding equilibria with suboptimal play on the off equilibrium paths. Besides strategy profiles, PBE requires that players have beliefs. A player has *beliefs* about which decision node within an information is reached. Beliefs are represented as a probability distribution over the set of decision nodes within an information set. Gibbons [13] defines perfect Bayesian equilibrium as strategy profiles and beliefs that satisfy the following four requirements.

Requirement 1: At each information set, the player with the move must have a *belief* about which decision node in the information set has been reached.

Requirement 2: The action taken at an information set by the player with the move must be optimal given the player's belief at the information set and the player's subsequent moves. This is referred to as *sequential rationality*.

Requirement 3: At information sets on the equilibrium path, beliefs are determined by Bayes' law and player's equilibrium strategies.

Requirement 4: At information sets off the equilibrium path, beliefs are determined by Bayes' law and the players' equilibrium strategies where *possible*.

Definition 2 (Perfect Bayesian equilibrium): A *perfect Bayesian equilibrium* consists of a strategy profile and beliefs that satisfy Requirements 1 to 4.

For signaling games, a PBE where the sender gives the same signal regardless of type is called a *pooling equilibrium*. Pooling equilibria require that Player 2 have two sets of beliefs, one for the on equilibrium path and one for the off equilibrium path. The on equilibrium path beliefs are given by probability distribution P and the off equilibrium path beliefs are given by either $p \in [0, 1]$ or $q \in [0, 1]$, depending on which path is off the equilibrium. A *separating equilibrium* is a PBE in which the sender gives an unique signal for each type. Player 2 has one set of beliefs for these equilibria which is given by P .

In the next sections, we examine the defender-attacker interaction by modeling it as a signaling game. Then, we quantify the perfect Bayesian equilibria of this game and draw conclusions on defender's strategies.

III. THE DECEPTION GAME

A defender protects her network by deploying honeypots, traps to detect unauthorized access. To improve efficacy, she camouflages her systems. She can disguise normal systems as honeypots and honeypots as normal systems. Rowe *et al.* [2] showed that this technique effectively enhances computer network security. After the network is created, an attacker then attempts to compromise systems. The attacker can successfully compromise normal systems, but not honeypots. If the attacker attempts to compromise a honeypot, the defender observes the actions and can later improve her defenses. We model this interaction between defender and attacker as the signaling game presented in Fig. 2. The notation used in this paper is summarized in Tab. I.

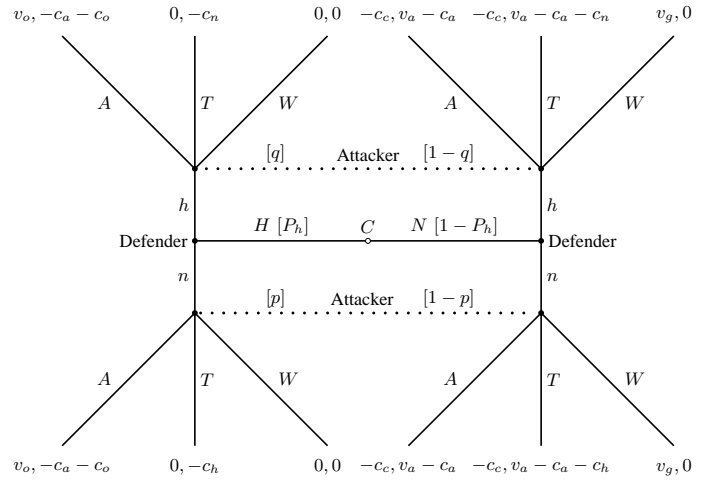


Fig. 2. The deception signaling game.

TABLE I
DECEPTION GAME NOTATION.

Notation	Description
Defender	
H	System is a honeypot
N	System is normal (<i>i.e.</i> , not a honeypot)
h	Signal that the system is a honeypot
n	Signal that the system is normal
v_o	Benefit of observing an attack on a honeypot ($v_o > 0$)
v_g	Benefit of avoiding an attack on a normal system ($v_g \geq 0$)
c_c	Cost due to a compromised normal system ($c_c > 0$)
Attacker	
A	Attack without determining the system type
T	Condition the attack on determining the system type
W	Do not attempt an attack
P_h	Belief that the system is a honeypot on the equilibrium path ($P_h \in [0, 1]$)
q, p	Belief that the system is a honeypot on the off-equilibrium path ($p, q \in [0, 1]$)
c_a	Cost of compromising a system ($c_a \geq 0$)
c_n	Cost of testing if a system is normal ($c_n \in [0, c_a]$)
c_h	Cost of testing if a system is a honeypot ($c_h \in [0, c_a]$)
c_o	Cost due to being observed ($c_o > 0$)
v_a	Benefit of compromising a normal system ($v_a \geq c_a$)

The game begins when the attacker attempts to compromise a system on the defender's network. Nature (C) chooses the type of the system that the defender will protect. Nature chooses either type honeypot (H) or normal system (N) with probability P_h and $1 - P_h$, respectively. We interpret Nature assigning the type as the attacker randomly selecting a system within the network to compromise. The defender can signal that the system is a honeypot (h) or a normal system (n), independent of the system's actual type. A "fake honeypot" [2] is a system of type N for which the defender signals h . The attacker receives the signal and then chooses an action. The attacker will either attack the system without determining the system type (A), condition the attack on determining the system type (T), or retreat (W).

Figure 2 shows that there are twelve potential outcomes. We

TABLE II
THE PLAYERS' PURE STRATEGIES.

		Defender			
type	s_1	s_2	s_3	s_4	
N	n	n	h	h	
H	n	h	n	h	

		Attacker								
signal	t_1	t_2	t_3	t_4	t_5	t_6	t_7	t_8	t_9	
n	A	A	A	T	T	T	W	W	W	
h	A	T	W	A	T	W	A	T	W	

assume that the costs of setting up the systems, incorporating camouflage, and the management of the systems are fixed costs and therefore are not considered in the defender's decision making process. We further assume that normal systems generate revenue for the defender and that honeypots do not. In this work, we exclude the possibility of a firm running honeypots for profit. If one of the normal systems is compromised, the defender incurs a loss of c_c . This includes both the costs due to loss of business and the cost of restoring the system. If the attacker attempts to compromise a honeypot, the defender observes these actions, learning how to improve her defenses. In this case, the defender gains v_o , the benefit from observing the attacker. If the attacker retreats before attacking a normal system, the defender gains v_g . The attacker incurs cost c_a when attacking a system, irrespective of success. If the attack succeeds, she gains v_a , leading to a "profit" of $v_a - c_a$. The attacker loses c_o when attacking honeypots. The attacker has tests to determine if a system is normal or not. The tests to determine if a system is normal or a honeypot costs c_n and c_h , respectively. After receiving the results of tests, the attacker either moves forward with her attack or abandons the attempt. If she tests for a honeypot, then she incurs either c_n or c_h , both of which are less than $c_a + c_o$. If she tests for a normal system, her payoffs are either $v_a - c_a - c_n$ or $v_a - c_a - c_h$.

The strategies of each player are given in Tab. II. The defender has four pure (non-deterministic) strategies. Each strategy has a contingency for each of its types. Strategy s_1 and s_4 is to signal normal system (n) and honeypot (h), respectively, regardless of the actual system type. In s_2 the defender signals the actual type of the system, while in s_3 she signals the opposite of the actual type, *e.g.*, the defender signals honeypot if the system is normal. The attacker has nine pure strategies, each strategy giving a plan of action for every signal that she can receive. In the case of strategy t_4 , the attacker plays T if she receives signal n and plays A if she receives h .

We now examine the deception game for (pure) perfect Bayesian equilibrium (PBE). We first examine the existence of separating equilibria. This reduces to examining the possible equilibria involving defender's strategies s_2 and s_3 . We examine if any equilibria contain strategy s_2 (reveal the true system type). When faced with this strategy, the attacker chooses A if she receives n as the signal, since the payoff obtained by playing A , $v_a - c_a$, is greater than the payoff obtained by

playing T , $v_a - c_a - c_h$, and the one obtained by playing W , 0. When faced with h , she responds with W as the resultant payoff of zero is greater than A 's payoff of $-c_a - c_o$ and T 's payoff of $-c_n$. The attacker plays A and W upon receiving n and h , respectively. This is strategy t_3 . The defender's best response to t_3 is s_4 : by sending h instead of n , the defender increases her payoff from $-c_c$ to v_g . Examining the defender's strategy s_3 , the attacker's best response is t_7 : she selects action W when she receives n and A when h . In response to t_7 , the defender switches to strategy s_2 as, again, her payoff for normal systems increases from $-c_c$ to v_g . Consequently, this game does not have any separating equilibria as neither s_2 nor s_3 result in a steady state.

We now investigate the existence of pooling equilibria. Potential equilibria should involve either strategy s_1 or s_4 . If the defender involves strategy s_1 , the attacker receives signal n and will choose A if the action results in an expected payoff greater than the expected payoff from her other two actions. The attacker's expected payoff for playing A must be greater than or equal to the expected payoff for playing T ,

$$P_h(-c_a - c_o) + (1 - P_h)(v_a - c_a) \geq P_h(-c_h) + (1 - P_h)(v_a - c_a - c_h)$$

which gives,

$$P_h \leq \frac{c_h}{c_a + c_o}. \quad (1)$$

The attacker's expected payoff for playing A must also be greater than or equal to the expected payoff for playing W ,

$$P_h(-c_a - c_o) + (1 - P_h)(v_a - c_a) \geq 0$$

which gives,

$$P_h \leq \frac{v_a - c_a}{v_a + c_o}. \quad (2)$$

We now need to determine the beliefs and actions for the off-equilibrium path of sending signal h . First, strategy t_3 does not lead to a steady state as the defender's best response is s_3 , not s_1 . The defender changes her response as signal h results in a payoff of v_g which is greater than the payoff of $-c_c$ when she sends n . The attacker responds with t_1 if she has a set of beliefs for the off-equilibrium path, q , that gives the maximum expected payoff when playing A . Thus, the expected payoff obtained by playing A should be greater than or equal to the expected payoff obtained by playing T . Therefore, q must satisfy,

$$q(-c_a - c_o) + (1 - q)(v_a - c_a) \geq q(-c_n) + (1 - q)(v_a - c_a - c_n)$$

which gives,

$$q \leq \frac{c_n}{c_a + c_o}. \quad (3)$$

The expected payoff obtained by playing A should also be greater than or equal to the expected payoff obtained by playing W . Therefore q must also satisfy,

$$q(-c_a - c_o) + (1 - q)(v_a - c_a) \geq 0$$

which gives,

$$q \leq \frac{v_a - c_a}{v_a + c_o} \quad (4)$$

Similarly, the attacker responds with t_2 if she has a set of beliefs that gives the maximum expected payoff when playing T . The attacker's expected payoff obtained by playing T is greater than or equal to the payoff obtained by playing A if

$$\begin{aligned} q(-c_n) + (1-q)(v_a - c_a - c_n) \\ \geq q(-c_a - c_o) + (1-q)(v_a - c_a) \end{aligned}$$

which is equivalent to,

$$q \geq \frac{c_n}{c_a + c_o}. \quad (5)$$

The attacker's expected payoff obtained by playing T is greater than or equal to the payoff obtained by playing W if

$$q(-c_n) + (1-q)(v_a - c_a - c_n) \geq 0$$

which is equivalent to,

$$q \leq 1 - \frac{c_n}{v_a - c_a}. \quad (6)$$

The attacker responds with T if the expected payoff obtained by playing T is greater than or equal to the one obtained by playing A ,

$$\begin{aligned} P_h(-c_h) + (1 - P_h)(v_a - c_a - c_h) \\ \geq P_h(-c_a - c_o) + (1 - P_h)(v_a - c_a) \end{aligned}$$

which gives,

$$P_h \geq \frac{c_h}{c_a + c_o}, \quad (7)$$

and if the expected payoff obtained by playing T is greater than or equal to the one obtained by playing W ,

$$P_h(-c_h) + (1 - P_h)(v_a - c_a - c_h) \geq 0$$

which gives,

$$P_h \leq 1 - \frac{c_h}{v_a - c_a}. \quad (8)$$

Neither t_4 nor t_6 lead to a steady state. If the attacker plays t_4 , the defender switches to s_2 as her payoff increases from 0 to v_g . If the attacker plays t_6 , the defender's best response is s_3 . Strategy t_5 leads to equilibrium if the off-equilibrium path beliefs satisfy conditions (5) and (6). The attacker responds with W when she receives signal n , if the expected payoff obtained by playing W is at least as large as the expected payoff obtained by playing A

$$0 \geq P_h(-c_a - c_o) + (1 - P_h)(v_a - c_a)$$

which gives,

$$P_h \leq \frac{v_a - c_a}{v_a + c_o}. \quad (9)$$

The payoff obtained by playing W should also be at least as large as the expected payoff obtained by playing T ,

$$0 \geq P_h(-c_h) + (1 - P_h)(v_a - c_a - c_h)$$

which gives,

$$P_h \geq 1 - \frac{c_h}{v_a - c_a}. \quad (10)$$

Strategies s_1 and t_8 and beliefs that satisfy (3) and (4) results in a PBE. The attacker plays W on the off-equilibrium path (strategy t_9) if she has a set of beliefs in which the expected payoff obtained by playing W is at least as large as the one obtained by playing A ,

$$0 \geq q(-c_a - c_o) + (1-q)(v_a - c_a)$$

$$q \geq \frac{v_a - c_a}{v_a + c_o}, \quad (11)$$

and at least as large as the payoff obtained by playing T ,

$$0 \geq q(-c_n) + (1-q)(v_a - c_a - c_n)$$

$$q \geq 1 - \frac{c_n}{v_a - c_a}. \quad (12)$$

If A is played on the off-equilibrium path, the defender switches to strategy s_2 . Thus, t_7 is not part of a steady state.

We performed a similar analysis for strategy s_4 , but because of the space restrictions we will not present it here. We summarize the potential equilibria and their conditions in Tab. III. Some of the equilibria have identical on equilibrium paths. A unique path is needed to have a unique outcome. Six outcomes are given by the ten equilibria: the on equilibrium path of E_1 and E_2 are identical, as well as the path of E_4 and E_5 , E_6 and E_7 , and E_9 and E_{10} .

IV. CASE STUDIES AND DISCUSSION

The question now is how the defender can best use the above analysis. Normally multiple equilibria pose problems as there does not exist any coordination to designate an equilibrium. But since the defender moves first, she dictates the resulting equilibrium by selecting her strategy (the "first mover advantage"). The attacker observes the signal and then chooses her best response.

The defender chooses her strategy as follows. She determines which equilibria exist by evaluating the conditions for E_1 - E_{10} . She then computes the attacker's payoff for each equilibria. Then for each strategy s_1 and s_4 , the defender selects the equilibria that maximizes the attacker's payoff, discarding the rest. At most two outcomes will be represented. Finally, the defender chooses the equilibrium strategy that maximizes her payoff.

In the following, we investigate two case studies to illustrate the above decision process.

Case Study 1. Assume a defender has 10 percent ($P_h = 0.10$) of the machines in her network setup as honeypots. Suppose the defender believes that the attack costs are $c_a = 3.00$, $c_h = 0.25$, $c_n = 0.50$, and $c_o = 0.10$ and that the attacker values a compromised system at $v_a = 5.00$. We begin by evaluating the equilibria conditions set forth in Tab. III. Only equilibria E_3 , E_6 , and E_7 are possible as the others have conditions

TABLE III
EQUILIBRIA AND THEIR CONDITIONS.

	PBE	conditions	
		equilibrium	off-equilibrium
E_1	(s_1, t_1, q)	$P_h \leq \frac{c_h}{c_a+c_o}$ $P_h \leq \frac{v_a-c_a}{v_a+c_o}$	$q \leq \frac{c_n}{c_a+c_o}$ $q \leq \frac{v_a-c_a}{v_a+c_o}$
E_2	(s_1, t_2, q)	$P_h \leq \frac{c_h}{c_a+c_o}$ $P_h \leq \frac{v_a-c_a}{v_a+c_o}$	$q \geq \frac{c_n}{c_a+c_o}$ $q \leq 1 - \frac{c_n}{v_a-c_a}$
E_3	(s_1, t_5, q)	$P_h \geq \frac{c_h}{c_a+c_o}$ $P_h \leq 1 - \frac{c_h}{v_a-c_a}$	$q \geq \frac{c_n}{c_a+c_o}$ $q \leq 1 - \frac{c_n}{v_a-c_a}$
E_4	(s_1, t_8, q)	$P_h \geq \frac{v_a-c_a}{v_a+c_o}$ $P_h \geq 1 - \frac{c_h}{v_a-c_a}$	$q \geq \frac{c_n}{c_a+c_o}$ $q \leq 1 - \frac{c_n}{v_a-c_a}$
E_5	(s_1, t_9, q)	$P_h \geq \frac{v_a-c_a}{v_a+c_o}$ $P_h \geq 1 - \frac{c_h}{v_a-c_a}$	$q \geq \frac{v_a-c_a}{v_a+c_o}$ $q \geq 1 - \frac{c_n}{v_a-c_a}$

	PBE	conditions	
		equilibrium	off-equilibrium
E_6	(s_4, t_1, p)	$P_h \leq \frac{c_n}{c_a+c_o}$ $P_h \leq \frac{v_a-c_a}{v_a+c_o}$	$p \leq \frac{c_h}{c_a+c_o}$ $p \leq \frac{v_a-c_a}{v_a+c_o}$
E_7	(s_4, t_4, p)	$P_h \leq \frac{c_n}{c_a+c_o}$ $P_h \leq \frac{v_a-c_a}{v_a+c_o}$	$p \geq \frac{c_h}{c_a+c_o}$ $p \leq 1 - \frac{c_h}{v_a-c_a}$
E_8	(s_4, t_5, p)	$P_h \geq \frac{c_n}{c_a+c_o}$ $P_h \leq 1 - \frac{c_n}{v_a-c_a}$	$p \geq \frac{c_h}{c_a+c_o}$ $p \leq 1 - \frac{c_h}{v_a-c_a}$
E_9	(s_4, t_6, p)	$P_h \geq \frac{v_a-c_a}{v_a+c_o}$ $P_h \geq 1 - \frac{c_n}{v_a-c_a}$	$p \geq \frac{c_h}{c_a+c_o}$ $p \leq 1 - \frac{c_h}{v_a-c_a}$
E_{10}	(s_4, t_9, p)	$P_h \geq \frac{v_a-c_a}{v_a+c_o}$ $P_h \geq 1 - \frac{c_n}{v_a-c_a}$	$p \geq \frac{v_a-c_a}{v_a+c_o}$ $p \geq 1 - \frac{c_h}{v_a-c_a}$

that require $P_h \neq 0.10$, which is not the case for the system we consider. Equilibrium E_3 involves strategy s_1 ; E_6 an E_7 involve strategy s_4 . For each of the defender's strategies, we compute the attacker's maximum expected payoff. Equilibria E_6 and E_7 have an identical equilibrium path, resulting in the same outcome. Thus, the attacker is indifferent between her equilibrium strategies in E_6 and E_7 . Strategy s_1 earns the defender $-0.90c_c$, while s_4 earns $0.10v_o - 0.90c_c$. Therefore, the defender will choose to implement strategy s_4 , a deceptive strategy in which all normal systems need to be camouflaged as honeypots.

Case Study 2. Most attacks are heavily automated. Networks of compromised systems ("botnets") provide inexpensive processing for highly parallel automated attacks. This essentially reduces the cost associated with testing and attacks to zero, and therefore, $c_a = c_h = c_n = 0$. With this in mind, the defender evaluates the equilibria conditions. Equilibria E_1 , E_2 , E_6 , and E_7 occur only when $P_h = 0$; equilibria E_4 , E_5 , E_9 , and E_{10} occur only when $P_h = 1$. Supposing that the defender has honeypots and normal systems within her network, these equilibria are not possible. Equilibria E_3 and E_8 are possible and both result in the defender losing $(1 - P_h)c_c$. Consequently, the defender will choose to implement either strategy s_1 , have all honeypot systems disguised as normal, or s_4 , have all normal systems disguised as honeypots.

V. CONCLUSION

In this paper, we investigated the use of deception in the interaction between a defender of a network and an attacker. We modeled the interaction between the defender and the attacker as a signaling game. We examined the game and determined and characterize its potential equilibria. We then show how the defender can use this analysis to better protect her network using two case studies. For future work, we plan to examine hybrid strategies, which are non-deterministic choices of actions. In the game presented in this paper, the defender chooses a network in which either all normal systems are camouflaged as honeypots, or all honeypots are camouflaged as normal systems depending on the costs involved. A hybrid

strategy will allow the defender to have a network in which a mix of all types of systems coexists (honeypots, camouflaged honeypots, normal systems, and camouflaged normal systems) allowing a much richer set of equilibria, leading to more effective deception strategies.

ACKNOWLEDGMENTS

This research was supported, in part, by NSF grant DGE-00654014.

REFERENCES

- [1] F. Cohen and D. Koike, "Misleading attackers with deception," in *Proc. of the 5th IEEE SMC Information Assurance Workshop*, 2004, pp. 30–37.
- [2] N. C. Rowe, E. John Custy, and B. T. Duong, "Defending cyberspace with fake honeypots," *J. Comput.*, vol. 2, no. 2, pp. 25–36, 2007.
- [3] The Honeynet Project, *Know Your Enemy: Learning about Security Threats*. Addison-Wesley Professional, 2004.
- [4] B. McCarty, "The honeynet arms race," *IEEE Security Privacy*, vol. 1, no. 6, pp. 79–82, 2003.
- [5] X. Fu, W. Yu, D. Cheng, X. Tan, K. Streff, and S. Graham, "On recognizing virtual honeypots and countermeasures," in *Proc. 2nd IEEE Int. Symp. on Dependable, Autonomic and Secure Computing (DASC '06)*, 2006, pp. 211–218.
- [6] J. Bowyer Bell and B. Whaley, *Cheating and Deception*. Transaction Publishers, 1991.
- [7] L. Spitzner, "The honeynet project: Trapping the hackers," *IEEE Security Privacy*, vol. 1, no. 2, pp. 15–23, 2003.
- [8] F. Cohen, "A note on the role of deception in information protection," *Computers and Security*, vol. 17, no. 6, pp. 483–506, 1998.
- [9] N. C. Rowe, "Measuring the effectiveness of honeypot counter-counterdeception," in *Proc. of the 39th Annual Hawaii Int. Conf. on System Sciences (HICSS '06)*, 2006, p. 129.3.
- [10] J. Grossklags, N. Christin, and J. Chuang, "Secure or insure? a game-theoretic analysis of information security games," in *Proc. of the 17th Intl. World Wide Web Conference*, April 2008.
- [11] N. Garg and D. Grosu, "Deception in honeynets: a game-theoretic analysis," in *Proc. of the 8th IEEE Information Assurance Workshop (IAW '07)*, 2007, pp. 107–113.
- [12] A. Patcha and J.-M. Park, "A game theoretic approach to modeling intrusion detection in mobile ad hoc networks," in *Proc. of the 5th IEEE SMC Information Assurance Workshop*, 2004, pp. 280–284.
- [13] R. Gibbons, *Game Theory for Applied Economists*. Princeton, NJ, USA: Princeton University Press, 1992.
- [14] J. F. Nash, "Equilibrium points in n -person games," *Proc. Nat'l Academy of Sciences of the United States of Am.*, vol. 36, no. 1, pp. 48–49, Jan. 1950.