

Game Theory for Adaptive Defensive Cyber Deception

Kimberly Ferguson-Walter*
Department of Defense
Washington D.C., USA

Justin Mauger*
SPAWAR System Center Pacific
San Diego, CA, USA

Sunny Fugate*
SPAWAR System Center Pacific
San Diego, CA, USA

Maxine Major
SPAWAR System Center Pacific
San Diego, CA, USA

ACM Reference Format:

Kimberly Ferguson-Walter, Sunny Fugate, Justin Mauger, and Maxine Major. 2019. Game Theory for Adaptive Defensive Cyber Deception. In *Hot Topics in the Science of Security Symposium (HotSoS)*, April 1–3, 2019, Nashville, TN, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3314058.3314063>

1 INTRODUCTION

1.1 Deception for Cyber Defense

As infamous hacker Kevin Mitnick describes in his book *The Art of Deception*, “the human factor is truly security’s weakest link” [18]. Deception has been widely successful when used by hackers for social engineering and by military strategists in kinetic warfare [26]. Deception affects the human’s beliefs, decisions, and behaviors. Similarly, as cyber defenders, deception is a powerful tool that should be employed to protect our systems against humans who wish to penetrate, attack, and harm them.

The cyber defender’s role is notoriously unfair since a defender aims to prevent intrusions at every possible location, and the attacker only needs to discover and exploit a single vulnerability in order to breach defenses. Similar to moving target defenses [14], the use of deception for cyber defense provides the promise of re-balancing this asymmetric disadvantage.

While many techniques have been developed to increase the speed and accuracy of detecting adversarial activity with the aim of making a defender’s job easier, beyond a priori hardening of systems, less research has been done on techniques to make the attacker’s job fundamentally more difficult. Moving target defenses help make an attacker’s job harder by adding unpredictability to the attack space by quickly changing information. Deception can add more uncertainty by including misinformation and masking true information. This further impacts the decision-making of attackers, causing them to waste both time and effort. Moreover, cyber deception can be used by a defender to impart an incorrect belief in the attacker. This incorrect belief can cause ripple effects into every stage of the cyber kill chain [12] and can interrupt multiple attacks over a long time period.

*These authors contributed equally to the paper.

1.1.1 Adaptive Cyber Deception. Advanced cyber defenses cannot solely rely on perimeter defenses and must be able to respond to attacker activity in cyber time—at the same speed as network traffic and cyber attacks. This requires intelligent defensive systems that can automatically react to malicious behavior and evolve over time as attacks change. The artificial intelligence method used for the defensive system must be able to look ahead and *dynamically* consider how an attacker might behave in the future before taking a defensive action. The concept of adaptive or active cyber defense [5]—where a system automatically prepares and implements predictive defensive strategies or reacts to detected suspicious activity without human intervention, is gaining acceptance, but has not yet been widely put into practice. Cyber deception is also an emerging research area in cyber defense [11, 22]. Adaptive defensive cyber deception combines these two concepts to strategically present misinformation which automatically changes as it observes changes in the network or attacker behaviors. Adaptive cyber deception is an altogether new, but inevitable extension of prior work, which cuts across the computer security, behavioral science, and artificial intelligence communities.

There are many reasons why cyber deception techniques should be adaptive. For example, surprise is one important element that can affect the attacker’s decision processes and actions. When an attacker experiences unexpected results, they may decide to change strategies or retry the same techniques, either of which will disrupt or delay their progress, giving defenders more time and opportunity to react appropriately. Static cyber deception techniques may cause surprise at first, but over time their effect will wear off, as the attackers become familiar with these techniques and learn to expect them. If the techniques are adaptive, they will detect when the attacker has developed a response to the deception, and will alter the method of deception accordingly. Surprise is only one example of how adaptive cyber deception can negatively impact an attacker and disrupt their progress. There are many more ways to affect an attacker which we are currently investigating, such as causing frustration, confusion, and self-doubt. These can cause an attacker to increase the number of errors they perform and make them easier to detect, delay their attack until further defenses are in place or a critical task is complete, and even deter an attacker from pursuing a particular target. The scenario presented in Section 4 of this paper assumes a defender goal of delaying an attacker.

1.2 Decoy Systems for Cyber Deception

There are a variety of cyber deception techniques discussed in cyber security research, including honeypots and honey-tokens, replay attacks, packet crafting and altered payloads, tar-pitting, false documents, decoy systems, and others. Due to their simplicity, both

conceptually and in regards to implementation, this work focuses on the use of decoys for adaptive deception. An increasing variety of decoy systems and approaches have become commercially available by companies including: [1, 8, 13, 25]. To our knowledge, these systems are not yet *adaptive* as defined above, but rather static, preconfigured defenses which we hope will evolve into the dynamic defenses that we foresee as the necessary future of defensive deception. Although related works primarily focus on the use of honeypots, we subscribe to the distinction between honeypot and decoy provided in [6], and will use decoy systems as a working example throughout this paper.

A decoy environment consists of *realistic*, *lightweight* virtual systems that appear to be real systems running real services from the perspective of an attacker scanning the network. These are deployed on a real network alongside real systems in order to maximize the chance of an attacker being detected and mitigated quickly. The large number of false assets helps provide an asymmetric advantage for cyber defenders by reducing the chance of a real asset being attacked, as well as distracting an attacker from real assets and content. This forces an attacker to take additional actions, thus slowing them down and increasing the likelihood of revealing themselves. This cyber deception can be taken even further, leading an attacker towards a specific incorrect belief.

1.2.1 An Adaptive Decoy System. The 2015 Gartner Report on Deception Techniques included the following key finding “Deception as an automated responsive mechanism represents a sea change in the capabilities of the future of IT security that product managers or security programs should not take lightly” [20]. However, adaptive cyber defense systems are still in their infancy, and cyber deception is just a small piece in the cyber defense landscape. We observe both a need to focus on adaptive cyber deception systems and a gap in current research, and thereby propose using game theory to pursue autonomous cyber deception systems which can decide when, where, and how to best use deception based on attacker behavior.

Pilot studies performed by [6] using red teamers as human subjects suggest that decoy systems can be highly effective at disrupting network reconnaissance, confusing an attacker by using their cognitive bias against them, and causing self-doubt, which then increases the attacker’s cognitive load. We claim that these effects can be multiplied by allowing the decoys to be adaptive to each adversary’s specific strategies and preferences. Furthermore, these initial pilot studies indicate that cyber deception may be as or more effective when the attacker is actually informed that there is deception being used on the network for defensive purposes. While rigorous studies addressing this question have been completed [7], the final analysis of the data is still forthcoming.

Implementing an adaptive cyber deception strategy in a real-world cyber environment necessitates capabilities that may not be deployed in a typical network. In particular, it requires sensors, actuators, and a means of logically connecting inputs to outputs, making decisions as to how and when to adapt.

(1) **Sensors** collect information to detect behavioral-based adversarial activity such as detecting scanning activity and logon attempts. More advanced sensors could detect activity such as the attacker attempting to use stolen passwords and could extend to post-exploitation activities, particularly where decoys contain honey-tokens.

(2) **Actuators** take an automated action on the network or host as directed. Actuation of decoys involve configuration changes, creating new decoys, changing decoy parameters, modifying service banners, and other deceptive activities. Further decoy adaptations could include changing the IP address, opening or closing ports, adding or removing services, or even spoofing a different operating system. Not only are these specialized tasks not normally managed by modern enterprise network management tools, but these tasks must be automated in order to rapidly respond to suspicious activity.

Furthermore, these same cyber deception techniques can be used to do more than delay, confuse and surprise an attacker. Cyber deception can be used to influence the attack in more direct ways. For example, the defender may want to learn something specific about an attacker or collect information about a specific type of attack. Deception can be used to entice or convince an attacker to take an action that, unknown to the attacker, actually benefits the defender in some way. This is important for cyber defenders, since as we move forward into more adaptive cyber defensive systems, we must consider the natural co-evolution of multi-step, multi-stage attack/defense situations. These advanced defenses must take a strategic view, where moves are considered many steps ahead of both attacker and defender actions; this is called cyber co-evolution [27].

A co-evolution strategy supports advanced defender goals such as *preference elicitation* (where the main goal is to gather information about the attacker, his preferences, and goals) and *topological misinformation* (where the main goal is to provide the attacker with a specific incorrect belief in the network topology). Strategies using cyber deception techniques to support these goals will be addressed in future work.

2 BACKGROUND

Game theory studies decision-making problems amongst a group of players, and is applicable to situations where two or more players have conflicting goals. It provides a quantitative framework for reasoning about decisions given scenarios where the players are either unaware or uncertain about the intent of opposing players. Therefore, game theory can provide insight into when and how strategies should be adopted by a cyber defender, or in our case an adaptive cyber defense system using automated deception techniques.

2.1 Game Theory Terminology

Briefly, a game consists of players, actions, payoffs, and strategies. In sequential games, players alternate turns, choosing from a set of available actions at each point. We assume games are finite in the sense that all action sequences end after a fixed amount of moves. A strategy is a complete description by a player of what actions to take at all possible decision points. Given a set of strategies, one for each player, there is a utility function assigning a numerical value to each player as a payoff for the outcome of everyone following their chosen strategy. A traditional analysis of games is finding equilibrium strategies. The most commonly calculated are *Nash equilibria*, in which players have no incentive to unilaterally deviate from their strategy, given the other players’ equilibrium strategy.

The following are terms used to describe game theory concepts presented in this paper.

- (1) **Perfect information:** All players know the previous actions taken by other players.
- (2) **Imperfect information:** There exists at least one player for whom other players' moves are partially hidden.
- (3) **Complete information:** Strategies and payoffs are known by all players.
- (4) **Incomplete information:** There exists at least one player who does not know all strategies and payoffs of the other players.
- (5) **Bayesian game:** A game of incomplete information can be converted to a game of complete but imperfect information in which some players have different *types* defined by their set of available actions, strategies, and payoffs. The other players maintain beliefs about these types, updating them as the game progresses according to Bayes' rule.
- (6) **Bayesian Equilibrium:** A version of the Nash equilibrium for Bayesian games. A Perfect Bayesian Equilibrium (PBE) is a further refinement.
- (7) **Hypergame:** A complex game in which at least one player has a misperception about the model of the game being played. Players may a) be unaware that they are playing the game, and b) be unaware of the possible moves in the game.
- (8) **Zero-sum game:** A game in which the total gains and losses for all players sum to zero. In a two player game, one player's payoff results in the opposing player receiving an equivalent negative payoff.

2.2 Previous Work

In this section we examine three game-theoretical implementations of deception in network security. Research in the field considers sequential Bayesian games in both the one-shot case where the game ends after a single iteration, and the repeated case in which players keep alternating turns.

The primary model developed in each of these works is that of a defender deploying honeypots to detect an attacker and obtain information on the attacker's intentions. The defender can disguise normal systems as honeypots and honeypots as normal systems. The attacker observes a system without being able to detect its real type and is uncertain whether to attempt to compromise the system. Similarly, the defender may be uncertain about how to interpret the actions of the attacker.

In [3] a one-shot scenario is considered in which the defender moves first by choosing whether or not to disguise a system, after which the attacker decides whether to compromise the system. They determine and characterize the Perfect Bayesian Equilibria (PBE) for this game. The authors conclude that camouflage is an equilibrium strategy for the defender and that these deceptive equilibrium actions are beneficial in defending a network. The paper includes two case studies exemplifying their approach and sets the way for further research.

In a similar approach, the paper by [4] applies these techniques to mitigate Denial of Service (DoS) attacks on a computer network by deploying honeypots as a means to attract the attacker and retrieve information about his real intentions. The authors observe that defense against DoS attacks turns out to be an optimization problem from the defender's point of view, where the defender is allocating limited

resources to minimize cost while maximizing deterrence. They then proceed to model this problem using signaling as a dynamic game with incomplete information. Solving for the PBE suggests a cost effective mitigation of DoS attacks through deception.

An extension of these concepts from the one-shot version to repeated scenarios that also include false information is explored in [17]. Here the application area is a honeypot-enabled network for the Internet of Things. Among their results for a repeated game, the Bayesian belief update scheme was shown to converge. The proof of their results was complemented by numerical simulations verifying their analyses. This paper presents many directions for the analysis of deception in games focused on network security.

3 GAME THEORY FOR ADAPTIVE CYBER DECEPTION

Cyber deception game models require additional complexity in that they are best modeled as hybrid games of both *imperfect* and *incomplete* information. The games themselves are non-cooperative and non-symmetric – defenders and attackers usually having very different strategies available. The goals of defenders and attackers are often in opposition and as such, many games can be structured as zero-sum games. If the payoff values of a particular strategy are not easily comparable to the payoffs for alternative strategies, then we propose that these strategies should be placed in different game trees (see Figure 1) and analyzed independently, or within the context of a hypergame.

As described by [16], deception and misperception games are well-suited for representation as hypergames. From a hypergame perspective we can naturally and directly represent the interrelationship between defender goals, observations, subgames, and individual strategies. For this purpose, we define *game contexts* as the differing perspectives of the game by the players: *adversary context* and *defender context*. In Section 3.1 we will introduce our hypergame cyber deception model, and in Section 4 we will describe a simple scenario illustrating how the defender's overall goal of delaying an attacker provides a context for determining useful estimates of strategy payoffs.

Additionally, in our model we will define an attacker as being "naive" or "sophisticated" according to whether they are aware that deception may be a component of the game and strategies in play. For purposes of simplicity, the illustrative scenario in this paper assumes a naive adversary who is unwitting of deception.

3.1 Cyber Deception as Hypergames

Hypergames can provide a solution for modeling conflict where misperception or intentional deception exists between players [16]. While hypergame theory has been discussed for human deception and cyber defense, no formal notation has been standardized to model cyber deception.

Hypergames are well-suited to model cyber deception, particularly given the broad set of potential defensive deception goals, strategies, and implementations. As a cyber defender, although fully in control of the game-board, we cannot know all the possible actions that the attacker will make. This is even more true for the attacker, who may not even know a game is being actively played, and even if they were made aware of the certainty of deception, would not

know what types of deceptive moves were available to the defender. In a cyber deception game, the defender's game tree may look very different from that of the attacker (see Figure 1), and the hypergame model can encompass all of the sub-game trees as they are played out for each individual player's perception of the game.

In our proposed hypergame model, game tree differs depending on the defender's goal. This has implications for the hypergame as well as for individual sub-games. In particular, payoff values may be significantly different in the context of different goals. For example, payoffs for an *obfuscation* goal are related to the likelihood of an attacker successfully finding real systems or information and may be purely based on how such information is hidden. The probability of an attacker randomly interacting with a realistic-looking decoy is $\frac{\text{decoy}}{\text{decoy} + \text{real}}$. Making decoys look more enticing than real systems will increase this probability and the chance of receiving the associated payoff.

However, the same sub-game may have different payoffs for different goals. If instead of hiding real systems the goal of the hypergame is to delay attackers while some critical activity is completed on the system, then payoffs may be based on the amount of time an attacker spends interacting with individual decoys. While the set of strategies may be identical, the differences in the game's goal can result in different payoffs capable of informing the selection of wildly different strategies.

3.2 Formal Definition of Cyber Deception Games

We first give a short definition of a regular game, then introduce our concept of a cyber deception game. In the following, all sets are finite.

Definition 3.1. A finite, sequential game $G = (\mathcal{P}, \mathcal{M}, \Theta, u, T)$ consists of the following:

- (1) A set of n players \mathcal{P} , traditionally written as a set of integers $[n] = \{1, 2, \dots, n\}$.
- (2) A collection $\mathcal{M} = \{\mathcal{M}^i\}$ of sets of moves/actions each player can take. Not all moves in \mathcal{M}^i are available to player i at all times.
- (3) Players take turns in sequence, and each sequence of moves is bounded in length by T . For $t \leq T$, a sequence of moves $m = (m_1^{i_1}, \dots, m_t^{i_t})$, where i_j references the player who made a move at time j , and $1 \leq j \leq t$, is called a history. By convention, player 1 moves first.
- (4) A collection $\Theta = \{\Theta^i\}$ of sets of strategies for each player, where a strategy is a complete description of moves to take in all contingencies.
- (5) A *strategy profile* is a tuple $\theta = (\theta_1, \dots, \theta_n)$ of strategies, one from each Θ^i . Each strategy profile results in an *outcome*, which is the game played out according to θ .
- (6) A set $u = \{u^i\}$ of utility functions for each player. The utility $u^i(\theta)$ is a numerical score representing the payoff to player i of the outcome of θ .

Note that player designators are superscripted in this definition. This detour from standard notation is necessary as we introduce the concept of player perceptions of other players' moves in Definition 3.3.

The game G can also be described in graphical form as a tree. While player moves and strategies may be repeated to achieve player goals, the finiteness condition ensures the tree eventually stops.

Definition 3.2. A game tree (G, V, E) is a representation of G as a directed acyclic graph with nodes V and edges E , loosely defined in the following way:

- (1) Internal nodes are decision points for the player whose turn it is. The root node decision belongs to player 1.
- (2) At each node belonging to player i , there is an outgoing edge for each possible move in \mathcal{M}^i .
- (3) Each move history defines a path through the game tree.
- (4) A strategy in Θ^i is given by a choice of outgoing edge from each node belonging to player i .
- (5) The outcome of a strategy profile is a unique terminal node of a path from the root node.
- (6) Each outcome associated with a terminal node provides a payoff to each associated player.

Note that in Figure 1 the decision node is shown as a small black circle. The boxes at the end of the outgoing edge from that decision represent states, and are labeled with the player move leading to that state.

We now define cyber deception games as an extension of regular games that allows us to formulate deception in the framework of hypergame theory. Specifically, we introduce the concept of a player's *perception* of the game. In Definition 3.3 we focus on two-player games with an attacker A and a defender D . For simplicity, we index by $\{A, D\}$ instead of integers.

Definition 3.3. Let $G = (\mathcal{P}, \mathcal{M}, \Theta, u, T)$ be a regular game as previously defined, with player set $\mathcal{P} = \{A, D\}$. A *cyber deception game* is a triple (G, G^A, G^D) , where G^A and G^D are *derived* games defined in the following manner.

- (1) In G , let $m^A = (m_1^A, m_2^A, \dots, m_r^A)$ be a move history in \mathcal{M}^A taken by A , and $m^D = (m_1^D, m_2^D, \dots, m_s^D)$ a sequence of moves in \mathcal{M}^D taken by D , with $r + s \leq T$. We allow different move length sequences, and do not require that moves be made in alternating fashion. We define a special move ϵ to be a null move. If desired, one can add the appropriate number of ϵ 's to equalize the sequence lengths. In what follows, we shall use the convention that $m = (m^A, m^D)$ refers to the interleaved sequence of moves.
- (2) For any two players $X, Y \in \{A, D\} = \mathcal{P}$ (with replacement), we define $m^{X|Y}$ as player Y 's perception of the sequence of moves m^X taken by player X . We use conditional probability notation to suggest this can be read as "beliefs about player X 's moves given that player Y holds these beliefs". Each element of $m^{X|Y}$ is in the set $\mathcal{M}^{X|Y}$, which is defined as Y 's perception of the set of X 's available moves \mathcal{M}^X . Note that it is not necessarily true that X 's perception of their own moves is correct, so $m^{X|X} = m^X$ is not necessarily true.
- (3) Still using our convention, let $m^{*|A} = (m^{A|A}, m^{D|A})$ be player A 's beliefs about the move sequence $m = (m^A, m^D)$.
- (4) In analogy with moves, $u^{*|A} = (u^{A|A}, u^{D|A})$ is A 's perception of utility $u = (u^A, u^D)$. The same notation is used for strategies, i.e. $\Theta^{*|A} = (\Theta^{A|A}, \Theta^{D|A})$.

- (5) The *derived* game $G^A = (\mathcal{P}, \mathcal{M}^{*|A}, \Theta^{*|A}, u^{*|A}, T)$ is A 's *perception* of G . The terms G^D , $\mathcal{M}^{*|D}$, $\Theta^{*|D}$, and $u^{*|D}$ are similarly defined.

3.3 Key Characteristics of Cyber Deception Games

We define several concepts which differentiate our model of adaptive cyber deception from traditional game theory models. In particular, we differentiate between a player's perception of possible moves, outcomes, and utilities and the true parameters of the game.

Similarly, we define *perceived moves* as an individual player's perception of the full move sequence. For example, player A perceives the full move sequence to be $m^{*|A} = (m^{A|A}, m^{D|A})$. If this were to differ in any way from the actual move sequence $m = (m^A, m^D)$, the game G^A would be one of imperfect information. In the scenario discussed in Section 4 we assume an *omniscient defender* as this best exemplifies the potential advantage the defender can gain through leveraging cyber deception. The omniscient defender has perfect and complete information. We will adopt this simplification as a first step in analysis of cyber deception.

We define *perceived utility* as player X 's utility function, $u^{X|X}$, which may differ from the true utility u^X (where X can be either player). In addition, they may also misperceive the other player's utility, i.e. $u^{Y|X} \neq u^Y$, thereby making the game G^X one of incomplete information. This is particularly true for player A when the perceived value of the targeted system is manipulated by player D .

In general, games of cyber deception tend to rely on manipulation of game payoffs. From a defender's perspective, the true payoffs are less relevant than the attacker's perceived payoffs in relation to the true payoffs. This attacker perception of system value is the opportunity that we are manipulating through deception techniques.

A key conceptual problem in former analyses of cyber deception games is the lack of fully addressing that perceived payoffs themselves can be affected by the deception strategies selected by defenders. Defenders seek to cause adversaries to believe that decoy systems have a high payoff and that real systems have a low payoff. In many scenarios, manipulation of perceived payoffs is the key parameter in optimizing defender advantage. By formally describing and modeling the manipulation of attacker estimations of payoffs and by controlling the observability of defender strategies we can better capture the dynamics with which cyber deception games provide increased defender advantage.

Finally, we define *adversary context* as the derived game G^A , which is the attacker's view of the game G , according to their perception of moves $\mathcal{M}^{*|A}$, utilities $u^{*|A}$, and strategies $\Theta^{*|A}$. Similarly, G^D is the defender's view of the game G in the *defender context*, according to their perception of move sequences $\mathcal{M}^{*|D}$, utilities $u^{*|D}$, and strategies $\Theta^{*|D}$.

4 EXAMPLE SCENARIO

As an illustrative example, consider a highly simplified scenario in which the defender has pre-deployed decoys on the network, and the attacker has just initiated a port scan of a single system. They believe the system is a database server containing possibly valuable information, and would like to gain access. However, the system is actually a decoy and holds no valuable information. In our model of

cyber deception, there are three game trees: one for the true game G , and one for each perceived game G^A and G^D . Figure 1 depicts this scenario by showing the adversary context tree on the left and the defender context tree on the right (since the defender is omniscient, $G^D \equiv G$). These trees are not exactly in one-to-one correspondence, but there exist mappings between them. We will describe how a specific move sequence is represented in each context tree. Payoffs are notional and zero sum.

In each tree, the root node belongs to the attacker, and the first move consists of the port scan. The first layer of outgoing edges represent the defender's possible moves in reaction to the scan. The next layer of edges represent the attacker's possible subsequent moves. Recall that ϵ is meant to represent a null move, similar to a player not doing anything. In deception games, we also use ϵ to denote two additional types of moves. One type is when the attacker believes that the defender did not detect him and thus took no action. Another is when the defender takes a move, but the attacker does not detect it. In either case, the move is effectively equivalent to `Do nothing`, as denoted in Figure 1.

As seen in attacker context tree on the left side, the attacker believes the defender's possible first moves are $\mathcal{M}^{D|A} = \{\epsilon, \text{Block IP}\}$. Moreover, the attacker believes that if they are able to reach the terminal node labeled e , their final payoff will be high: $u^{A|A} = 10$. A greedy strategy is to wait for the defender's reaction, then make the move `Attempt login` if possible.

The defender context tree on the right side depicts the true situation. In addition to the moves in $\mathcal{M}^{D|A}$, the defender's possible moves \mathcal{M}^D include `{Disable login, Launch new decoy}`, which are unknown to and undetectable by the attacker. The dotted curves between node a and nodes b, c, d represent an information set illustrating that these are all equivalent to the result of an ϵ move in the attacker's eyes. Let us say the defender decides to modify the database service running on the decoy to disallow all logins. The attacker then attempts to login, and fails. The double line indicates the true move sequence $m^{*|D} = (m_1^{A|D} = \text{Port scan of decoy}, m_2^{D|D} = \text{Disable login}, m_3^{A|D} = \text{Attempt login (failure)})$.

$m^{*|A} = (m_1^{A|A} = \text{Port scan of decoy}, m = (m_2^{D|D} = \text{Disable login}, m_3^{A|D} = \text{Attempt login (failure)}))$.

Returning to the left tree, the attacker interprets the lack of visible response as indicating the defender did not detect or did not react to the port scan, which is equivalent to $m_2^{D|A} = \epsilon$. They then proceed to attempt to login. The double lines indicates the perceived full move sequence $m^{*|A} = (m_1^{A|A} = \text{Port scan}, m_2^{D|A} = \epsilon, m_3^{A|A} = \text{Attempt login})$.

The attacker's true payoff for the outcome as illustrated in game G^D as terminal node f is $u^{A|D} = -20$, benefiting the defender, as the resulting outcome is a failed attempt to log into a decoy machine. The reason for the low payoff is that targeting the decoy wastes the attacker's time and effort and keeps them away from important information (at least temporarily), so provides a negative final payoff for the attacker and a positive payoff for the defender. The net result is the attacker wastes a chance to log on to a real server, the defender is alerted to the attacker's location and as to what login credentials were stolen and is given a chance to modify and protect the real

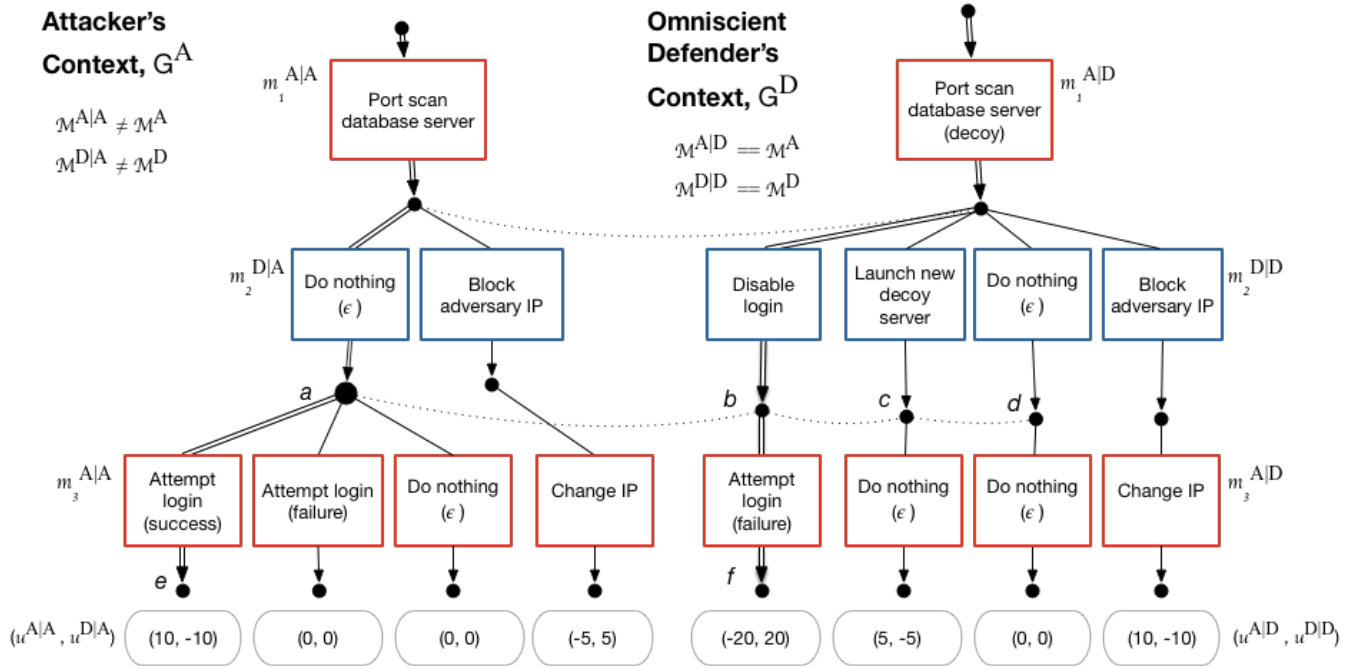


Figure 1: The left tree shows the adversary context (the game tree for G^A). The double lines indicate the attacker's perception $m^{*|A}$ of the move sequence leading to their optimal perceived payoff pair $u^{*|A} = (10, -10)$. The right tree shows the defender's context (the game tree for G^D). The double lines indicate the defender's perception $m^{*|D}$ of the move sequence leading to their optimal perceived payoff pair $u^{*|D} = (-20, 20)$. The dotted horizontal curves represent the correspondence between the players' information sets.

server. The defender was able to manipulate the game board and the attacker's perceived payoffs, causing the attacker to select a strategy that did not produce the highest possible payoff, despite the attacker's attempt to use a greedy strategy.

4.1 Online Learning for Adaptive Cyber Deception

Since cyber games are massively iterated, online learning is a necessary component of an adaptive cyber deception system. The defender must attempt to infer the attacker's beliefs over time and apply them to its future decision making. As the defender observes the attacker interacting with the network through the information collected by its sensors, the defender will need to use this information to model the state of the attacker and estimate the attacker perceived payoffs. With knowledge of the current game tree and estimation of attacker perceptions, the defender can now dynamically manipulate the game board to change the payoffs associated with the next possible actions. This is an iterative process where the defender must continue to learn about the attacker through observation and update the models accordingly. The decisions made and actions taken by the defender both manipulates the payoffs the attacker can receive, and limits the strategies available to the attacker at the next time step. The hypergame model presented in this paper lends itself well to an online learning solution.

5 DISCUSSION

In this work, we demonstrate a straightforward extension of prior game theory models of cyber defense through the use of individual player models of the game environment where the game structure and payoffs may be manipulated by another player. Based on hypergame theory we introduce a formal notation to describe cyber deception scenarios which are both imperfect and incomplete. Our work combines aspects of cyber security research, cyber deception techniques, and game theory. As described in Section 2.2, prior research concerning game theoretic models for deception in network security have primarily been concerned with defender manipulation of the veracity of signals sent to an attacker. In these prior models the defender's primary manipulation is deciding which machines in the network should be fake and which should be real and whether or not to send true or false signals regarding whether a system is real or not. This approach assumes the attacker is fully cognizant of the nature of the deception and true parameters of the game environment for both players.

In our work, we consider the setup where the defender becomes aware of the presence of the attacker through their interaction with a decoy and is therefore able to manipulate the true payoffs and game structure using cyber deception technology. Unlike previous models such as [3, 9, 19], we consider the setup where the attacker is unaware of proactive and reactive defensive deception moves taken by the defender. Considerations of the differences between naive and informed attackers has been studied in work such as [23] in which an attacker's strategy set depends on their knowledge of a

defender's strategy for masking system and service identities. We expand on prior approaches by modeling an extensive form game using independent game trees for attacker and defender in order to convert a complex game of incomplete information into two independent games where both players believe they have perfect and complete information, but where only the omniscient defender knows the true structure and payoffs for the game.

Prior models also generally only address the initial interaction of an attacker and deceptive defender. By assuming a naive attacker, our model allows defenders to make several choices relating to a desired attack scenario and for game parameter manipulations by the defender. While the naive attacker will rationally optimize their choices and make moves according to their own game model, they will falsely assume perfect and complete information. The deceiving defender is then able to perform additional optimization steps within the defender's tree and optimize defender advantage by minimizing attacker payoffs.

The notation for the hypergame model presented in this paper provides a framework to quantify how cyber deception can be used to influence players' perceptions of available moves and potential payoffs in a game with active misinformation. The scenario illustrated here is limited to depicting attacker and defender perceptions of available moves and potential payoffs in a misinformation game, but could be expanded to include online learning to fully implement an adaptive cyber deception solution.

6 FUTURE WORK

In future work we intend to develop a richer model of player behaviors and payoffs, including the development of a learning model for attacker behaviors and utility. In particular, prior work on attack trees has applied game theory to analyzing the behaviors of attackers and defenders [2, 15, 28]. Traditional attack trees are modeled according to the goals (or initial actions) of the attacker with the leaves of the tree containing one or more defensive countermeasures. Future models of our hypergame concepts may investigate an alternative form of attack trees rooted in defender goals rather than attacker actions. Given a defender game context and goal, such an approach can provide a set of priors to guide initial strategy selection for defense. In this way, we believe that defensive deception is an enabling assumption for improving defender advantage in cyber security scenarios. In essence, deception provides defenders with the freedom of maneuver currently only enjoyed by cyber attackers.

The cyber deception framework presented in this paper can apply to more complicated versions of cyber deception games and we plan to explore the potential of this model more thoroughly in future work. Examples include scenarios where:

- the defender does not have perfect knowledge
- there are multiple attackers either working cooperatively or unaware of each other
- there are resource allocation and costs associated with various player actions
- the attacker is aware that there is deception but unaware of the details
- the attacker is using counter-deception

Practical applications of game theory deployed for physical security at airports and ship ports have shown success [21, 24], and

we hope to achieve similar realistic demonstrations in cyber security. We are currently creating a practical implementation of our adaptive cyber deception techniques using the Rainbow autonomies framework [10]. This framework is agnostic to the problem domain, but allows for a modular implementation to connect to sensors and actuators in decoy systems (or any system). This allows us to implement our game theoretic models of conflict, make decisions based on information collected by sensors, implement strategy execution through automated actuations, and eventually validate both our current model and future improvements in online learning and adaptation. Finally, while we have chosen a cyber security domain as the context for our research, the general structure of our model has practical applications in many other domains. If our model proves useful for cyber security it will likely bear fruit in non-cyber domains where adversarial scenarios can benefit from defensive deception.

REFERENCES

- [1] Attivo Networks. viewed October 2017. Deception-Based Threat Detection and Continuous Response Platform. <https://attivonetworks.com/product/deception-technology/>.
- [2] Stefano Bistarelli, Marco Dall'Aglia, and Pamela Peretti. 2006. Strategic Games on Defense Trees. In *Formal Aspects in Security and Trust*. Springer, Berlin, Heidelberg, Berlin, Heidelberg, 1–15.
- [3] Thomas E Carroll and Daniel Grosu. 2009. A Game Theoretic Investigation of Deception in Network Security. In *2009 Proceedings of 18th International Conference on Computer Communications and Networks - ICCCN 2009*. IEEE, 1–6.
- [4] Hayreddin Çeker, Jun Zhuang, Shambhu Upadhyaya, Quang Duy La, and Boon-Hee Soong. 2016. Deception-Based Game Theoretical Approach to Mitigate DoS Attacks. In *Decision and Game Theory for Security*. Springer International Publishing, Cham, 18–38.
- [5] Dorothy E Denning. 2014. Framework and principles for active cyber defense. *Computers and Security* 40 (2014), 108–113.
- [6] Kimberly J Ferguson-Walter, Dana S LaFon, and Temmie B Shade. 2017. Friend or Faux: Deception for Cyber Defense. *Journal Of Information Warfare* 16, 2 (Aug. 2017), 28–42.
- [7] Kimberly J Ferguson-Walter, Temmie B Shade, Andrew V Rogers, Elizabeth M Niedbala, Michael C Trumbo, Kevin Nauer, Kristin M Divis, Aaron P Jones, Angela Combs, and Robert G Abbott. 2019. The Tularosa Study: An Experimental Design and Implementation to Quantify the Effectiveness of Cyber Deception. In *Proceedings of the 32nd Hawaii International Conference on System Sciences (HICSS)*. 10.
- [8] Galois Inc. viewed October 2017. CyberChaff™: Confounding and Detecting Adversaries. <https://galois.com/project/cyberchaff/>.
- [9] Nandan Garg and Daniel Grosu. 2007. Deception in Honeynets: A Game-Theoretic Analysis. In *2007 IEEE SMC Information Assurance and Security Workshop*. IEEE, 107–113.
- [10] David Garlan, S-W Cheng, A-C Huang, Bradley Schmerl, and Peter Steenkiste. 2004. Rainbow: Architecture-based self-adaptation with reusable infrastructure. *Computer* 37, 10 (2004), 46–54.
- [11] Kristin E Heckman, Frank J Stech, Roshan K Thomas, Ben Schmoker, and Alexander W Tsow. 2015. *Cyber Denial, Deception and Counter Deception*. Springer International Publishing, Cham.
- [12] Eric M. Hutchins, Michael J. Cloppert, and Rohan M. Amin. 2011. Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains. *Leading Issues in Information Warfare & Security Research* 1, 1 (2011), 80.
- [13] Illusive Networks. viewed October 2017. Deception Management System. <https://www.illusivenetworks.com/solutions>.
- [14] Sushil Jajodia, Anup K. Ghosh, Vipin Swarup, Cliff Wang, and X. Sean Wang. 2011. *Moving Target Defense: Creating Asymmetric Uncertainty for Cyber Threats* (1st ed.). Springer Publishing Company, Incorporated.
- [15] Christopher Kiekintveld, Viliam Lisý, and Radek Pibil. 2015. Game-Theoretic Foundations for the Strategic Use of Honeypots in Network Security. In *Cyber Warfare*. Springer, Cham, Cham, 81–101.
- [16] Nicholas S Kovach, Alan S Gibson, and Gary B Lamont. 2015. Hypergame Theory: A Model for Conflict, Misperception, and Deception. *Game Theory* 2015, 2 (Aug. 2015), 1–20.
- [17] Q. D. La, T. Q. S. Quek, J. Lee, S. Jin, and H. Zhu. 2016. Deceptive Attack and Defense Game in Honeypot-Enabled Networks for the Internet of Things. *IEEE Internet of Things Journal* 3, 6 (Dec 2016), 1025–1035.

- [18] Kevin Mitnick and William Simon. 2003. *The Art of Deception*. John Wiley & Sons, Inc., New York, NY.
- [19] Radek Píbil, Viliam Lisý, Christopher Kiekintveld, Branislav Bošanský, and Michal Pěchouček. 2012. Game Theoretic Model of Strategic Honeypot Selection in Computer Networks. In *Decision and Game Theory for Security*. Springer, Berlin, Heidelberg, Berlin, Heidelberg, 201–220.
- [20] Lawrence Pingree. 2015. Emerging Technology Analysis: Deception Techniques and Technologies Create Security Technology Business Opportunities. *Gartner, Inc* (2015).
- [21] James Pita, Manish Jain, Janusz Marecki, Fernando Ordóñez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, and Sarit Kraus. 2008. Deployed ARMOR Protection: The Application of a Game Theoretic Model for Security at the Los Angeles International Airport. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems: Industrial Track (AAMAS '08)*. 125–132.
- [22] Neil C Rowe and Julian Rrushi. 2016. *Introduction to Cyberdeception*. Springer International Publishing, Cham.
- [23] Aaron Schlenker, Omkar Thakoor, Haifeng Xu, Fei Fang, Milind Tambe, Long Tran-Thanh, Phebe Vayanos, and Yevgeniy Vorobeychik. 2018. Deceiving cyber adversaries: A game theoretic approach. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 892–900.
- [24] Eric Shieh, Bo An, Rong Yang, Milind Tambe, Craig Baldwin, Joseph DiRenzo, Ben Maule, and Garrett Meyer. 2012. PROTECT: A Deployed Game Theoretic System to Protect the Ports of the United States. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS '12)*. 13–20.
- [25] TrapX Security. viewed October 2017. Deception in Depth – The Architecture of Choice. <https://trapx.com/product/>.
- [26] Barton Whaley. 1969. *Stratagem: deception and surprise in war*. Artech House, Cambridge, Massachusetts.
- [27] Gerald Willard. 2014. Understanding the Co-Evolution of Cyber Defenses and Attacks to Achieve Enhanced Cybersecurity. *Journal Of Information Warfare* 14, 2 (April 2014).
- [28] Saman Zonouz, Himanshu Khurana, William Sanders, and Timothy Yardley. 2014. RRE: A Game-Theoretic Intrusion Response and Recovery Engine. *IEEE Transactions on Parallel and Distributed Systems* 25 (Feb. 2014), 395–406. Issue 2.