SPAWAR®

Systems Center
PACIFIC

TECHNICAL REPORT 3141
July 2018

Game Theory for Adaptive Defensive Cyber Deception

Kimberly Ferguson-Walter
U.S. Department of Defense

Sunny Fugate
Justin Mauger
Maxine Major
Space and Naval Warfare Systems Center Pacific

Approved for public release.

SSC Pacific
San Diego, CA 92152-5001

M. K. Yokoyma, CAPT, USN
Commanding Officer

W. R. Bonwit
Executive Director

ADMINISTRATIVE INFORMATION

The work described in this report was performed for NSA Research by the Information Operations and Cybersecurity Divisions of the Cyber/Science & Technology Department, Space and Naval Warfare Systems Center Pacific (SSC Pacific), San Diego, CA.

Released by
Gregory Settelmayer, Head
Information Operations Division

Under authority of
Stephen Russell, Head
Cyber/Science & Technology Department

EXECUTIVE SUMMARY

OBJECTIVE

In prior game theory work, cyber deception games have generally been modeled as non-cooperative, sequential games, where the quality and authenticity of the signal is being manipulated. This work relies on a signal, which can be deceptive, being actively pushed to the attacker. We have formulated a new framework for cyber deception games in which each player has their own perception of the game being played and the moves being taken. A particular player may manipulate other players' perceived payoffs to induce them to take sub-optimal actions. This model of deception seeks to address defender interactions with an attacker following the initial deployment of cyber defenses addressed by previous work. Our primary contribution is a model of defensive cyber deception that incorporates defender control of attacker perception of the cyber environment. Our ultimate goal it to inform future cyber defense systems, enabling more sophisticated responses to attacker behaviors and improving defensive posture.

CONTENTS

Figures

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” [16]. Deception has been widely successful when used by hackers for social engineering and by military strategists in kinetic warfare [23]. 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 [12], 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 and can interrupt multiple attacks over a long time period.

1.1.1 Adaptive Cyber Deception

Advanced cyber defenses need to 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 controlling 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 [10, 20]. 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 an 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.

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. Within the last few years a variety of decoy systems and approaches have become commercially available by companies including: [1, 7, 11, 22]. 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. We will use decoy systems as a working example throughout this paper. Though not critical to this paper, we subscribe to the distinction between honeypot and decoy provided in [6].

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” [18]. 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, 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 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, but more formal studies are needed.

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 control system to connect inputs to outputs, making decisions as to how and when to adapt. Sensors collect information to detect behavioral-based adversarial activity such as detecting scanning activity and logon attempts. Actuators take an automated action on the network or host as directed by a control system. Actuation of decoys can involve configuration changes, creating new decoys, changing decoy parameters, modifying service banners, and other deceptive activities. Such tasks must be automated in order to rapidly respond to suspicious activity. However, these specialized tasks are not normally performed by modern enterprise network management tools.

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 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 have a strategic view, where moves are considered many steps ahead of both attacker and defender actions. This is called cyber co-evolution [24].

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 automated cyber defense system using deception techniques.

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 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.

In our model and analysis we use the well known definitions of perfect information, complete information, Bayesian games, information sets, and Bayesian Equilibrium.

2.1 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.

[3] considers the one-shot scenario 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 recent 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 [15]. 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 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 [14], 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" which we use to define a high-level hypergame described in the next section. In our hypergame cyber deception model, the defender's overall goal of interference provides a context for determining useful estimates of strategy payoffs.

Additionally, in our model We 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 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 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 $[n] = \{1, 2, \cdots, n\}$.*
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. By convention, player 1 moves first. For $t \leqslant T$, a sequence of moves $m = \left(m_{i_1}^1, \cdots, m_{i_t}^t\right)$, where $m_{i_j}^j \in \mathcal{M}_{i_j}$, is called a history.*
4. *A collection $\Theta = \{\Theta_i\}$ of sets of strategies for each player. As previously indicated, a strategy is a complete description of moves to take in all contingencies. A strategy profile is a tuple $\theta = (\theta_1, \cdots, \theta_n)$ of strategies, one from each Θ_i.*
5. *Each strategy profile results in an outcome, which is the game played out according to θ.*
6. *A tuple $u = (u_1, \cdots, u_n)$ 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.*

The game G can also be described in graphical form as a tree. The finiteness condition ensures the tree eventually stops.

Definition 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 labeled by the player whose turn it is.*

2. *At each node belonging to player i, there is an outgoing edge for each possible move in $\mathcal{M}_i$. The root node belongs to player 1.*

3. *Each move history defines a path to a unique node.*

4. *A strategy in Θ_i is given by a choice of outgoing edge from each node belonging to player i.*

5. *An outcome is a path from the root node to a unique terminal node.*

6. *Each terminal node is labeled with the payoff to each player of that outcome.*

We now define cyber deception games as an extension of regular games that allows us to formulate deception in the framework of game theory. Specifically, we introduce the concept of a player's *perception* of the game he is participating in. We focus on two-player games with an attacker A and a defender D, with A moving first. We index by $\{A, D\}$ instead of integers, and we use superscripts for readability.

Definition 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 $m^A = (m_1^A, m_2^A, \cdots, m_r^A)$ be a move history in $\mathcal{M}^A$ taken by A, and $m^D = (m_1^D, m_2^D, \cdots, m_s^D)$ a sequence of moves in $\mathcal{M}^D$ taken by D, with $r + s \leqslant T$. We allow different length sequences, and do not require that moves be made in alternating fashion. The actual sequence of moves in the order in which they occurred will be some interleaving of m^A and m^D, in other words, a shuffle. 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 $X, Y \in \mathcal{P}$ (possibly identical), 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 $m^{X|X} = m^X$.*

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 *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.*

3.2 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. 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

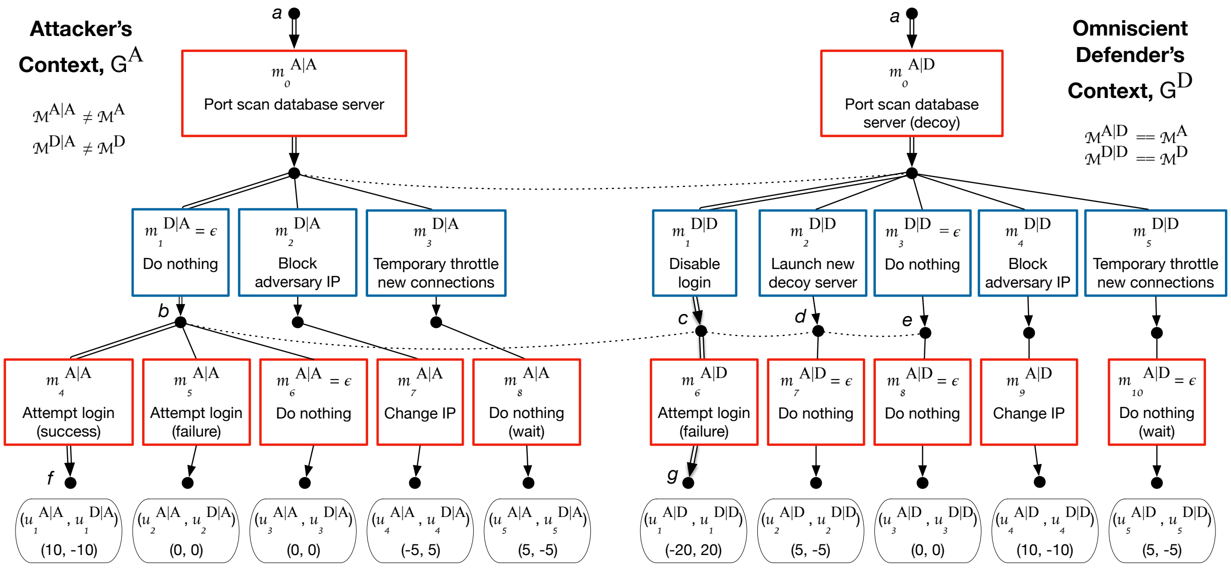

Figure 1. The left tree shows the attacker's context (the game tree for game G^A). The double lines indicate the attacker's perception $m^{*|A}$ of the move sequence leading to their optimal perceived payoff pair $u_1^{*|A} = (u_1^{A|A}, u_1^{D|A})$. The right tree shows the defender's context (the game tree for game G^D). The double lines indicate the defender's perception $m^{*|D}$ of the move sequence leading to their optimal perceived payoff pair $u_1^{*|D} = (u_1^{A|D}, u_1^{D|D})$. Since we assume an omniscient defender, they control (or know) the actual payoffs and moves, i.e. $u^{*|D} = (u^A, u^D)$ and $m^{*|D} = (m^A, m^D)$. The dotted horizontal curves represent the correspondence between the players' information sets.

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 attacked system is manipulated by player D.

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 our current model we assume an *omniscient defender* as this exemplifies the control that deception enables for defender advantage. The omniscient defender has perfect and complete information. We will adopt this simplification as a first step in analysis of cyber deception.

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's 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 scenario in which the defender has pre-deployed decoys on the network, and the attacker has just initiated a port scan of a particular system. They believe the system is a database server containing possibly valuable information, and would like to break in to it. 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 particular 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 top layer of outgoing edges represent the defender's possible moves in reaction to the scan. The lower 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 essentially equivalent to `Do nothing`.

As seen in attacker context tree on the left side, the attacker believes the defender's possible first moves are $\mathcal{M}^{D|A} = \{\epsilon, \texttt{Block IP, Throttle}\}$. They believe that if they are able to reach the terminal node labeled f, their final payoff will be high: $u_1^{A|A} = 10$. 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 actual situation. In addition to the moves in $\mathcal{M}^{D|A}$, the defender's possible moves $\mathcal{M}^D$ include $\{\texttt{Disable login, Launch new decoy}\}$, which are undetectable by the attacker. The dotted curves between node a and nodes b, c, d are meant to illustrate 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 $m = (m_1^{D|D} = \texttt{Disable login}, m_6^{A|D} = \texttt{Attempt login (failure)})$.

Returning to the left tree, the attacker interprets the lack of visible response as indicating the defender did not detect the port scan, which is equivalent to $m_1^{D|A} = \epsilon$. They then proceed to attempt login. The double lines indicates the perceived full move sequence $m^{*|A} = (\epsilon, m_4^{A|A} = \texttt{Attempt login})$.

The attacker's payoff for the outcome in game G^D is $u_1^{A|D} = -20$, benefiting the defender. The state is `failed login`. 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. The net result is the attacker wastes a chance to logon to a real server, the defender is alerted as to what login credentials were stolen and is given a chance to modify and protect the real server. The defender was able to manipulate the game board and the attacker's *perceived* payoffs, causing the attacker to select a branch in the tree that did not produce the highest possible payoff.

5. CONCLUSIONS

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. Our work combines aspects of cyber security research, network-based deception techniques, and game theory. As described in Section 2.1, 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 decoy. 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 identity of the attacker through their interaction with a decoy and is therefore able to manipulate the true payoffs and game structure using deception. Unlike previous models such as [3, 8, 17], we consider the setup where the attacker is unaware of pro-active and reactive defensive (deceptive) moves taken by the defender. We use 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 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 falsely assume perfect and complete information. The deceiving defender is then able perform additional optimization steps within the defender's tree and optimize defender advantage and minimize attacker payoffs.

In most realistic scenarios, neither player knows each other's strategy sets nor do they have an accurate notion of each other's payoffs. However, games of deception are unique in that the defender has significant control over the game environment and can be assumed to be able to hide their true strategies and payoffs from the adversary while simultaneously being able to perfectly observe attacker behaviors, and even farther, to manipulate payoffs and the set of strategies available to the attacker. In essence, we assume deception gives defenders ultimate control of the game environment with the exception of being certain if the adversary is aware or unaware of deception. However a player can choose to disclose (or signal) that deception is in play, should that be deemed beneficial.

6. FUTURE WORK

Our current work uses illustrative examples of attacker and defender utility and game structure. In future work we intend to involve 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]. However, this has been done primarily for purposes of optimizing resources used to patch vulnerabilities to secure networks against future attacks [13]. 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 defender countermeasures. In our work, we consider trees rooted with defender goals. We intend to propose an alternative form of attack trees which are 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 defensive maneuver. 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.

We assert that the cyber deception framework presented in this paper can apply to more complicated versions of cyber deception games and plan to explore the potential of this model more thoroughly in future work. Examples include: when the defender does not have perfect knowledge; when there are multiple attackers either working cooperatively or unaware of each other; when there are resource allocation and costs associated with various player actions; when the attacker is aware that there is deception but unaware of the details; and when the attacker is using counter-deception.

Practical applications of game theory deployed for physical security at airports and ship ports have shown success [19, 21]. 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 autonomics framework [9]. This framework is agnostic to the problem domain, but allows for easy implementation of shims 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, we feel that 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. Deception-based threat detection and continuous response platform. https://attivonetworks.com/product/deception-technology/, viewed October 2017.

2. Stefano Bistarelli, Marco Dall'Aglio, and Pamela Peretti. Strategic Games on Defense Trees. In *Formal Aspects in Security and Trust*, pages 1–15. Springer, Berlin, Heidelberg, Berlin, Heidelberg, August 2006.

3. Thomas E Carroll and Daniel Grosu. A Game Theoretic Investigation of Deception in Network Security. In *2009 Proceedings of 18th International Conference on Computer Communications and Networks - ICCCN 2009*, pages 1–6, 2009.

4. Hayreddin Çeker, Jun Zhuang, Shambhu Upadhyaya, Quang Duy La, and Boon-Hee Soong. Deception-Based Game Theoretical Approach to Mitigate DoS Attacks. In *Decision and Game Theory for Security*, pages 18–38. Springer International Publishing, Cham, October 2016.

5. Dorothy E Denning. Framework and principles for active cyber defense. *Computers and Security*, 40:108–113, 2014.

6. Kimberly Ferguson-Walter, D S LaFon, and T B Shade. Friend or Faux: Deception for Cyber Defense. *Journal Of Information Warfare*, 16(2):28–42, August 2017.

7. Galois Inc. CyberChaff$^{\text{TM}}$: Confounding and Detecting Adversaries. https://galois.com/project/cyberchaff/, year=viewed October 2017.

8. Nandan Garg and Daniel Grosu. Deception in Honeynets: A Game-Theoretic Analysis. In *2007 IEEE SMC Information Assurance and Security Workshop*, pages 107–113, 2007.

9. David Garlan, S-W Cheng, A-C Huang, Bradley Schmerl, and Peter Steenkiste. Rainbow: Architecture-based self-adaptation with reusable infrastructure. *Computer*, 37(10):46–54, 2004.

10. Kristin E Heckman, Frank J Stech, Roshan K Thomas, Ben Schmoker, and Alexander W Tsow. *Cyber Denial, Deception and Counter Deception*. Advances in Information Security. Springer International Publishing, Cham, 2015.

11. Illusive Networks. Deception Management System. https://www.illusivenetworks.com/solutions, viewed October 2017.

12. Sushil Jajodia, Anup K. Ghosh, Vipin Swarup, Cliff Wang, and X. Sean Wang. *Moving Target Defense: Creating Asymmetric Uncertainty for Cyber Threats*. Springer Publishing Company, Incorporated, 1st edition, 2011.

13. Christopher Kiekintveld, Viliam Lisý, and Radek Píbil. Game-Theoretic Foundations for the Strategic Use of Honeypots in Network Security. In *Cyber Warfare*, pages 81–101. Springer, Cham, Cham, 2015.

14. Nicholas S Kovach, Alan S Gibson, and Gary B Lamont. Hypergame Theory: A Model for Conflict, Misperception, and Deception. *Game Theory*, (2):1–20, 2015.

15. Q. D. La, T. Q. S. Quek, J. Lee, S. Jin, and H. Zhu. Deceptive Attack and Defense Game in Honeypot-Enabled Networks for the Internet of Things. *IEEE Internet of Things Journal*, 3(6):1025–1035, Dec 2016.

16. Kevin Mitnick and William Simon. *The Art of Deception*. Controlling the Human Element of Security. John Wiley & Sons, Inc., New York, NY, 2003.

17. Radek Píbil, Viliam Lisý, Christopher Kiekintveld, Branislav Bošanský, and Michal Pěchouček. Game Theoretic Model of Strategic Honeypot Selection in Computer Networks? In *Decision and Game Theory for Security*, pages 201–220. Springer, Berlin, Heidelberg, Berlin, Heidelberg, 2012.

18. Lawrence Pingree. Emerging Technology Analysis: Deception Techniques and Technologies Create Security Technology Business Opportunities. *Gartner, Inc*, 2015.

19. James Pita, Manish Jain, Janusz Marecki, Fernando Ordóñez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, and Sarit Kraus. 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, pages 125–132, Richland, SC, 2008. International Foundation for Autonomous Agents and Multiagent Systems.

20. Neil C Rowe and Julian Rrushi. *Introduction to Cyberdeception*. Springer International Publishing, Cham, 2016.

21. Eric Shieh, Bo An, Rong Yang, Milind Tambe, Craig Baldwin, Joseph DiRenzo, Ben Maule, and Garrett Meyer. 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 - Volume 1*, AAMAS '12, pages 13–20, 2012.

22. TrapX Security. Deception in Depth – The Architecture of Choice. https://trapx.com/product/, viewed October 2017.

23. Barton Whaley. *Stratagem: deception and surprise in war*. Artech House, Cambridge, Massachusetts, 1969.

24. Gerald Willard. Understanding the Co-Evolution of Cyber Defenses and Attacks to Achieve Enhanced Cybersecurity. *Journal Of Information Warfare*, 14(2), April 2014.

INITIAL DISTRIBUTION

84300	Library	(1)
85300	Archive/Stock	(1)
71000	S. Fugate	(1)
56150	J. Mauger	(1)
58230	M. Major	(1)

Defense Technical Information Center
Fort Belvoir, VA 22060–6218 (1)

This page left blank intentionally

REPORT DOCUMENTATION PAGE	*Form Approved* *OMB No. 0704-01-0188*

1. REPORT DATE *(DD-MM-YYYY)* July 2018	**2. REPORT TYPE** Final	**3. DATES COVERED** *(From - To)*

4. TITLE AND SUBTITLE

Game Theory for Adaptive Defensive Cyber Deception

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHORS

Sunny Fugate
Justin Mauger
Maxine Major
Space and Naval Warfare Systems Center Pacific

Kimberly Ferguson-Walter
U.S. Department of Defense

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

SSC Pacific
53560 Hull Street
San Diego, CA 92152–5001

8. PERFORMING ORGANIZATION REPORT NUMBER

TR 3141

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

National Security Agency (NSA)
9800 Savage Rd., Suite 6272
Ft. George G. Meade, MD 20755-6000

10. SPONSOR/MONITOR'S ACRONYM(S)

R23

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release.

13. SUPPLIMENTARY NOTES

This is work of the United States Government and therefore is not copyrighted. This work may be copied and disseminated without restriction.

14. ABSTRACT

In prior game theory work, cyber deception games have generally been modeled as non-cooperative, sequential games, where the quality and authenticity of the signal is being manipulated. This work relies on a signal, which can be deceptive, being actively pushed to the attacker. We have formulated a new framework for cyber deception games in which each player has their own perception of the game being played and the moves being taken. A particular player may manipulate other players' perceived payoffs to induce them to take sub-optimal actions. This model of deception seeks to address defender interactions with an attacker following the initial deployment of cyber defenses addressed by previous work. Our primary contribution is a model of defensive cyber deception that incorporates defender control of attacker perception of the cyber environment. Our ultimate goal is to inform future cyber defense systems, enabling more sophisticated responses to attacker behaviors and improving defensive posture.

15. SUBJECT TERMS

Game theory for adaptive defensive cyber deception; adaptive cyber deception; adaptive decoy system;

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			Sunny Fugate
					19B. TELEPHONE NUMBER *(Include area code)*
U	U	U	U	22	(619) 553-7978

This page left blank intentionally

This page left blank intentionally

Approved for public release.

SPAWAR
Systems Center
PACIFIC

UNITED STATES NAVY

SSC Pacific
San Diego, CA 92152-5001