Calhoun: The NPS Institutional Archive
DSpace Repository

Reports and Technical Reports

Faculty and Researchers' Publications

2012-01

A Game Theoretic Model of Strategic Conflict in Cyberspace

Schramm, Harrison; Alderson, David L.; Carlyle, W.
Matthew; Dimitrov, Nediako B.

<https://hdl.handle.net/10945/25616>

Downloaded from NPS Archive: Calhoun

Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>

NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

**A GAME THEORETIC MODEL OF STRATEGIC CONFLICT
IN CYBERSPACE**

by

Harrison Schramm
David L. Alderson
W. Matthew Carlyle
Nedialko B. Dimitrov

January 2012

Approved for public release; distribution is unlimited

Prepared for: Office of Naval Intelligence N2N6CS, 4251 Suitland Road, Washington,
D.C. 20395 and Office of Naval Research, 875 N. Randolph Street, Suite
1425, Arlington, VA 22203

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 31-01-2012		2. REPORT TYPE Technical Report		3. DATES COVERED (From-To) June 2011 – January 2012	
4. TITLE AND SUBTITLE A Game Theoretic Model of Strategic Conflict in Cyberspace				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER N0001412WX20058	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Harrison C. Schramm David L. Alderson W. Matthew Carlyle Nedialko B. Dimitrov				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943				8. PERFORMING ORGANIZATION REPORT NUMBER NPS-OR-11-005	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Schramm was sponsored by Office of Naval Intelligence N2N6CS, 4251 Suitland Road, Washington, D.C. 20395 Alderson and Carlyle were sponsored by the Office of Naval Research, 875 N. Randolph Street, Suite 1425, Arlington, VA 22203 Dimitrov was sponsored by the Naval Postgraduate School Research Initiation Program, Monterey, CA 93943				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT We study cyber conflict as a two-person, zero-sum game in discrete time, where each player discovers new exploits according to an independent random process. Upon discovery, the player must decide if and when to exercise a munition based on that exploit. The payoff from using the munition is a function of time that is (generally) increasing. These factors create a basic tension: the longer a player waits to exercise a munition, the greater his payoff because the munition is more mature, but also the greater the chance that the opponent will also discover the exploit and nullify the munition. Assuming perfect knowledge, and under mild restrictions on the time-dependent payoff function for a munition, we derive optimal exercise strategies and quantify the value of engaging in cyber conflict. Our analysis also leads to high-level insights on cyber conflict strategy.					
15. SUBJECT TERMS Cyber Conflict, Markov Game, Deterrence, Game Theory					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 46	19a. NAME OF RESPONSIBLE PERSON Harrison C. Schramm
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

THIS PAGE INTENTIONALLY LEFT BLANK

NAVAL POSTGRADUATE SCHOOL
Monterey, California 93943-5000

Daniel T. Oliver
President

Leonard A. Ferrari
Executive Vice President and
Provost

The report entitled “*A Game Theoretic Model of Strategic Conflict in Cyberspace*” was prepared for and funded by the Office of Naval Research and Office of Naval Intelligence (N2/N6).

Further distribution of all or part of this report is authorized.

This report was prepared by:

Harrison C. Schramm
CDR, USN

David L. Alderson
Assistant Professor of
Operations Research

W. Matthew Carlyle
Associate Professor of
Operations Research

Nedialko B. Dimitrov
Assistant Professor of
Operations Research

Reviewed by:

Ronald D. Fricker
Associate Chairman for Research
Department of Operations Research

Robert F. Dell
Chairman
Department of Operations Research

Released by:

Douglas Fouts
Interim Vice President and
Dean of Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

We study cyber conflict as a two-person, zero-sum game in discrete time, where each player discovers new exploits according to an independent random process. Upon discovery, the player must decide if and when to exercise a munition based on that exploit. The payoff from using the munition is a function of time that is (generally) increasing. These factors create a basic tension: the longer a player waits to exercise a munition, the greater his payoff because the munition is more mature, but also the greater the chance that the opponent will also discover the exploit and nullify the munition. Assuming perfect knowledge, and under mild restrictions on the time-dependent payoff function for a munition, we derive optimal exercise strategies and quantify the value of engaging in cyber conflict. Our analysis also leads to high-level insights on cyber conflict strategy.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I. INTRODUCTION.....	1
II. RELATED WORK.....	3
III. ANALYSIS	5
A. FOUNDATION	5
1. Discovery of the Exploit.....	5
2. Development of Munition.....	5
3. Employment.....	5
4. Obsolescence.....	5
B. FORMULATION.....	6
1. Discovery.....	6
2. Development	6
3. Employment.....	7
C. ZERO-SUM GAME WITH PERFECT INFORMATION.....	7
1. Markov Game Transitions	7
2. Both Players Know the Exploit.....	9
3. Only One Player Knows the Exploit	10
4. Neither Player Has the Exploit	12
IV. NUMERICAL ANALYSIS	15
A. SCENARIO 1: CONSTANT a_i FUNCTIONS	15
B. SCENARIO 2: LINEARLY INCREASING a_1	15
C. SCENARIO 3: NONMONOTONE a_1	17
V. EXTENSIONS AND APPLICATIONS.....	21
A. DELAYED ACTION.....	21
1. One Player Delayed Action	21
B. DETERRENCE.....	23
1. A Short Review of Strike Stability	23
2. Strike Stability for Cyber Conflict	24
VI. CONCLUSIONS AND FUTURE WORK.....	25
LIST OF REFERENCES.....	27
INITIAL DISTRIBUTION LIST	29

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF FIGURES

- Figure 1. Diagram of states in the Markov game. The arrows in the diagram show the possible transitions from one state to another, as described in Table 1. The horizontal axis describes increases in holding time for Player 1, τ_1 , and the vertical axis describes increases in holding time for Player 2, τ_2 8
- Figure 2. Value of Scenario 2 from Player 1's point of view. The vertical axis plots the value, $v_\tau(h)$, as a function of the number of time periods Player 1 waits before attacking, h . The value function increases to the point $h=5$, and decreases afterward. By Theorem 2, this implies that Player 1's optimal attack time, k_1^* , is 5. 16
- Figure 3. Value of Scenario 2 as a function of Player 1's probability of discovering the Exploit, p_1 . Here we see that the value of the game is a concave function of Player 1's probability of detecting the exploit. Increases in detection probability at low detection values provide a bigger increase in the game value than increases in detection probability at high detection values. 17
- Figure 4. The function a_1 for Scenario 3. Unlike our previous examples, the value of Player 1's attack has a dip at $\tau_1 = 6$. In this scenario, Theorem 2 no longer applies in finding the optimal attack time, k_1^* 18
- Figure 5. Player 1's value as a function of waiting time, h in Scenario 3. We see that the payoff for waiting to $h=7$ is less than executing at $h=5$ 18
- Figure 6. Player 1's utility curve as a function of waiting time, w , against an evenly matched opponent. We see that Player 1's utility drops off rapidly from an expected value of zero, with the implication that waiting is costly. 22
- Figure 7. Player 1's required detection probability, p_1 , required to achieve $V^w\langle 0,0,0 \rangle = 0$ as a function of waiting time, w . Player 1's required capability increases rapidly and, because p_1 may never be greater than 1, parity is unachievable after $w=9$ 23

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

Table 1.	Markov game action sets and state transitions as a function of $\langle T, \tau_1, \tau_2 \rangle$, the state of the game. The game always starts in $\langle T, 0, 0 \rangle$. As Player i discovers the exploit, τ_i , becomes greater than zero and Player i 's action set includes attack.	8
Table 2.	Payoff matrix for the Markov game when both players know the exploit. The payoff associated with "Wait, Wait" depends on the future evolution of the game... 9	9
Table 3.	Payoffs for the case where Player 1 knows the exploit and Player 2 does not. By definition, Player 2 has a singleton action set and the matrix reduces to a single column. Player 1 prefers to attack if $Y \leq a_1(\tau)$	10

THIS PAGE INTENTIONALLY LEFT BLANK

EXECUTIVE SUMMARY

Conflict in cyberspace is difficult to analyze; methods developed for other dimensions of conflict, such as land warfare, war at sea, and missile warfare, do not adequately address cyber conflict. A characteristic that distinguishes cyber conflict is that actors do not know the true state of their arsenal(s)—i.e., an opponent may negate a potential attack by discovering and fixing the vulnerability in their system; they may do this without knowledge of their adversary's intentions.

Our analysis focuses on the *National* level, with decisions and actions that would be available to a Unified Commander. This is fundamentally different than analyses that are focused on the defense of a specific technological system.

In this report, we develop a rigorous game-theoretic description of two players and a single vulnerability. We do so under an assumption of *perfect information*, in the sense that as soon as a player discovers a vulnerability he knows with certainty if the adversary has also discovered it. We consider the decisions facing a Commander with limited resources who has a single decision: Upon discovering vulnerability, he may:

- ***Wait:*** Waiting increases the damage of a munition based on the vulnerability; however, it also risks the adversary's discovery of the vulnerability, negating the munitions' effectiveness.
- ***Attack:*** Attacking exercises an available munition.

From these assumptions, we gain the following insights:

- ***Success requires rapid action.*** Our model shows that delays in taking action reduce the chance of a player's success in cyber conflict. Such delays can come from a variety of sources, including bureaucratic or command restrictions. A byproduct of our model is the calculation of how proficient a player must be in other areas to make up for delays in taking action; in most cases, the required capability is unattainable. The immediate consequence of this is that command structures in cyberspace should be agile with the correct level of delegation of authority.
- ***Prospects for deterrence in cyber conflict may be limited.*** The ability of players to deter their opponents from attacking depends on an assured second strike. In cyber conflict, opposing players may have munitions based on the same exploit, and the first player to use the exploit effectively removes second strike munitions from the opponent's arsenal. Complicating factors to the cyber conflict game, such as an inability to identify the player who performed a cyber attack, or a player's ability to respond with kinetic munitions, also have an effect on deterrence in cyber conflict.

The framework contained herein not only informs the decision facing a commander in conflict, but also allows for exploratory analysis—particularly in the

trade-offs between speed of detection and speed of attack development. Therefore, this model may be useful for both cyber warriors and budget analysts.

I. INTRODUCTION

Conflict in Cyberspace, or *cyber conflict*, is important at both the strategic and tactical levels. In this paper, we consider the strategic decisions made by states or other groups about when and how to engage in cyber conflict. The increasing dependency on interconnected networks, both in military and civilian life, means that little is beyond the reach of cyberspace. Cyberspace plays a central role in our social, economic, and civic welfare. It is, therefore, not surprising that the United States “has identified cyber security as one of the most serious economic and national security challenges we face as a nation” (National Security Council, 2010). Consequently, security and defense in cyberspace has become an increasingly large part of the defense budget (Stervstein, 2011).

A defining characteristic of cyber conflict is the way in which weapons in cyberspace are discovered, developed, and employed. Players search for mechanisms that can cause cyber systems to perform in ways not intended in their original design, called *exploits*, and, once found, develop them into one or more *cyber munitions*. These munitions can then be used as part of a *cyber attack*. In searching for exploits to use against an adversary, a player may also discover flaws in their own system and decide to *patch* them so an adversary cannot use them. Moreover, a player could develop munitions based on an exploit that the adversary independently fixes, thereby making the munitions obsolete. Thus, collections of cyber munitions, or *arsenals*, are dynamic and their effectiveness depends on the relative state of knowledge of the opponents.

In this context, apparently simple questions such as “how long should we hold a munition in development before using it in an attack?” and “how should we allocate limited resources to offense versus defense?” require novel, analytical models. Moreover, the dynamic nature of cyber weapons development and obsolescence makes it difficult to assess the potency of an arsenal; this is true for assessing our own arsenal as well as an arsenal belonging to an adversary. Clear, useful analysis at the national level is important both for making sound future investment decisions and for creating informed strategic and policy guidance.

To analyze the strategic decisions involved in cyber conflict, we use a game theoretic framework—we view cyber warfare as a game consisting of attacks that opposing players exercise at a time of their choosing. Each player discovers, develops, and chooses to exercise attacks to maximize the value of their cyber operations. Our analysis is independent of specific technologies, and does not assume an explicit cyber system or exploit.

Using minimal assumptions, our model leads to two fundamental insights:

- ***Success requires rapid action.*** Our model shows that delays in taking action reduce the chance of a player’s success in cyber conflict. Such delays can come from a variety of sources, including bureaucratic or command restrictions. A byproduct of our model is the calculation of how proficient a player must be in other areas to make up for delays in taking action; in most cases, the required capability is unattainable. The

immediate consequence of this is that command structures in cyberspace should be agile with the correct level of delegation of authority.

- ***Prospects for deterrence in cyber conflict may be limited.*** The ability of players to deter their opponents from attacking depends on an assured second strike. In cyber conflict, opposing players may have munitions based on the same exploit, and the first player to use the exploit effectively removes second strike munitions from the opponent's arsenal. Complicating factors to the cyber conflict game, such as an inability to identify the player who performed a cyber attack, or a player's ability to respond with kinetic munitions, also have an effect on deterrence in cyber conflict.

II. RELATED WORK

The JASON 2010 report, *The Science of Cyber-Security* (JASON, 2010), recommends a variety of analytic approaches and suggests borrowing ideas from other sciences such as physics, cryptography, and biological sciences, including epidemiology. The JASON report introduces a two-player, stationary, discrete-time model called the Forwarder's Dilemma as an example of what a game-theoretic analysis might look like. This game considers whether an administrator should forward another system's messages on their network and is similar both in format and solution to the well-known Prisoner's Dilemma (e.g. Fudenberg & Tirole, 1991). Lye and Wing (2002) and Shen, Chen, Blasch, and Tadda (2007) also consider cyber attacks in the context of a game. The most comprehensive survey of game theory and cyberspace is by Shiva, Dasgupta, and Wu (2010). They develop a taxonomy of game theoretic models with two broad categories:

- **Static versus Dynamic.** A "one shot" cyber conflict game, where players choose plans of action and then execute them simultaneously, is a static game. A cyber conflict game with multiple stages and sequential decisions is a dynamic game.
- **Available Information.** Players may have exact, imperfect, or no knowledge about their opponent's intentions or capabilities. If the players know the actions of other players once taken, this is called a game with perfect information. If the players know the structure of the game and payoffs, but not the actions, this is called a game with complete information. Finally, a game in which the payoffs evolve in time in a random process is a stochastic game.

While game theory considers both cooperative and noncooperative games, work to date on cyber conflict deals only with noncooperative games. In the taxonomy of Shiva et al. (2010) our proposed model is a noncooperative, dynamic, stochastic game with perfect information.

The previous study that has the most commonality with our approach is that of Lye and Wing (2002). They consider a two-player, stochastic game between an attacker and administrator. Their model is at the machine level; it focuses on an attacker attempting to find the best policy among a portfolio of several attacks to damage a university computer network. This game theoretic model of Lye and Wing maps to the tactical level of conflict, as opposed to our model that is focused at the strategic level between two players engaged in cyber conflict.

Our work differs from previous work by abstracting cyber conflict away from individual machines and instruction sets in the same manner that Lanchester equations (Washburn & Kress, 2009) abstract physical conflict away from soldiers and weapons. The goal of this paper is to provide a foundation from which to build more complex models, towards the ultimate goal of integrating the cyber domain into the spectrum of conflict analysis in order to support strategic models for decision makers at the national level.

THIS PAGE INTENTIONALLY LEFT BLANK

III. ANALYSIS

A. FOUNDATION

As defined previously, a computer system may contain exploits. These are unknown until discovered, after which they can be fixed in the form of a *patch* or weaponized into a *munition*. We model the life-cycle of a single cyber exploit as a four-stage process.

1. Discovery of the Exploit

We model the discovery of a single exploit by each player as a random process, occurring independently for each player, which may depend on factors such as training, investment, experience and luck.

2. Development of Munition

Once an exploit is discovered, a player can develop a munition based on the exploit. We assume that there is a relationship between the length of time that a player knows about an exploit and the effectiveness of the munition he develops based on that exploit. Munitions may only be developed for known exploits.

3. Employment

Once a munition is developed, it can be employed at will against an adversary in an attack.

4. Obsolescence

Consider a game between two players, Player 1 and Player 2. If Player 1 discovers an exploit in his system and patches it before Player 2 can develop and employ a munition based on that exploit, then that munition becomes obsolete.

Uncertainties about the obsolescence of a player's own arsenal are a key dimension in the analysis of cyber conflict. For the purposes of this analysis, we assume that a player who is aware of an exploit also knows whether the other player(s) are aware of the same exploit; this removes one type of uncertainty. For a player who is unaware of an exploit, we assume neither player knows the time until the unaware player discovers the exploit. This uncertainty in discovery times is the fundamental tension that our model seeks to explore.

We model cyber warfare as a Markov game (Thie, 1983; Fudenberg & Tirole, 1991) where the choices available to each player depend on the number of exploits known by each player and the strength of each player's munitions. In general, there may be multiple exploits that each player discovers, develops into munitions, and uses to attack, but we choose to focus our analysis on a scenario where there is only a single exploit to be discovered. At the beginning of this scenario, neither player knows the

exploit. Each player probabilistically discovers the exploit, and when either player chooses to attack, then payoffs are determined and the game terminates.

B. FORMULATION

Our model focuses on a strategic cyber conflict between two players, where there is a single exploit to be discovered. Let i index the players $i \in \{1, 2\}$. The mathematical notation used to describe the game falls into three broad categories: Discovery, Development, and Employment.

1. Discovery

Let T be the duration of time that an exploit has existed, which we also call the *clock time*. Without loss of generality, we assume that the game starts when the exploit is created. We create a discrete-time model, with T increasing over the set of positive integers. If the exploit was part of the original system, then T is the age of the system. If the exploit was introduced as part of a software upgrade, then T is the age of the upgrade. Let d_i be the moment in clock time that Player i discovers the exploit. We define $\tau_i = \max(0, T - d_i)$ to be the relative time that player i has known about the exploit; we call this Player i 's *holding time*. By definition, if Player i is not aware of the exploit, then $\tau_i = 0$. We define a state of the cyber game, S , as:

$$S = \langle T, \tau_1, \tau_2 \rangle,$$

where the elements of this three-tuple represent how long the exploit has existed, how long Player 1 has known the exploit, and how long Player 2 has known the exploit, respectively.

2. Development

A player's success in cyber conflict depends on both his ability to discover exploits and his ability to develop effective munitions. We assume that at any moment following the discovery d_i , Player i has the ability to create and deploy a perfectly effective patch. However, we assume that the act of deploying the patch effectively announces it to the adversary; so patching nullifies everyone's munitions based on that exploit, and this ends the game for both sides. Let $p_i(T)$ denote the probability that Player i discovers an exploit as clock time progresses from period T to period $T + 1$. For convenience, let $q_i(T) = 1 - p_i(T)$. Let $a_i(\tau_i)$ be the value of an attack by Player i using a munition developed using a holding time of τ_i . The value of an attack is a function of τ instead of T because we assume that once the exploit is known, the effectiveness of the munition depends on holding time and not clock time. We impose two constraints on $a_i(\tau_i)$. First, we assume $a_i(0) = 0$; namely, that if an exploit is not known, then an attack based on it has no value. Additionally, we assume $0 \leq a_i(\tau) \leq B_i$, where B_i is an arbitrary upper bound, thus disallowing cyber attacks with either a negative value or an infinite value.

3. Employment

Once a player has a cyber munition, he may choose to use it. Let $\theta_i(T)$ denote the action set of Player i at time T . We define $\theta_i(T) \subseteq \{W, A\}$ where:

- **W : Wait.** While a player is waiting, he is either waiting to discover the exploit ($\tau_i = 0$) or he may know about the exploit ($\tau_i > 0$) and be working to make his munitions more effective.
- **A : Attack.** When a player attacks he receives the value of his attack at that time. Attacking also broadcasts the attack's underlying exploit to all players.

A player who does not know the exploit has a singleton action set, $\{W\}$, and a player that does know the exploit has the full action set, $\{W, A\}$.

C. ZERO-SUM GAME WITH PERFECT INFORMATION

To fully specify the game, we must define action sets for each player, and the utilities for player's actions. We assume a zero-sum, strategic conflict; i.e., that any utility gain by one player results in an equal utility loss by the opponent. We use the convention that Player 1 is a maximizing player and Player 2 is a minimizing player. We assume that each player knows the state of the Markov game, S . But this perfect information assumption does not mean that a player knows the exploit. A player is still limited by his action set. For example, if the state of the game is $\langle T, 1, 0 \rangle$, it means that: Player 1 knows the exploit, has a holding time of 1, and has an action set of $\{W, A\}$; while Player 2 does not know the exploit, has a holding time of 0, and therefore has an action set of solely $\{W\}$.

1. Markov Game Transitions

The discovery and development of attacks is modeled as transitions in the state of the Markov game. The game begins in the state $\langle 0, 0, 0 \rangle$ and proceeds in discrete rounds. In each round, the clock time T increases deterministically. For each Player i , the holding time $\tau_i = 0$ until the player discovers the exploit. Exploit discovery happens with probability $p_i(T)$ for Player i in round T . Once an exploit is discovered by a player, the player's holding time increases deterministically. The resulting transitions of the Markov game state are summarized in Table 1. A visual depiction of the states of the game is presented in Figure 1.

	$\tau_2 = 0$	$\tau_2 > 0$
	$\theta_1 = \{W\}; \theta_2 = \{W\}$	$\theta_1 = \{W\}; \theta_2 = \{A, W\}$
$\tau_1 = 0$	$\langle T, 0, 0 \rangle \left\{ \begin{array}{l} \xrightarrow{(1-p_1(T))(1-p_2(T))} \langle T+1, 0, 0 \rangle \\ \xrightarrow{p_1(T)(1-p_2(T))} \langle T+1, 1, 0 \rangle \\ \xrightarrow{(1-p_1(T))p_2(T)} \langle T+1, 0, 1 \rangle \\ \xrightarrow{p_1(T)p_2(T)} \langle T+1, 1, 1 \rangle \end{array} \right.$	$\langle T, 0, \tau_2 \rangle \left\{ \begin{array}{l} \xrightarrow{(1-p_1(T))} \langle T+1, 0, \tau_2+1 \rangle \\ \xrightarrow{p_1(T)} \langle T+1, 1, \tau_2+1 \rangle \end{array} \right.$
$\tau_1 > 0$	$\theta_1 = \{A, W\}; \theta_2 = \{W\}$ $\langle T, \tau_1, 0 \rangle \left\{ \begin{array}{l} \xrightarrow{(1-p_2(T))} \langle T+1, \tau_1+1, 0 \rangle \\ \xrightarrow{p_2(T)} \langle T+1, \tau_1+1, 1 \rangle \end{array} \right.$	$\theta_1 = \{A, W\}; \theta_2 = \{A, W\}$ $\langle T, \tau_1, \tau_2 \rangle \xrightarrow{1} \langle T+1, \tau_1+1, \tau_2+1 \rangle$

Table 1. Markov game action sets and state transitions as a function of $\langle T, \tau_1, \tau_2 \rangle$, the state of the game. The game always starts in $\langle T, 0, 0 \rangle$. As Player i discovers the exploit, τ_i , becomes greater than zero and Player i 's action set includes attack.

Figure 1. Diagram of states in the Markov game. The arrows in the diagram show the possible transitions from one state to another, as described in Table 1. The horizontal axis describes increases in holding time for Player 1, τ_1 , and the vertical axis describes increases in holding time for Player 2, τ_2 .

Let $V\langle T, \tau_1, \tau_2 \rangle$ define the *value* of the game in state $\langle T, \tau_1, \tau_2 \rangle$; this value represents the expected value to the players if they play the game starting at that state. Because the game is zero sum, payoffs for both players can be described by a single value. To analyze the game, we seek to characterize this value function. In particular, $V\langle 0, 0, 0 \rangle$ is the value of engaging in cyber conflict. We seek to characterize $V\langle T, \tau_1, \tau_2 \rangle$

for every state of the Markov game. We proceed in our analysis by considering three cases on τ_1, τ_2 .

2. Both Players Know the Exploit

In this case, we have $\tau_1 > 0, \tau_2 > 0$ and both players have full action sets, meaning each may attack or wait. Table 2 represents the payoffs of the Markov game in such a state in matrix form. Each entry in the matrix contains a single real number, since the game is zero sum. If both players wait, the value is determined by future play. If one player attacks and the other waits, the attacking player receives the full value of his munition. If both players attack simultaneously, the sum of the munition values gives the result of the game.

	Player 2 plays: W	Player 2 plays: A
Player 1 plays: W	$V\langle T+1, \tau_1+1, \tau_2+1 \rangle$	$-a_2(\tau_2)$
Player 1 plays: A	$a_1(\tau_1)$	$a_1(\tau_1) - a_2(\tau_2)$

Table 2. Payoff matrix for the Markov game when both players know the exploit. The payoff associated with “Wait, Wait” depends on the future evolution of the game.

This leads to the following observation.

Theorem 1. For any game state $\langle T, \tau_1, \tau_2 \rangle$ such that $\tau_1 > 0$ and $\tau_2 > 0$, “Attack, Attack” is an iterated elimination of dominated strategies equilibrium with a value of $a_1(\tau_1) - a_2(\tau_2)$.

Proof. Suppose $V\langle T+1, \tau_1+1, \tau_2+1 \rangle \geq 0$. Then $V\langle T+1, \tau_1+1, \tau_2+1 \rangle \geq -a_2(\tau_2)$ and $a_1(\tau_1) \geq a_1(\tau_1) - a_2(\tau_2)$. Therefore, “Attack” is a dominating strategy for Player 2. Given Player 2 chooses “Attack,” Player 1 must also play “Attack” and “Attack, Attack” is an equilibrium. A symmetric argument holds if $V\langle T+1, \tau_1+1, \tau_2+1 \rangle \leq 0$.

□

Theorem 1 results in the following corollary.

Corollary 1. If the game starts in state $\langle T, \tau_1, \tau_2 \rangle$, with $\tau_1 > 0$ and $\tau_2 > 0$, the game terminates immediately and

$$V\langle T, \tau_1, \tau_2 \rangle = a_1(\tau_1) - a_2(\tau_2).$$

Interpreting the results of Theorem 1 and the above corollary, a game starting in $\langle T, 0, 0 \rangle, T \geq 0$ ends, optimally, no later than when one of the following states is reached: $\langle T, 1, \tau_2 \rangle$ or $\langle T, \tau_1, 1 \rangle$. However, the game may also end earlier if a player who discovers the exploit chooses to attack before the second player has discovered the exploit. Because each i , $a_i(\cdot)$ has a unique, associated τ_i , for ease of exposition we drop the index i from

future uses of τ . For the remainder of this paper, statements like $a_2(\tau)$ should be understood to mean $a_2(\tau_2)$.

3. Only One Player Knows the Exploit

For simplicity, we develop the theory from a state where Player 1 has the exploit and Player 2 does not. The analysis follows identical lines in the opposing situation. In this case, Player 1 has a full action set and Player 2 may only wait to discover the exploit $\theta_1 = \{A, W\}, \theta_2 = \{W\}$. Suppose the state of the game is $\langle T, \tau, 0 \rangle$. We define

$$Y = (1 - p_2(T))V\langle T+1, \tau_1+1, 0 \rangle + p_2(T)V\langle T+1, \tau_1+1, 1 \rangle$$

to be the expected utility if both players choose to wait at time T . Table 3 displays the payoffs in matrix form.

	Player 2 Plays: Wait
Player 1 Plays: Wait	Y
Player 1 Plays: Attack	$a_1(\tau)$

Table 3. Payoffs for the case where Player 1 knows the exploit and Player 2 does not. By definition, Player 2 has a singleton action set and the matrix reduces to a single column. Player 1 prefers to attack if $Y \leq a_1(\tau)$.

The fundamental analytic question is “from which states does Player 1 prefer to attack?” If Player 2 discovers the exploit, the game transitions to the scenario described previously and immediately concludes as specified in Theorem 1. We characterize states $\langle T, \tau, 0 \rangle$ from which Player 1 prefers to attack as follows. We define $v_\tau(h)$ as the expected utility to Player 1 if he waits h time periods before attacking, starting in state $\langle T, \tau, 0 \rangle$.

In particular, we have:

$$\begin{aligned} v_\tau(0) &= a_1(\tau) \\ v_\tau(1) &= q_2(T) \cdot a_1(\tau+1) + p_2(T) \cdot (a_1(\tau+1) - a_2(1)) \\ v_\tau(2) &= q_2(T+1)q_2(T) \cdot a_1(\tau+2) + p_2(T+1)q_2(T) \cdot (a_1(\tau+2) - a_2(1)) + \\ &\quad p_2(T) \cdot (a_1(\tau+1) - a_2(1)) \end{aligned}$$

$$v_\tau(h) = a_1(\tau+h) \cdot \prod_{k=0}^{h-1} q_2(T+k) + \sum_{k=0}^{h-1} \left[(a_1(\tau+k+1) - a_2(1)) \cdot p_2(T+k) \prod_{j=0}^{k-1} q_2(T+j) \right]. \quad (1)$$

The definition of $v_\tau(h)$ allows us to evaluate the states from which Player 1 prefers to attack. Player 1 prefers to attack rather than wait in state $\langle T, \tau, 0 \rangle$ if and only if the following holds:

$$a_1(\tau) = v_\tau(0) \geq v_\tau(h) \text{ for all } h \geq 1. \quad (2)$$

This statement mirrors our intuition that a player should attack only if an immediate attack results in a higher utility than waiting for any number of turns before attacking.

Theorem 2. If $a_1(\tau)$ is concave and nondecreasing, and $p_2(T)$ is nondecreasing, then $v_\tau(0) \geq v_\tau(1)$ implies that Player 1 should attack in state $\langle T, \tau, 0 \rangle$ (i.e., Player 1 can never do better by waiting).

Proof. We proceed by showing that the theorem assumptions imply that

$$v_\tau(0) \geq v_\tau(h) \text{ for all } h \geq 2.$$

Consider the quantity

$$\begin{aligned} v_\tau(h+1) - v_\tau(h) &= a_1(\tau+h+1) \prod_{k=0}^h q_2(T+k) - \\ &\quad a_1(\tau+h) \prod_{k=0}^{h-1} q_2(T+k) + (a_1(\tau+h+1) - a_1(\tau+h)) p_2(T+h) \prod_{j=0}^{h-1} q_2(T+j) \\ &= \prod_{k=0}^{h-1} q_2(T+k) [a_1(\tau+h+1) - a_1(\tau+h) - a_2(1) p_2(T+h)] . \end{aligned}$$

We know that $v_\tau(0) \geq v_\tau(1)$, which implies that

$$\begin{aligned} 0 &\geq v_\tau(1) - v_\tau(0) \\ &= a_1(\tau+1) - a_1(\tau) - p_2(T) a_2(1) \\ &\geq a_1(\tau+h+1) - a_1(\tau+h) - p_2(T) a_2(1), \end{aligned}$$

where the last inequality came from the fact that $a_1(\cdot)$ is concave and nondecreasing. Continuing with the last expression above, we have

$$\begin{aligned} 0 &\geq a_1(\tau+h+1) - a_1(\tau+h) - p_2(T) a_2(1) \\ &\geq a_1(\tau+h+1) - a_1(\tau+h) - p_2(T+h) a_2(1), \end{aligned}$$

where the last inequality came from the fact that $p_2(\cdot)$ is nondecreasing and $a_2(1)$ is nonnegative. Finally, multiplying both sides of the inequality by the positive number

$$\prod_{k=0}^{h-1} q_2(T+k), \text{ gives}$$

$$\begin{aligned}
0 &\geq \prod_{k=0}^{h-1} q_2(T+k)[a_1(\tau+h+1)-a_1(\tau+h)-p_2(T+h)a_2(1)] \\
&= v_\tau(h+1)-v_\tau(h).
\end{aligned} \tag{3}$$

We can complete the proof as follows:

$$\begin{aligned}
v_\tau(h)-v_\tau(0) &= v_\tau(h)-v_\tau(h-1)+ \\
&\quad v_\tau(h-1)-v_\tau(h-2)+ \\
&\quad v_\tau(h-2)\dots \\
&\quad v_\tau(1)-v_\tau(0).
\end{aligned}$$

Each of the paired terms on the right-hand side is smaller than zero, by Equation (2); thus, we have

$$v_\tau(h)-v_\tau(0) \leq 0,$$

completing the proof.

□

For the remainder of this paper, we assume stationary probabilities $p_i(T)=p_i \forall T$. Theorem 2 shows that $v_\tau(0) \geq v_\tau(1)$ is sufficient to prefer Attack at a holding time of τ , while Equation (1) shows that $v_\tau(0) \geq v_\tau(1)$ is necessary to prefer Attack at τ . Therefore, from state $\langle T, 1, 0 \rangle$ Player 1 waits for $k^* = \min_k \{v_k(0) \geq v_k(1)\}$ turns before attacking. Substituting the definition of $v_\tau(\cdot)$, we can write this as $k^* = \min_k \{a_1(k+1)-a_1(k) \leq p_2 a_2(1)\}$. The set in the definition of k^* is never empty when $a_1(\cdot)$ is bounded, concave, and nondecreasing, and $p_2 a_2(1)$ is not identically zero, meaning that Player 1 will eventually prefer to attack. We conclude that:

$$V\langle T, 1, 0 \rangle = v_0(k^*). \tag{4}$$

While we presume that most cases will have nondecreasing a_1, a_2, p_1, p_2 functions, there is no reason that it must be so. Nondecreasing functions model situations where the passage of time brings increased capability, both in development and detection. However, there may be interesting, and operationally relevant, cases where the functions are decreasing. Although we do not present detailed results here, the value functions in these alternate situations may be evaluated directly by using Equations (1) and (2).

4. Neither Player Has the Exploit

In this case, the game has been in play for an unknown amount of time and $\tau_1 = \tau_2 = 0$; therefore, both players have singleton action sets,

$$\begin{aligned}
\theta_1 &= \{W\} \\
\theta_2 &= \{W\}.
\end{aligned}$$

Using the theory previously developed, the value of the game, given that Player 1 discovers the exploit first, is: $V\langle T, 1, 0 \rangle$. Similarly, if Player 2 discovers the exploit first, the value is: $V\langle T, 0, 1 \rangle$. In the case where both players simultaneously discover the exploit: $V\langle T, 1, 1 \rangle = a_1(1) - a_2(1)$. Because the state $\langle T, 0, 0 \rangle$ transitions into previously analyzed states, we are only concerned with the first transition. For stationary discovery probabilities, the next state transition probabilities out of $S = \langle T, 0, 0 \rangle$ are:

$$\begin{aligned}\Pr\{\text{next state is } \langle T, 1, 0 \rangle\} &= \gamma_{1,0} = \frac{p_1(1-p_2)}{p_1(1-p_2) + p_2(1-p_1) + p_1p_2} \\ \Pr\{\text{next state is } \langle T, 0, 1 \rangle\} &= \gamma_{0,1} = \frac{p_2(1-p_1)}{p_1(1-p_2) + p_2(1-p_1) + p_1p_2} \\ \Pr\{\text{next state is } \langle T, 1, 1 \rangle\} &= \gamma_{1,1} = \frac{p_1p_2}{p_1(1-p_2) + p_2(1-p_1) + p_1p_2},\end{aligned}$$

where we have introduced the γ values for brevity.

The value of the game starting from $\langle T, 0, 0 \rangle$ is

$$\begin{aligned}V\langle T, 0, 0 \rangle &= \gamma_{1,0}V\langle T, 1, 0 \rangle - \gamma_{0,1}V\langle T, 0, 1 \rangle + \gamma_{1,1}V\langle T, 1, 1 \rangle \\ &= \gamma_{1,0}v_0^1(k_1^*) - \gamma_{0,1}v_0^2(k_2^*) + \gamma_{1,1}(a_1(1) - a_2(1)),\end{aligned}\tag{5}$$

where the negative sign comes from the fact that Player 1 is a maximizing player and Player 2 is a minimizing player, $v_0^1(\cdot), k_1^*$ denotes the results of Equations (3) and (4) if Player 1 is the first to discover the exploit, while $v_0^2(\cdot), k_2^*$ denotes the results of Equations (3) and (4) if Player 2 is the first to discover the exploit.

THIS PAGE INTENTIONALLY LEFT BLANK

IV. NUMERICAL ANALYSIS

In this section, we consider some concrete examples of the theory developed in the previous section. Unless otherwise specified, we assume $p_i(T) = p_i \forall T$ and $p_i \neq 0$. As a notational convenience, we denote the value of any particular example as V^n , where n is the example number.

A. SCENARIO 1: CONSTANT a_i FUNCTIONS

Suppose that Players 1 and 2 both have attack value functions such that:

$$\begin{aligned} a_i(0) &= 0 \\ a_i(\tau) &= c_i \quad \forall \tau \geq 1 \end{aligned}$$

Because $a_i(\tau)$ is concave and increasing for both players, we can use Theorem 2 to compute the optimal attack time for each player, k_i^* for $i=1,2$, which is 1 for both players. We may directly compute the value of the game using Equation (5):

$$V^1 = \frac{p_1(1-p_2)a_1(1) - p_2(1-p_1)a_2(1) + p_1p_2(a_1(1) - a_2(1))}{p_1(1-p_2) + p_2(1-p_1) + p_1p_2}.$$

In particular, Player 1 has a positive expected payoff if and only if:

$$p_1a_1(1) > p_2a_2(1).$$

In this case, a player may make up for a deficiency in either discovery or development by being strong in the other area. Because $0 \leq p_i \leq 1$, these trade-offs are implicitly limited.

B. SCENARIO 2: LINEARLY INCREASING a_i

Suppose Players 1 and 2 have attack functions such that:

$$\begin{aligned} a_i(0) &= 0 \\ a_1(\tau) &= \tau \quad 1 \leq \tau \leq 5 \\ a_1(\tau) &= 5 \quad \forall \tau \geq 5 \\ a_2(\tau) &= c \quad \forall \tau \geq 1 \end{aligned}$$

This function is also concave and increasing, and we may use Theorem 2 to determine the optimal attack time, k_i^* , for both players. Specifically, $k_2^* = 1$ and k_1^* is dependent on the values of p_2 and c as follows:

$$k_1^* = \begin{cases} 1 & \text{if } p_2c \geq 1 \\ 5 & \text{otherwise} \end{cases}.$$

As verification, we compute the values of $v_\tau(h)$ for $h=1,2,\dots,5$. We see in Figure 2 that the maximizing value is $h=5$. For example, if $a_2(1)=1, p_2=0.2$.

Figure 2. Value of Scenario 2 from Player 1's point of view. The vertical axis plots the value, $v_\tau(h)$, as a function of the number of time periods Player 1 waits before attacking, h . The value function increases to the point $h=5$, and decreases afterward. By Theorem 2, this implies that Player 1's optimal attack time, k_1^* , is 5.

Knowing k^* for both players, we may compute the value of the game, $V^2\langle T, 0, 0 \rangle$ as a function of p_1 ; see Figure 3.

Figure 3. Value of Scenario 2 as a function of Player 1's probability of discovering the Exploit, p_1 . Here we see that the value of the game is a concave function of Player 1's probability of detecting the exploit. Increases in detection probability at low detection values provide a bigger increase in the game value than increases in detection probability at high detection values.

C. SCENARIO 3: NONMONOTONE a_1

Suppose that $a_2(1)=1$, $p_2=0.3$, and Player 1's value function has a single dip, specifically $a_1(\tau)=(1, 2, 3, 4, 5, 3, 6,)$, as shown in Figure 4. In this case, we cannot use Theorem 2 to compute the optimal attack time. However, we may compute the optimal attack time directly, by computing the value of holding for each possible holding period, as depicted in Figure 5.

Figure 4. The function a_1 for Scenario 3. Unlike our previous examples, the value of Player 1's attack has a dip at $\tau_1 = 6$. In this scenario, Theorem 2 no longer applies in finding the optimal attack time, k_1^* .

Because $a_1(\tau)$ is not concave and increasing, we cannot apply to Theorem 2. Here we need to actually compute the numeric values of $v_\tau(h)$. Performing this calculation, we see that $k_1^* = 5$ and it is not advisable to wait through the nonincreasing region.

Figure 5. Player 1's value as a function of waiting time, h in Scenario 3. We see that the payoff for waiting to $h = 7$ is less than executing at $h = 5$.

A decision maker may want to know what value of $a_1(7)$ would change Player 1's decision? We answer this question by performing a line search on $a_1(7)$ and determine the threshold value is ≈ 6.6 .

THIS PAGE INTENTIONALLY LEFT BLANK

V. EXTENSIONS AND APPLICATIONS

In this section, we explore the operationally relevant implications of our model.

A. DELAYED ACTION

It may be the case that a player discovers an exploit and cannot take action; specifically, he is unable (or not allowed) to attack, patch, or work towards development of a munition for some predetermined fixed time after discovery of an exploit. This may be due to legal, policy, or organizational limitations.

1. One Player Delayed Action

Suppose Player 1 has a rule where he must wait w time periods after discovery before any attack, patch, or development of a munition. Consistent with our previous definition of perfect information, if Player 2 has the exploit, he learns if Player 1 knows the exploit. Player 2 also knows the existence and duration of Player 1's delay rule.

We wish to understand the value of this delayed version of our game, which we denote as $V^w \langle \cdot \rangle$.

If both players have the exploit, Player 2 can wait and exercise his munition the turn before Player 1 is able to begin work; therefore,

$$V^w \langle T, 1, 1 \rangle = -a_2 (w-1).$$

If Player 2 has the exploit and Player 1 does not, Player 2 may continue developing his munition until Player 1 discovers the exploit, and an additional $(w-1)$ time periods before attacking; therefore,

$$V^w \langle T, 0, 1 \rangle = - \sum_{i=0}^{\infty} p_1 (1-p_1)^i a_2 (i+w).$$

Finally, if Player 1 has the exploit and Player 2 does not, there are two possibilities. First, Player 1 may retain sole knowledge of the exploit until the end of the waiting period, or, second, Player 2 may discover the exploit during Player 1's forced delay time; therefore,

$$V^w \langle T, 1, 0 \rangle = (1-p_2)^w V \langle T, 1, 0 \rangle - \sum_{i=1}^{w-1} p_2 (1-p_2) a_2 (w-i).$$

We may combine these expressions to write:

$$V^w \langle T, 0, 0 \rangle = \gamma_{1,0} \left[(1-p_2)^w V \langle T, 1, 0 \rangle - \sum_{i=1}^{w-1} p_2 (1-p_2) a_2(w-i) \right] - \gamma_{0,1} \left[\sum_{i=0}^{\infty} p_1 (1-p_1)^i a_2(i+w) \right] - \gamma_{1,1} a_2(w-1). \quad (6)$$

The implication of this is that unproductive waiting times are damaging to a player's prospects in cyber conflict.

Consider the specific example of two evenly matched players with bounded, linear development functions; thus: $p_1 = p_2 = 0.1$, $a_1(\tau) = a_2(\tau) = \tau$ for $0 < \tau \leq 10$ and $a_1(\tau) = a_2(\tau) = 10$ for $\tau > 10$. By symmetry, $V \langle T, 0, 0 \rangle = 0$ for this game when neither player is forced to wait.

Now consider the case where Player 1 has a waiting time, w . We plot Player 1's expected payoff as a function of w in Figure 6.

Figure 6. Player 1's utility curve as a function of waiting time, w , against an evenly matched opponent. We see that Player 1's utility drops off rapidly from an expected value of zero, with the implication that waiting is costly.

We can also ask "How good does Player 1's detection probability p_1 need to be in order to make up for a given waiting time w ?" Figure 2 shows the adjustment required in this example; for waiting times longer than five periods, even perfect detection does not achieve parity.

Figure 7. Player 1's required detection probability, p_1 , required to achieve $V^w\langle 0,0,0\rangle=0$ as a function of waiting time, w . Player 1's required capability increases rapidly and, because p_1 may never be greater than 1, parity is unachievable after $w=9$.

The lesson of Figures 7 and 8 is that waiting times are costly and adversely affect one's prospects in cyber conflict.

B. DETERRENCE

In the preceding subsection, we advise belligerents in cyber conflict to develop and execute their attacks quickly—a stance that is incompatible with the notion of “crisis stability” (Kent & Thayler, 1989) of classical deterrence theory. Can deterrence in cyber conflict be achieved and, if so, how? Several scholars ask this question, notably (Stern, 2011). In this paper, we consider one aspect of cyber deterrence.

1. A Short Review of Strike Stability

The concept of strike stability was developed during the Cold War to understand which sets of circumstances would lead to nuclear conflict. The original papers describe the development and application of this theory to nuclear arms. Kent and Thayler (1989) describe a game that has many similarities with the one described herein; two players are faced with the decision of “attacking” or “not attacking.” They make this decision by weighing the benefits of going “first” or “second,” with the assumption that the other player will surely retaliate with whatever force he has left. The closer the ratio of costs of going second to going first is to one, the more stable the system is because the decision maker is indifferent to striking first or striking second and may be deterred. Low values of strike stability indicate a large disadvantage to attacking second and therefore lead to

instability. Deterrence requires both sides to choose non-action (“Wait” in our model) at each decision epoch.

2. Strike Stability for Cyber Conflict

The analysis of Section 4 shows that if a player has the ability to attack, he eventually does with certainty. This means that cyber conflict with perfect information and a single exploit is deterrence unstable. Intuitively, this is because there is no second strike. Theorem 1 is sufficient to demonstrate that the single-attack case is deterrence unstable; the first player to attack receives the reward of his development to date, and the nonattacking player is left with an empty arsenal.

Other considerations may provide some degree of deterrence in reality. For example, military, economic, or diplomatic consequences, or large cyber munition arsenals, may provide some guarantee of a second strike. Such guarantees, while important to deterrence, are outside the bounds of our current work. Nevertheless, without these external guarantees, deterrence in cyber conflict does not exist.

VI. CONCLUSIONS AND FUTURE WORK

We have developed and exercised a limited, stylized model. Real situations, of course, have many differences from the idealized mathematics; the utility of this work is to define the cyber conflict problem with perfect information. Additionally, we:

- Demonstrate a framework for analyzing the problem;
- Demonstrate that in cyber conflict, idle wait times are damaging, and provide a means to calculate their disutility; and
- Show implications for deterrence in cyber conflict.

This paper considered a single attack in discrete time with perfect information—three idealizations that help us begin to tackle the problem of cyber conflict. Of these three, the perfect information assumption appears to be the richest area to explore in the future, and with this exploration come considerations of credibility, reputations, and risk taking. Also ripe for future work is consideration of cases with multiple attacks.

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

- Fudenberg, D., & Tirole, J. (1991). *Game theory*. Cambridge: MIT Press.
- JASON. (2010). *The science of cyber security* (Technical Report: JSR-10-102). McLean: MITRE Corporation.
- Kent, G., & Thayler, D. (1989). *First strike stability: A methodology for evaluating strategic forces*. Santa Monica: RAND.
- Lye, K., & Wing, J. (2002). Game strategies in network security. *International Journal of Information Security*, 4, 71-86.
- National Security Council. (2010). *The comprehensive national cybersecurity initiative*. Retrieved from <http://www.whitehouse.gov/cybersecurity/comprehensive-national-cybersecurity-initiative>
- Shen, D., Chen, G., Blasch, E., & Tadda, G. (2007, 9-13 April). *A Markov game theoretic approach for cyber situational awareness*. SPIE's Defense and Security Symposium, Orlando, FL.
- Shiva, S., Dasgupta, D., & Wu, O. (2010). *Game theoretic approaches to protect cyberspace* (Technical Report: CS-10-001). Memphis: University of Memphis.
- Sterner, E. (2011). Deterrence in cyberspace. *Strategic Studies Quarterly*, Spring, 68-80.
- Sternstein, A. (2011). *The White House's 2012 budget devotes a greater percentage of IT funds to cybersecurity*. Retrieved from http://www.nextgov.com/nextgov/ng_20110216_3295
- Thie, P. (1983). *Markov decision processes*. Lexington: UMAP expository monograph series.
- Washburn, A., & Kress, M. (2009). *Combat modeling*. New York, NY: Springer Press.

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California
3. Research Sponsored Programs Office, Code 41
Naval Postgraduate School
Monterey, California
4. Richard Mastowski (Technical Editor)1
Graduate School of Operational and Information Sciences (GSOIS)
Naval Postgraduate School
Monterey, California
5. CDR Harrison Schramm, USN1
Graduate School of Operational and Information Sciences (GSOIS)
Naval Postgraduate School
Monterey, California
6. Assistant Professor David Alderson1
Graduate School of Operational and Information Sciences (GSOIS)
Naval Postgraduate School
Monterey, California
7. Associate Professor W. Matt Carlyle1
Graduate School of Operational and Information Sciences (GSOIS)
Naval Postgraduate School
Monterey, California
8. Assistant Professor Nedialko Dimitrov1
Graduate School of Operational and Information Sciences (GSOIS)
Naval Postgraduate School
Monterey, California
9. RADM Andrew Singer, USN (Ret.)1
Information Dominance Center for Excellence
Naval Postgraduate School
Monterey, California

10. Dr. Alexander Kott.....1
Network Science Division
Army Research Laboratory
Adelphi, Maryland
11. CAPT Jeff Kline, USN (Ret.)1
Graduate School of Operational and Information Sciences (GSOIS)
Naval Postgraduate School
Monterey, California
12. CDR Doug Burton, USN1
Graduate School of Operational and Information Sciences (GSOIS)
Naval Postgraduate School
Monterey, California