



NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

THE ROLE OF DECEPTIVE DEFENSE IN CYBER STRATEGY

by

Mathieu Couillard

June 2023

Co-Advisors:

Britta Hale

Jon Lindsay

Second Reader:

John J. Arquilla

Approved for public release. Distribution is unlimited.

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the 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 Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC, 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE June 2023	3. REPORT TYPE AND DATES COVERED Master's thesis	
4. TITLE AND SUBTITLE THE ROLE OF DECEPTIVE DEFENSE IN CYBER STRATEGY			5. FUNDING NUMBERS	
6. AUTHOR(S) Mathieu Couillard				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release. Distribution is unlimited.			12b. DISTRIBUTION CODE A	
13. ABSTRACT (maximum 200 words) This thesis examines the role of deceptive defense in cyber strategy and contributes a new practical concept to aid in its implementation. In <i>The Art of War</i>, Sun Tzu stated, "All warfare is based on deception." Cyber conflict should pose no exception. The ability to operate anonymously in the cyber domain enables attackers to operate with some degree of impunity, as the attribution of their identities is a complex, costly, and imperfect process. However, defenders can utilize cyber deception to improve network security and enhance both forensic and potential retaliatory cyber operations. This work introduces Deceptive Resistance to Adversary Cyber Operations (DRACO), a concept that simulates active network resources while remaining outside of a secure network perimeter. In this respect it differs from traditional "honeypot" deployments that are either within the network or completely detached. DRACO offers a "deception as a service" architecture, requiring limited integration effort and offering cloud hosting potential. During experimentation on a live network, DRACO successfully redirected attacks that were intended for authentic endpoints and revealed insights pertaining to the attack patterns on active and inactive addresses within a network. Additionally, DRACO presents an opportunity to engage with the adversary, which in turn can enable the identification of an ongoing attack, adversary attribution, and assist in the preparation of retaliatory options.				
14. SUBJECT TERMS cyber, strategy, intelligence, cyberwar, warfare, information, dominance			15. NUMBER OF PAGES 117	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UU	

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release. Distribution is unlimited.

THE ROLE OF DECEPTIVE DEFENSE IN CYBER STRATEGY

Mathieu Couillard
Lieutenant-Colonel, Canadian Army
BCE, Université Laval, 2008

Submitted in partial fulfillment of the
requirements for the degrees of

**MASTER OF SCIENCE IN DEFENSE ANALYSIS
(IRREGULAR WARFARE)**

and

MASTER OF SCIENCE IN COMPUTER SCIENCE

from the

**NAVAL POSTGRADUATE SCHOOL
June 2023**

Approved by: Britta Hale
Co-Advisor

Jon Lindsay
Co-Advisor

John J. Arquilla
Second Reader

Carter Malkasian
Chair, Department of Defense Analysis

Gurminder Singh
Chair, Department of Computer Science

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

This thesis examines the role of deceptive defense in cyber strategy and contributes a new practical concept to aid in its implementation. In *The Art of War*, Sun Tzu stated, “All warfare is based on deception.” Cyber conflict should pose no exception. The ability to operate anonymously in the cyber domain enables attackers to operate with some degree of impunity, as the attribution of their identities is a complex, costly, and imperfect process. However, defenders can utilize cyber deception to improve network security and enhance both forensic and potential retaliatory cyber operations. This work introduces Deceptive Resistance to Adversary Cyber Operations (DRACO), a concept that simulates active network resources while remaining outside of a secure network perimeter. In this respect it differs from traditional “honeypot” deployments that are either within the network or completely detached. DRACO offers a “deception as a service” architecture, requiring limited integration effort and offering cloud hosting potential. During experimentation on a live network, DRACO successfully redirected attacks that were intended for authentic endpoints and revealed insights pertaining to the attack patterns on active and inactive addresses within a network. Additionally, DRACO presents an opportunity to engage with the adversary, which in turn can enable the identification of an ongoing attack, adversary attribution, and assist in the preparation of retaliatory options.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
II.	STRATEGIC CONTEXT	5
A.	DECEPTION IN WARFARE.....	5
1.	A Brief History of Deception.....	5
2.	The Theory of Deception	7
B.	CYBER CONFLICT	9
1.	Information Dominance through Other Means	9
2.	Conceptions of the Cyber Domain.....	10
C.	DECEPTIVE CYBER DEFENSE.....	12
1.	Deception in Cyberspace	12
2.	Cyber Deception Research	14
D.	CONCLUSION	16
III.	THE ROLE OF DECEPTIVE DEFENSE IN CYBER STRATEGY: LESSONS FROM DECOY VESSELS OF THE GREAT WAR	17
A.	CASE STUDY: Q-SHIPS	18
1.	Background: The U-Boat Campaign	18
2.	Q-ships	19
3.	Cyber Q-ships?.....	22
B.	EFFECT DESCRIPTIONS.....	23
1.	Identification	23
2.	Misdirection.....	24
3.	Retaliation.....	25
4.	Summary.....	26
C.	NAVIGATING THE SECURITY-DECEPTION TRADE-OFF	27
D.	CONCLUSION	29
IV.	DRACO: SECURE AND COMPELLING DECEPTION	31
A.	INTRODUCTION.....	31
B.	DECEPTIVE RESISTANCE TO ADVERSARY CYBER OPERATIONS	31
1.	Architecture.....	31
2.	Insider Threat Variant	33
C.	ALTERNATIVE HONEYPOT ARCHITECTURES	35
1.	Operational Scenarios	35

2.	External Honeypot	37
3.	DMZ Honeypot.....	38
4.	Intranet Honeypot.....	39
5.	DRACO	40
D.	CONCLUSION	41
V.	EXPERIMENTATION	43
A.	METHODOLOGY	43
1.	Production Network and Collaboration with the U.S. Naval Postgraduate School.....	43
2.	DRACO: Deception Network Architecture.....	45
B.	EXPERIMENT 1: SIMULATING AN ENDPOINT WITH DRACO	48
1.	Design	49
2.	Analysis and Results	52
3.	Conclusion	66
C.	EXPERIMENT 2: COMPARISON OF ATTACK PATTERNS ON ACTIVE AND INACTIVE ENDPOINTS	67
1.	Design	67
2.	Analysis and Results	68
3.	Summary.....	76
D.	LIMITATIONS.....	76
E.	CONCLUSION	77
VI.	CONCLUSION	79
A.	SUMMARY	79
B.	FUTURE WORK.....	80
	LIST OF REFERENCES	83
	INITIAL DISTRIBUTION LIST	95

LIST OF FIGURES

Figure 1.	Lockheed Martin’s cyber kill chain	26
Figure 2.	DRACO architecture.....	32
Figure 3.	DRACO–Insider Threat Variant (DRACO-ITV)	34
Figure 4.	Typical network security zones with corresponding options for honeypot deployment.....	35
Figure 5.	Cyber threat actors and their primary motivation	37
Figure 6.	High-level architecture of the U.S. Naval Postgraduate School’s network entry point	44
Figure 7.	Deception network architecture, a) with DRACO and b) without.....	46
Figure 8.	Control Group network architecture	49
Figure 9.	Architecture of the deception network, DRACO.....	50
Figure 10.	Line plot, daily HTTP unique source addresses	53
Figure 11.	Box plot, daily HTTP unique source addresses	53
Figure 12.	Line plot, daily HTTP requests	54
Figure 13.	Box plot, daily HTTP requests.....	54
Figure 14.	Line plot, daily SSH unique source addresses	55
Figure 15.	Box plot, daily SSH unique source addresses.....	56
Figure 16.	Line plot, daily SSH connections.....	57
Figure 17.	Box plot, daily SSH connections	57
Figure 18.	ACF plot, daily HTTP unique source addresses	59
Figure 19.	ACF plot, daily HTTP requests	59
Figure 20.	ACF plot, daily SSH unique source addresses.....	60
Figure 21.	ACF plot, daily SSH connections	60

Figure 22.	Top countries by proportion of estimated geolocation of source addresses on the DRACO server.....	65
Figure 23.	Proportion of unique addresses by hour of the day for each server.....	66
Figure 24.	Architecture of the active and inactive simulated endpoints	68
Figure 25.	Total new connections, unique source addresses, and commands on the active and inactive simulated endpoints.....	69
Figure 26.	Line plot, daily SSH connections.....	70
Figure 27.	Line plot, daily SSH unique source addresses	70
Figure 28.	Daily commands issued by attacks on the simulated SSH shell.....	71
Figure 29.	Proportion of unique source addresses by time of the day (local time)....	72
Figure 30.	Top countries by proportion of estimated geolocation of source addresses	73

LIST OF TABLES

Table 1.	Effects of deceptive cyber defense	30
Table 2.	Conceptual evaluation of DRACO and typical honeypot location options.....	41
Table 3.	Characteristics of the addresses assigned to the DRACO server and the control servers	52
Table 4.	Results of Shapiro-Wilk test for normality on daily statistics	62
Table 5.	Results of distribution tests on daily statistics	64
Table 6.	Sources IP address by estimated geolocation	73
Table 7.	Results Shapiro-Wilk test results for non-normality	74

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

APT	Advanced Persistent Threat
ARPANET	Advanced Research Projects Agency Network
AWS	Amazon Web Services
CCB	Configuration Control Board
CISA	Cybersecurity and Infrastructure Security Agency
DID	Different in Difference
DMZ	Demilitarized Zone
DRACO	Deceptive Resistance to Adversarial Cyber Operations
DRACO-ITV	Deceptive Resistance to Adversarial Cyber Operations–Insider Threat Variant
EC2	Elastic Compute Cloud
GB	Gigabyte
HTTP	Hypertext Transfer Protocol
ICS	Industrial Control System
ISP	Internet Service Provider
NAT	Network Address Translation
NIST	National Institute of Standards and Technology
SSH	Secure Shell
VPN	Virtual Private Network

THIS PAGE INTENTIONALLY LEFT BLANK

EXECUTIVE SUMMARY

This work examines the role of deceptive cyber defense through qualitative and quantitative analysis, including the development of an operational concept to aid in its implementation. We study key challenges to offensive and defensive operations and illustrate how deceptive cyber defense can address them. We then present a concept to aid in its operationalization: Defensive Resistance to Adversary Cyber Operations (DRACO). A conceptual evaluation as well as experimentation on live networks suggest that DRACO offers significant advantages relative to typical honeypot architectures and can provide valuable insights to network defenders.

The Role of Deceptive Defense in Cyber Strategy

We chose First World War Q-ships as a conceptual case illustrating deceptive defense that can identify a hidden attacker, misdirect an attack, and allow for retaliation. Q-ships were decoy vessels that lured stealthy U-boats to a position of vulnerability when surfaced, providing *identification* of threats. In the same way, an effective defensive cyber deception may compel the adversary to an engagement on the defender's terms. By offering a target of opportunity, both Q-ships and deceptive cyber defense *misdirect* adversaries, drawing them away from their intended targets. The ultimate purpose of Q-ships was to destroy or disable U-boats in *retaliation* for an initial engagement. Current cyber response practice is analogous to launching an anti-submarine ship to search for a U-boat that had sunk a defenseless merchant ship. Yet, by leveraging deception measures, cyber response can more closely mirror the retaliation measures of Q-ships, where the attacker is itself targeted while performing malicious activities.

DRACO: Secure and Compelling Deception

DRACO addresses key challenges associated with deception technology by introducing an architecture which presents an effective deception to cyber threat actors while reducing the integration burden for the end user. Use of network address translation or "spoofing" allows deceptive resources to remain separate from the end user's secure network, while providing responses that appear authentic to cyber threat actors. When

designated attack traffic (e.g., an unauthorized connection attempt) is received at the edge of the network, the router forwards it to a network address translation (NAT) service, which modifies the destination address to match a deception network, with a similar translation taking place for response traffic. This “deception as a service” approach enables hosting of the deceptive network component on the cloud for easy deployment, separation from operational network infrastructure, and enhanced availability.

Our evaluation of typical honeypot deployment locations and the DRACO concept show varying levels of performance that ultimately favor DRACO. External honeypots offer high ease of integration and maintenance and a very low security risk due to their isolation from an internal network; however, their deceptive effectiveness may be lower due to lack of activity within the internet-facing zone. Intranet honeypots, while offering very high deceptive effectiveness, come with a low ease of integration and maintenance. However, intranet honeypots also offer a potential foothold to cyber threat actors as they are operating within the system—representing a greater security risk. Demilitarized zone (DMZ) honeypots attempt to strike a balance by providing a medium ease of integration and maintenance, deceptive effectiveness, and a lower security risk than in intranet hosting, due to the separation of the DMZ from the most sensitive network resources. DRACO, on the other hand, achieves high ease of integration and maintenance, strong deceptive effectiveness, and a very low security risk (see Table 1 for comparison).

Table 1. Conceptual evaluation of DRACO and typical honeypot location option

Criterion	External Honeypot	DMZ Honeypot	Intranet Honeypot	DRACO
Ease of integration and maintenance	High	Medium	Low	High
Deceptive Effectiveness	Low	Medium	Very High	High
Security Risk	Very Low	Low	Medium	Very Low

Experimentation

Validation of DRACO's ability to simulate an authentic endpoint securely was shown via experiments on live networks. The first experiment evaluated DRACO's ability to simulate an active endpoint by comparing attack patterns on a simulated endpoint and on typical research honeypots. We inferred this ability by comparing the volume of attacks and unique source addresses between a control group and an experimental server which employs DRACO. During the experiment, the DRACO server logged significantly more attacks and unique source addresses than any control server. Thus, the first experiment provided an initial indication that DRACO may effectively redirect attacks that were intended for an authentic endpoint.

For the second experiment, we employed DRACO to study attack patterns on an active and an inactive address within a subnet. In this case, an active endpoint was one which hosted at least one service for legitimate users of the network (i.e., an attacker would receive a response on more ports than just the port [s] corresponding to the deception). Conversely, an inactive endpoint was one which responded only on the port(s) corresponding to the deception. Our results indicated that there were similarities between the attack patterns on inactive endpoints and active endpoints. However, attackers were far more engaged in their attempts to compromise the active endpoint, based on the number of commands issued on the simulated session. This demonstrates DRACO's potential for cybersecurity research and reinforces the need to study attack patterns on active endpoints.

Summary

This study identified three important aims of effective deceptive cyber defense: *identification*, *misdirection*, and *retaliation*. To achieve these effects, we presented a honeypot placement and integration approach (DRACO). We conceptually evaluated DRACO and typical honeypot placement options against two cyber actor profiles and quantitatively evaluated it through experimentation on the U.S. Naval Postgraduate School network. Our work provides an initial indication that DRACO can 1) effectively redirect attacks that were intended for an authentic endpoint, and 2) be useful for targeted research regarding specific threats to the production network.

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGMENTS

I arrived at the U.S. Naval Postgraduate School in May 2021 with the clear intent to contribute to the operationalization of deceptive cyber defense. I thank my former colleagues for opening my eyes in this regard.

While my mission was clear, the path to achieve it was far more ambiguous and challenging than I could have anticipated. I am deeply grateful to my thesis co-advisors, Dr. Britta Hale and Dr. Jon Lindsay, for their direction and unwavering patience along this journey of discovery. Additionally, it was a privilege to draw upon the wisdom and guidance of Dr. John Arquilla, my second reader. Special thanks to Mr. George Lober who acted interchangeably as writing coach and therapist throughout the process.

The chief information security officer, Mr. Chris Angelopoulos, and the information systems security manager, Mr. Gerald Holland, allowed a (Canadian) student to deploy an experimental technology on the U.S. Naval Postgraduate School's network. I'm very thankful for their trust and the support from their outstanding team. Thank you to Mr. Daryl Williams, Mr. Mike Williams, Mr. Edgar Mendoza, and Mr. Joseph Martineau, without whom this project would not have been possible.

My peers and faculty in the Defense Analysis and Computer Science departments represent some of the finest non-commissioned officers, officers, and public servants that I have ever worked with. My family and I will forever cherish the bonds that we have forged in all services of the U.S. military and across the world thanks to the U.S. Naval Postgraduate School.

Most of all, I wish to thank my family. My wife juggled her own very demanding job while doing far more than her share for our family. I'm grateful to my sons for their unrelenting encouragement to finish writing my "chapter book" so that we could spend more time on the beach. Merci pour votre appui, les grands. Maintenant, allons à la plage!

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

Sun Tzu stated, “All warfare is based on deception,” and cyber conflict poses no exception to this aphorism. Cyber threat actors exploit human and technical vulnerabilities to penetrate even the most secure and sensitive networks. Once inside, attackers gain a foothold and silently pivot across the network, employing yet more deceptive tradecraft to cloak themselves beneath the noise of network traffic. Notably, a 2022 report from IBM found that on average, network breaches were only detected 212 days after the initial compromise.¹ In many cases, the ambiguity of the cyber domain enables the attacker to thwart punishment, as the ultimate attribution of cyber attackers’ identities is a costly and imperfect process.

Numerous cyber attacks in recent history illustrate this deceptive pattern. For instance, hackers gained unauthorized access to the internal network of SolarWinds (a system management software developer) in September 2019.² Six months later, the attackers compromised up to 18,000 of SolarWinds’ customers by delivering malicious code through program updates. The victims included leading companies such as Microsoft and Intel and key government agencies such as the Cybersecurity and Infrastructure Security Agency (CISA). The attack was detected by the cybersecurity firm FireEye, but not until December 2020.³ The U.S. government eventually attributed the SolarWinds hack to Russia’s Foreign Intelligence Service (SVR) in an April 2021 press release, which also announced retaliatory economic sanctions.⁴ More than 18 months had passed since the initial compromise.

¹ “Cost of a Data Breach,” IBM, 2022, 14, <https://www.ibm.com/reports/data-breach>.

² Dina Temple-Raston, “A ‘Worst Nightmare’ Cyberattack: The Untold Story Of The SolarWinds Hack,” *NPR*, April 16, 2021, <https://www.npr.org/2021/04/16/985439655/a-worst-nightmare-cyberattack-the-untold-story-of-the-solarwinds-hack>.

³ “Highly Evasive Attacker Leverages SolarWinds Supply Chain to Compromise Multiple Global Victims with SUNBURST Backdoor,” Mandiant, December 13, 2020, <https://www.mandiant.com/resources/blog/evasive-attacker-leverages-solarwinds-supply-chain-compromises-with-sunburst-backdoor>.

⁴ “Imposing Costs for Harmful Foreign Activities by the Russian Government,” The White House, April 15, 2021, <https://www.whitehouse.gov/briefing-room/statements-releases/2021/04/15/fact-sheet-imposing-costs-for-harmful-foreign-activities-by-the-russian-government/>.

Defense leaders are keenly aware of the magnitude of the cyber defense challenge. Early data from the Russia-Ukraine conflict in 2022 demonstrate that the coordination of cyber and kinetic effects to achieve military objectives, once deemed a futurist concept, is now reality.⁵ In response to the growing threat, the Pentagon's cyber budget is expected to reach \$11.2B in fiscal year 2023.⁶ From a tactical standpoint, the Department of Defense is implementing a "Zero Trust" strategy to strengthen inevitably penetrable networks from within rather than rely on a leakproof perimeter defense.⁷ The U.S. Cyber Command deployed "hunt forward" cyber defense teams to Ukraine in the lead-up to the Ukraine-Russia conflict.⁸ The role of these teams was not to strengthen the perimeter defense of Ukrainian networks but rather to detect ongoing malicious cyber activity within these networks. This paradigm shift represents both acceptance of the reality that attackers *will* gain access to the network and a desire to engage with the adversary.

In the same way that deception can enable the cyber offense, deception can elevate current efforts to deter or detect cyber attacks. Defenders utilize cyber deception through many techniques, including one widely known as the *honeypot*. Through deceptive cyber defense, defenders seek to draw out the adversary to alter the balance of information, revealing the adversary's intent and tactics. This approach also presents a potential opportunity to access and counter-attack adversarial systems ("back-hacking"). Deceptive cyber defense can thus form a bridge between defensive and offensive cyber operations and forge an information advantage. While this stratagem shows great potential, the adoption of honeypots may be limited by challenges in creating an effective deception and

⁵ Microsoft Digital Security Unit, "Special Report: An Overview of Russia's Cyberattack Activity in Ukraine," April 27, 2022.

⁶ Colin Demarest, "Pentagon Seeks \$11.2 Billion for Cyber in FY23 Budget Request," C4ISRNet, March 28, 2022, 2, <https://www.c4isrnet.com/cyber/2022/03/28/pentagon-seeks-112-billion-for-cyber-in-fy23-budget-request/>.

⁷ C. Todd Lopez, "DOD Releases Path to Cyber Security Through Zero Trust Architecture," U.S. Department of Defense, November 38, 2022, <https://www.defense.gov/News/News-Stories/Article/Article/3229211/dod-releases-path-to-cyber-security-through-zero-trust-architecture/>.

⁸ Cyber National Mission Force Public Affairs, "Before the Invasion: Hunt Forward Operations in Ukraine," U.S. Cyber Command, November 28, 2022, <https://www.cybercom.mil/Media/News/Article/3229136/before-the-invasion-hunt-forward-operations-in-ukraine/>.

the perceived security risks of hosting honeypots internally to an operational network.⁹ This raises the question: *to what degree and under what conditions can deceptive cyber defense be used to achieve an information advantage?*

This thesis approaches the research through qualitative and quantitative analysis, including the development of a concept to aid in the operationalization of deceptive cyber defense. We examine key challenges with offensive and defensive operations and illustrate how deceptive cyber defense can address them. Having shown the potential effect of this technology on cyber conflict, we present a concept to aid in its operationalization: Deceptive Resistance to Adversary Cyber Operations (DRACO). A conceptual evaluation and experimentation on live networks suggest that DRACO offers significant advantages relative to typical honeypot architectures and can provide valuable insights to network defenders.

Chapter II presents the strategic context and recent advancements in deception technology. In Chapter III, we consider the role of deceptive cyber defense in cyber strategy, broadly identifying the ways in which deception can enable offensive and defensive cyber operations. We present the DRACO concept in Chapter IV, and in Chapter V conduct experiments to evaluate it. Chapter VI concludes with a summary of our findings and suggested pathways for future research.

⁹ Iyatiti Mokube and Michele Adams, “Honeypots: Concepts, Approaches, and Challenges,” in *Proceedings of the 45th Annual Southeast Regional Conference on – ACM-SE 45* (New York, NY: ACM Press, 2007), 321, <https://doi.org/10.1145/1233341.1233399>.

THIS PAGE INTENTIONALLY LEFT BLANK

II. STRATEGIC CONTEXT

This thesis draws upon two important bodies of knowledge: deception and cyber conflict. The area of overlap between these concepts is diverse and rich in offensive and defensive opportunities. The following sections provide a basis for the study of deceptive cyber defense and illuminate the research gaps that it can address. First, we present a historical overview of deception in warfare, including the theory of deception. Second, we introduce the two major theories of cyber conflict and identify their respective limitations. Finally, we bridge these concepts in a review of recent research on deceptive cyber defense.

A. DECEPTION IN WARFARE

1. A Brief History of Deception

The historian Jon Latimer wrote: “Deception on the battlefield is surely as old as warfare itself.”¹⁰ In the opening pages of *Deception in War*, Latimer describes some of the earliest known examples of military deception, from the Hittite ruse of sending “deserters” to lead Pharaoh Ramses II of Egypt to an ambush in 1294 BC, to Gideon’s use of torches to simulate a large force in a daring overnight attack against the Midianites around 800 BC.¹¹ Latimer reports that in the first century AD, Roman general and senator Julius Sextus Frontinus dedicated the second of his two volume study of warfare on the art of deception (*Stratagems*).¹² Deception was commonly employed throughout the ancient world, but is perhaps most closely linked to the teachings of Sun Tzu. *The Art of War*, thought to be written around the fifth century BC, holds numerous passages that reinforce the importance of deception and surprise:

All warfare is based on deception.

Therefore, when capable, feign incapacity; when active, inactivity.

¹⁰ Jon Latimer, *Deception in War: The Art of the Bluff, the Value of Deceit, and the Most Thrilling Episodes of Cunning in Military History, from the Trojan Horse to the Gulf War*, First Edition (New York, NY: Abrams, 2003).

¹¹ Latimer, 6.

¹² Latimer, 9.

When near, make it appear that you are far away; when far away, that you are near.¹³

Pursuant to the principles laid out by Sun Tzu, *The Thirty-Six Stratagems* illustrates various deceptive tactics employed in politics and war.¹⁴ This work has been attributed to Sun Tzu or Zhuge Liang but is generally presumed to be an *oeuvre collective* passed on in written and oral history. *The Thirty-Six Stratagems* was only widely distributed after it was published by the Chinese Communist Party in 1961.¹⁵ The stratagems are presented in six categories that describe the context for their use (e.g., whether in a winning or losing situation). *The Thirty-Six Stratagems* is still relevant today because of how simply it conveys a stratagem in relation to an objective.

More recently, the Second World War yields many examples of military deception. The most famous of these is Operation FORTITUDE, the Allied deception plan surrounding the invasion of Normandy in 1944. Lieutenant-Colonel Roger Hesketh of the British Army, director of FORTITUDE, wrote the official history of the operation in 1945.¹⁶ Hesketh described the planning process, beginning with the understanding of the strategic context, then with the ideation of narratives that would reinforce the adversary's biases, and lastly with the means for an effective deception.¹⁷ The allies used physical deception (including "phantom armies"), fake wireless traffic, double agents, and diplomatic misinformation to convince the German General Staff that the Allies intended to land either in Norway or the Pas-de-Calais. Intercepted German communications have shown that FORTITUDE, likely the most elaborate deception operation ever conducted, was successful.¹⁸

¹³ Sun Tzu (Translated by Samuel B. Griffith), *The Art of War* (New York, NY: Clarendon Press, 1963), 66.

¹⁴ Peter Taylor, *The Thirty-Six Stratagems: A Modern Interpretation of a Strategy Classic* (Oxford, UK: Infinite Ideas, 2013), 9.

¹⁵ Taylor, 10.

¹⁶ Though Hesketh wrote this in 1945, it was only published in 2002. See Roger Hesketh, *Fortitude: The D-Day Deception Campaign* (New York, NY: Abrams, 2002).

¹⁷ Hesketh, 364–70.

¹⁸ Thaddeus Holt, *The Deceivers: Allied Military Deception in the Second World War* (New York, NY: Scribner, 2014), 565–91.

2. The Theory of Deception

The rich body of data on military deception enabled scholar Barton Whaley to formulate a general theory of deception in 1982. Whaley posited that the most effective forms of deception are those which dissimulate a true intent while simulating an alternate reality, as opposed to those which “combine *dazzling* with *decoying* to achieve only a mere razzle-dazzle effect.”¹⁹ In other words, Whaley’s theory suggests that a successful deception will be specifically crafted to lead an opponent into believing an alternate reality, rather than be generically constructed to distract them. In later writing, Whaley analyzed the general benefits of deception and ultimately equated military deception to an effective method of generating surprise.²⁰ Based on a database of 232 military engagements between 1914 and 1973, Whaley verified the three following “consequences of deception”:

- the more intense the deception (i.e., the more types of ruses used), the greater the surprise—but only up to a point;
- the greater the surprise, the smaller the force ratio required;
- the greater the surprise, the fewer the casualties.²¹

Whaley has shown that deception has a success rate of 80% when attempted, yet this stratagem has not been adequately codified in many Western militaries and is rarely the subject of academic analysis.²² Whaley writes that deception has been “intermittently and idiosyncratically practiced” and is most often the result of “exceptionally imaginative leaders rather than a routine textbook practice.”²³ Scholar and retired brigadier general Walter Jajko noted that Operation BODYGUARD, the strategic allied deception plan during the Second World War, was perhaps the only U.S. case study of strategic

¹⁹ Barton Whaley, “Toward a General Theory of Deception,” *Journal of Strategic Studies* 5, no. 1 (March 1982): 187, <https://doi.org/10.1080/01402398208437106>.

²⁰ Barton Whaley, “The One Percent Solution: Costs and Benefits of Military Deception,” in *Information Strategy and Warfare* (London: Routledge, 2007), 127.

²¹ Whaley, 128.

²² Whaley, 127.

²³ Barton Whaley, *Stratagem: Deception and Surprise in War* (Norwood, MA: Artech House, 2007), 2, <https://us.artechhouse.com/Stratagem-Deception-and-Surprise-in-War-P1754.aspx>.

deception.²⁴ Jajko suggests that this may be a reflection of the American perception that “deception as a weapon is seen as intimidating, unacceptable, and even evil.”²⁵

The influence of Carl Von Clausewitz may also contribute to the perceived Western reluctance to embrace military deception. Professor Michael I. Handel of the U.S. Naval War College authored a comparative analysis of the works of Sun Tzu and Clausewitz.²⁶ According to Handel, Clausewitz saw deception “as the last resort of the weak and the desperate,” especially at the operational and strategic levels where he deemed surprise an unlikely possibility.²⁷ Indeed, Clausewitz wrote that “Plans and orders issued for appearances only ... have as a rule so little strategic value that they ... should not be considered as a significant independent field of action at the disposal of the commander.”²⁸ Clausewitz also stated that “it is rare that one state surprises another, either by an attack or by preparations for war,” due to the slow pace of strategic movements.²⁹ Handel argues that Clausewitz’s “lack of interest” in deception and surprise was a product of his era, prior to the Industrial Revolution, “which led to previously unimaginable improvements in mobility.”³⁰ Ultimately, however, the historical data on military deception speaks for itself: the benefits vastly outweigh the costs.³¹ There is no reason to believe that in the modern operating environment, rich with disinformation and manipulation, this statistic should be any different.

²⁴ Walter Jajko, “Deception: Appeal for Acceptance; Discourse on Doctrine; Preface to Planning,” *Comparative Strategy* 21, no. 5 (October 1, 2002): 351–63, <https://doi.org/10.1080/01495930290043092>.

²⁵ Jajko, 351.

²⁶ Michael I. Handel, “Sun Tzu and Clausewitz: The Art of War and on War Compared:” (Fort Belvoir, VA: Defense Technical Information Center, January 1, 1991), <https://doi.org/10.21236/ADA239084>.

²⁷ Handel, 43.

²⁸ Clausewitz, Carl von (Edited and translated by Michael Eliot Howard Peter Paret), *On War* (Princeton, NJ: Princeton University Press, 1989), 202–3.

²⁹ Clausewitz, 198–99.

³⁰ Handel, “Sun Tzu and Clausewitz,” 45.

³¹ This is Whaley’s thesis, supported by a cost/benefit analysis in: Whaley, “The One Percent Solution.”

B. CYBER CONFLICT

1. Information Dominance through Other Means

In 1994, John Arquilla described information dominance as “knowing everything about an adversary while keeping the adversary from knowing much about oneself.” The strength of this simple definition is that it extends beyond information technology and centers on *knowledge*. In the age of big data, it seems easy to forget that an information advantage is only meaningful if it enhances one’s knowledge or degrades that of the adversary.³² Conversely, militaries too often conflate technological superiority or a mass of unmonitored sensors with an information advantage. Information dominance is a concept that transcends technology, as Arquilla illustrated through both Hannibal’s use of signal mirrors in the Punic Wars and the superior intelligence network of the communist forces in the Vietnam War.³³ Nonetheless, this contest for information dominance has been fundamentally altered by the creation of the internet.

The first two nodes of the Advanced Research Projects Agency Network (ARPANET), the predecessor to the internet, were connected between the Network Measurement Center at the University of California, Los Angeles, and the Stanford Research Institute in 1969.³⁴ Robust security was a luxury deemed not only infeasible at the time due to the limited processing power available, but also non-essential given the rush to operationalize the internet. Vinton Cerf, one of the founders of the internet, stated, “We didn’t focus on how you could wreck this system intentionally.”³⁵ By 1989, the ARPANET consisted of 60,000 nodes and the academic Peter J. Denning claimed that it had a “more profound effect on the world than that of the more spectacular and expensive

³² For an appreciation of the difference between information and knowledge, consult Gene Bellinger, Durval Castro, and Anthony Mills. “Data, Information, Knowledge, and Wisdom,” *Computer Science*, 2011.

³³ John Arquilla, “The Strategic Implications of Information Dominance,” *Strategic Review* 22 (1994): 22–34.

³⁴ Barry M. Leiner et al., “A Brief History of the Internet,” *ACM SIGCOMM Computer Communication Review* 39, no. 5 (October 7, 2009): 22–31, <https://doi.org/10.1145/1629607.1629613>.

³⁵ “The Real Story of How the Internet Became so Vulnerable,” *Washington Post* (blog), accessed July 7, 2021, <http://www.washingtonpost.com/sf/business/2015/05/30/net-of-insecurity-part-1/>.

Apollo Missions.”³⁶ The internet is now a behemoth on course to connect 55.7 billion end points by 2025 according to the International Data Corporation.³⁷ The Internet of Things is in our homes, appliances, vehicles, and military systems.³⁸ The internet is thus of paramount and ever-increasing importance to nearly all facets of society. Consequently, state and non-state actors have thus long exploited the foundational insecurity of the internet to wage this latest iteration of the contest for information dominance.

2. Conceptions of the Cyber Domain

The military’s early understanding of the cyber domain was centered on the conduct of war. In 1993, John Arquilla and David Ronfeldt foresaw a revolution in military affairs that would result from the effective application of networked technologies. The authors predicted that the information technology revolution would give rise to *cyberwar*, the conduct of “military operations according to information-related principles.”³⁹ Academics have since debated whether cyberwar would materialize or has in fact already occurred.⁴⁰ Russia’s coordination of land, air, naval, and cyber operations against Georgia in 2008 has been posited as first the instance of cyberwar.⁴¹ However, the relative importance of cyber operations in this conflict has also been questioned.⁴² In the context of the current Russian invasion of Ukraine, academics have once again debated the impact that cyber operations

³⁶ Peter J. Denning, “The Science of Computing: The ARPANET after Twenty Years,” *American Scientist* 77, no. 6 (1989): 532, <https://www.jstor.org/stable/27856002.a>

³⁷ International Data Corporation, “IoT Growth Demands Rethink of Long-Term Storage Strategies, Says IDC,” IDC, accessed August 10, 2021, <https://www.idc.com/getdoc.jsp?containerId=prAP46737220>.

³⁸ For an overview of the Internet of Things, consult Samuel Greengard, *The Internet of Things, Revised and Updated Edition* (Cambridge, MA: MIT Press, 2021).

³⁹ John Arquilla and David Ronfeldt, “Cyberwar Is Coming!,” *Comparative Strategy* 12, no. 2 (1993): 30.

⁴⁰ For opposing viewpoints, see Thomas Rid, *Cyber War Will Not Take Place* (New York, NY: Oxford University Press, 2013); and John Arquilla, “Cyberwar Is Already Upon Us,” *Foreign Policy* (blog), February 27, 2012, <https://foreignpolicy.com/2012/02/27/cyberwar-is-already-upon-us/>.

⁴¹ David Hollis, “Cyberwar Case Study: Georgia 2008,” *Small Wars Journal*, January 6, 2011, 10, <https://smallwarsjournal.com/blog/journal/docs-temp/639-hollis.pdf>.

⁴² Erik Gartzke, “The Myth of Cyberwar: Bringing War in Cyberspace Back Down to Earth,” *International Security* 38, no. 2 (October 2013): 41–73, https://doi.org/10.1162/ISEC_a_00136.

may have on a potential conflict.⁴³ The argument about the scope and scale of cyberwar will go on. The true shortcoming of this discourse, however, is that it does not address the massive proportion of cyber operations that occurs below the threshold of war. In addition, it does not sufficiently explore the dependence of militaries on advanced information systems, the disruption of which can cause serious losses in combat.

Scholars have instead suggested that cyber operations can more broadly be understood as an *intelligence contest*. Erik Gartzke and Jon Lindsay pioneered this notion, noting that the “strategic manipulation of information” in cyberspace lends itself to intelligence operations.⁴⁴ Joshua Rovner observed that “most activities in cyberspace have little to do with the use of force,” and instead represent “a struggle to protect information, while corrupting and stealing information from rivals.”⁴⁵ Lindsay added that cyber operations are a form of “political statecraft” because they “rely on deception.”⁴⁶ This theory thus provides a sound basis to study cyber conflict in the “gray zone,” notoriously exemplified by Russian cyber-enabled information operations and China’s purported cyber espionage campaigns.⁴⁷ Critics, such as Michael Fischerkeller and Richard Harknett, have contended that the strategic concept of intelligence does not lead to an accurate assessment of the “potential strategic significance of states’ cyber behaviors short of armed conflict.”⁴⁸ Nevertheless, no alternative concept has been embraced by academia.

⁴³ Lennart Maschmeyer and Nadiya Kostyuk, “There Is No Cyber ‘Shock and Awe’: Plausible Threats in the Ukrainian Conflict,” *War on the Rocks*, February 8, 2022, <https://warontherocks.com/2022/02/there-is-no-cyber-shock-and-awe-plausible-threats-in-the-ukrainian-conflict/>.

⁴⁴ Erik Gartzke and Jon R. Lindsay, “Weaving Tangled Webs: Offense, Defense, and Deception in Cyberspace,” *Security Studies* 24, no. 2 (April 3, 2015): 316–48, <https://doi.org/10.1080/09636412.2015.1038188>.

⁴⁵ Joshua Rovner, “Cyber War as an Intelligence Contest,” *War on the Rocks*, September 16, 2019, <https://warontherocks.com/2019/09/cyber-war-as-an-intelligence-contest/>.

⁴⁶ Jon R. Lindsay, “Cyber Conflict vs. Cyber Command: Hidden Dangers in the American Military Solution to a Large-Scale Intelligence Problem,” *Intelligence and National Security* 36, no. 2 (February 23, 2021): 260–78, <https://doi.org/10.1080/02684527.2020.1840746>.

⁴⁷ Brandon Valeriano, Benjamin M. Jensen, and Ryan C. Maness, *Cyber Coercion: The Evolving Character of Cyber Power and Strategy* (New York, NY: Oxford University Press, 2018).

⁴⁸ Michael P. Fischerkeller and Richard J. Harknett, “Cyber Persistence Theory, Intelligence Contests and Strategic Competition” (Institute for Defense Analyses, June 1, 2020), <https://apps.dtic.mil/sti/pdfs/AD1118679.pdf>.

Neither the cyberwar nor intelligence-contest theories provide a complete understanding of the cyber domain. Beyond “Cyber Pearl Harbor” or nuclear strike analogies, the application of cyber operations to destroy or disrupt objectives (for purposes other than intelligence) is entirely possible and strategically impactful.⁴⁹ Yet, most cyber operations today are conducted below the threshold of war with an aim to alter the balance of information between actors. Tergiversation between these two models has led to sub-optimal outcomes from both perspectives. From an intelligence contest perspective, Lindsay has identified a mismatch between the tasks of intelligence and cybersecurity that fall upon U.S. Cyber Command, which is “emphatically not an intelligence organization,” and the emphasis the Command places on warfighting concepts.⁵⁰ From a cyberwar perspective, Arquilla has been critical of the excessive compartmentalization of offensive cyber operations in the U.S. military—that disallows cross-fertilization with insights from defensive concepts—resulting from the influence of the intelligence community from which they were born.⁵¹ The evidence suggests that a unifying theory of the cyber domain is needed to align the appropriate doctrinal, organizational, strategic, and technological concepts.

C. DECEPTIVE CYBER DEFENSE

1. Deception in Cyberspace

Gartzke and Lindsay have identified deception as a defining aspect of interactions in the cyber domain.⁵² In cyberspace, deception is best known from the defender’s viewpoint (the attacker knows he can resort to a wide range of deceptive tools to maintain anonymity) due to the attribution problem.⁵³ American Deputy Secretary of Defense

⁴⁹ James J Wirtz, “The Cyber Pearl Harbor,” *Intelligence and National Security* 32, no. 6 (2017): 758–67.

⁵⁰ Lindsay, “Cyber Conflict vs. Cyber Command,” 260.

⁵¹ John Arquilla, *Bitskrieg: The New Challenge of Cyberwarfare*, 1st edition (Cambridge, UK and Medford, MA: Polity, 2021).

⁵² Gartzke and Lindsay, “Weaving Tangled Webs.”

⁵³ David D. Clark and Susan Landau, “Untangling Attribution,” in *Proceedings of A Workshop on Deterring Cyberattacks: Informing Strategies and Developing Options for U.S. Policy* (Washington, D.C.: The National Academies Press, 2011), 25–40.

William Lynn famously wrote, “Whereas a missile comes with a return address, a computer virus generally does not.”⁵⁴ Attribution techniques are improving, but the process remains imperfect.⁵⁵ In the same way, cyber defenders can deceive to misdirect and even infect the adversary, increasing both the risk and cost of an attack. Scholar Neil C. Rowe has described cyber deception as an effective “third line of defense,” behind access controls and intrusion defense.⁵⁶ As Gartzke and Lindsay have noted, “modern soldiers wear camouflage whether they are taking or holding territory.”⁵⁷

One among many deceptive cyber defense techniques is known as the honeypot. The National Institute of Standards and Technology (NIST) defines a honeypot as “A system (e.g., a web server) or system resource (e.g., a file on a server) that is designed to be attractive to potential crackers and intruders, like honey is attractive to bears.”⁵⁸ Security researchers have attributed the first recorded use of the honeypot, now a widespread cybersecurity technique, to Bill Cheswick of AT&T Bell Laboratories in 1991.⁵⁹ Cheswick deliberately created a vulnerability on his network’s internet gateway and left a password file with credentials for a decoy enclave. Cheswick then led an attacker “on a merry chase in order to trace his location and learn his techniques.”⁶⁰ Today, many internet security researchers employ this method to discover and characterize emerging malware.⁶¹ Honeypots have notably enabled IBM, Symantec, and others to analyze the

⁵⁴ William J. Lynn III, “Defending a New Domain,” *Foreign Affairs* (September/October 2010).

⁵⁵ Herbert Lin, *Attribution of Malicious Cyber Incidents: From Soup to Nuts*, SSRN Scholarly Paper (Rochester, NY: Social Science Research Network, September 2, 2016).

⁵⁶ Neil C. Rowe and Julian Rrushi, *Introduction to Cyberdeception* (Cham, Switzerland: Springer International Publishing, 2016), 17, <https://doi.org/10.1007/978-3-319-41187-3>.

⁵⁷ Gartzke and Lindsay, “Weaving Tangled Webs.”

⁵⁸ NIST, “Honeypot,” Computer Security Resource Center, accessed February 10, 2022, <https://csrc.nist.gov/glossary/term/honeypot>.

⁵⁹ Bill Cheswick, “An Evening with Berferd in Which a Cracker Is Lured, Endured, and Studied,” 1991, <https://www.cheswick.com/ches/papers/berferd.pdf>.

⁶⁰ Cheswick, 1.

⁶¹ “Honeypots in Cybersecurity Explained,” CrowdStrike, March 9, 2022, <https://www.crowdstrike.com/cybersecurity-101/honeypots-in-cybersecurity-explained/>.

WannaCry ransomware attacks in 2017.⁶² NIST also lists honeypots as a recommended security control in its risk management framework.⁶³ Deceptive cyber defense thus represents a useful tool for cybersecurity research, network defense, and cyber tracking and counterattack.

2. Cyber Deception Research

Today, a variety of open-source and commercial honeypot platforms are available for research or production use. For instance, researchers have simulated management services, such as the Secure Shell Protocol, to study and visualize attack patterns.⁶⁴ As concern for the resilience of critical infrastructure has grown, so has research on industrial control system (ICS) honeypots.⁶⁵ Others have developed improvements to scalability and interaction or the efficiency and versatility of honeynet (a network of honeypots) architecture.⁶⁶ However, these developments do not directly address the operationalization of deception technology. In fact, researchers have suggested that key challenges to doing

⁶² Rafael Amado, “WannaCry: An Analysis of Competing Hypotheses,” Digital Shadows, May 18, 2017, <https://www.digitalshadows.com/blog-and-research/wannacry-an-analysis-of-competing-hypotheses/>.

⁶³ *Risk Management Framework*, NIST Special Publication (SP) 800–53 (Gaithersburg, MD: National Institute of Standards and Technology, November 30, 2016).

⁶⁴ Ioannis Koniaris, Georgios Papadimitriou, and Petros Nicopolitidis, “Analysis and Visualization of SSH Attacks Using Honeypots,” in *Eurocon 2013* (IEEE, 2013), 65–72; Daniel Reti et al., “Secure (S)Hell: Introducing an SSH Deception Proxy Framework,” *ArXiv:2104.03666 [Cs]*, April 8, 2021, <http://arxiv.org/abs/2104.03666>.

⁶⁵ Daniele Antonioli, Anand Agrawal, and Nils Ole Tippenhauer, “Towards High-Interaction Virtual ICS Honeypots-in-a-Box,” in *Proceedings of the 2nd ACM Workshop on Cyber-Physical Systems Security and Privacy*, 2016, 13–22; Amine Belqruch and Abdelilah Maach, “SCADA Security Using SSH Honeypot,” in *Proceedings of the 2nd International Conference on Networking, Information Systems & Security* (2nd International Conference on Networking, Information Systems & Security, Morocco: Association for Computing Machinery, 2019), 1–5.

⁶⁶ Juan David Guarnizo et al., “Siphon: Towards Scalable High-Interaction Physical Honeypots,” in *Proceedings of the 3rd ACM Workshop on Cyber-Physical System Security* (Conference on Computer and Communications Security, Abu Dhabi, 2017), 57–68; Wenjun Fan et al., “HoneyDOC: An Efficient Honeypot Architecture Enabling All-Round Design,” *IEEE Journal on Selected Areas in Communications* 37, no. 3 (2019): 683–97, <https://doi.org/10.1109/JSAC.2019.2894307>.

so include the effort required to generate effective deception and the potential security risk associated to production honeypots.⁶⁷

Lance Spitzner discusses the placement of honeypots in relation to typical network zones, assessing that external honeypots may offer little value due to the volume of indiscriminate attacks they would observe, while demilitarized zone (DMZ) and intranet honeypots can be useful for the detection or prevention of threats.⁶⁸ However, honeypot placement is rarely the subject of academic research. In a wide-ranging survey of honeypot research, Javier Franco et al. find that only 12 of 79 studies between 2002 and 2020 disclosed the location of the honeypot.⁶⁹ The authors note that “one does not see any study in the literature that aims to optimize the deployment location for the decoy system,” adding that “this is an important gap in the honeypot/honeynet research.” Devran Yener et al. share this observation and compare the traffic generated by honeypots served on two different internet service providers (ISP).⁷⁰ Others analyze attack patterns based on honeypot geolocation or when using different cloud providers, but not in relation to an operational network.⁷¹

⁶⁷ David Liebowitz et al., “Deception for Cyber Defence: Challenges and Opportunities,” in *2021 Third IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications (TPS-ISA)*, 2021, 173–82, <https://doi.org/10.1109/TPSISA52974.2021.00020>; Iyatiti Mokube and Michele Adams, “Honeypots: Concepts, Approaches, and Challenges,” in *Proceedings of the 45th Annual Southeast Regional Conference*, ACM-SE 45 (New York, NY: Association for Computing Machinery, 2007), 321–26, <https://doi.org/10.1145/1233341.1233399>.

⁶⁸ Lance Spitzner, *Honeypots: Tracking Hackers* (Boston, MA: Addison-Wesley, 2003), 281–83, <https://books.google.com/books?id=oNFQAAAAMAAJ>.

⁶⁹ Javier Franco et al., “A Survey of Honeypots and Honeynets for Internet of Things, Industrial Internet of Things, and Cyber-Physical Systems,” *IEEE Communications Surveys Tutorials* 23, no. 4 (2021): 2351–83, <https://doi.org/10.1109/COMST.2021.3106669>.

⁷⁰ Devran Yener, Tansel Özyer, and Emin Kuğu, “Analysis and Comparison of Honeypot Activities on Two ISP Networks,” in *2021 2nd International Informatics and Software Engineering Conference (IISEC)* (IEEE, 2021), 1–6.

⁷¹ Jay Thom, Yash Shah, and Shamik Sengupta, “Correlation of Cyber Threat Intelligence Data Across Global Honeypots,” in *2021 IEEE 11th Annual Computing and Communication Workshop and Conference (CCWC)*, 2021, 0766–72, <https://doi.org/10.1109/CCWC51732.2021.9376038>; Christopher Kelly et al., “A Comparative Analysis of Honeypots on Different Cloud Platforms,” *Sensors* 21, no. 7 (January 2021): 2433, <https://doi.org/10.3390/s21072433>.

Darknets or network telescopes monitoring large subnets of inactive internet addresses offer a complementary approach to the analysis of attack patterns.⁷² Where honeypot deployments typically simulate a small number of endpoints, darknets offer the benefit of scale but perform limited interaction with the attacker. The impact of darknet size and geolocation are the ongoing subject of research, but the relationship between attack patterns on dark networks and operational networks remains an open question.⁷³ Large scale deployment of honeypots on darknets shows the potential to narrow the gap by improving the deceptive effectiveness relative to monitoring darknets.⁷⁴ Even with a higher degree of interaction, darknets may not form an accurate representation of the threats that operational networks will face. For instance, Philipp Richter et. al suggest that the mere proximity of a darknet IP range to an operational network significantly affects attack patterns.⁷⁵

D. CONCLUSION

Deception has a proven track record developed over the millennia of recorded military affairs yet remains underappreciated in academic research and in practice. Similarly, deception is essential tradecraft to any cyber attack, yet its potential has not been fully harnessed by network defenders. Scholars have suggested that deceptive cyber defense could alter the balance of cyber conflict in favor of the defender, but there is a clear lack of research to support this claim and inform an optimal implementation of this stratagem.

⁷² Karyn Benson et al., “Leveraging Internet Background Radiation for Opportunistic Network Analysis,” in *Proceedings of the 2015 Internet Measurement Conference* (2015 Internet Measurement Conference, Tokyo: Association for Computing Machinery, 2015), 423–36.

⁷³ Francesca Soro et al., “Are Darknets All The Same? On Darknet Visibility for Security Monitoring,” in *2019 IEEE International Symposium on Local and Metropolitan Area Networks (LANMAN)*, 2019, 1–6, <https://doi.org/10.1109/LANMAN.2019.8847113>.

⁷⁴ Morteza Safaei Pour, Joseph Khoury, and Elias Bou-Harb, “HoneyComb: A Darknet-Centric Proactive Deception Technique For Curating IoT Malware Forensic Artifacts,” in *NOMS 2022–2022 IEEE/IFIP Network Operations and Management Symposium*, 2022, 1–9, <https://doi.org/10.1109/NOMS54207.2022.9789827>.

⁷⁵ Philipp Richter and Arthur Berger, “Scanning the Scanners: Sensing the Internet from a Massively Distributed Network Telescope,” in *Proceedings of the Internet Measurement Conference*, 2019, 144–57.

III. THE ROLE OF DECEPTIVE DEFENSE IN CYBER STRATEGY: LESSONS FROM DECOY VESSELS OF THE GREAT WAR

On March 3, 2023, the White House released the highly anticipated *National Cybersecurity Strategy*—formally asserting an active approach to cybersecurity.⁷⁶ In fact, a full pillar of the strategy is dedicated to preventive or pre-emptive actions (“Disrupt and Dismantle Threat Actors”).⁷⁷ Usually, proactive network defense describes threat hunting, which reflects the acceptance that a determined cyber adversary will find a way past the network perimeter.⁷⁸ Yet many organizations still struggle to implement this practice, primarily due to a lack of resources.⁷⁹ In addition, cyber threat actors easily relocate across borders, and proxies blur lines between state and non-state actors.⁸⁰ How then can government agencies efficiently identify—let alone disrupt or dismantle—these dynamic entities?

Deceptive cyber defense holds the key to an effective active network defense strategy and can enable a capability to deter cyber threat actors. In the previous chapter, we presented a brief history of deception in military history and identified it as a defining characteristic of cyber conflict, with implications for offensive and defensive operations. This chapter shows that deceptive cyber defense can have a profound impact on cyber strategy. We present a case study of Q-ships, anti-submarine warships disguised as commercial vessels during the First World War, to illustrate potential effects of deceptive cyber defense: *identification*, *misdirection*, and *retaliation*. We then develop these effects

⁷⁶ The White House, *National Cybersecurity Strategy* (Washington, DC: White House, 2013), <https://www.whitehouse.gov/wp-content/uploads/2023/03/National-Cybersecurity-Strategy-2023.pdf>.

⁷⁷ The White House, 14–18.

⁷⁸ “What Is Cyber Threat Hunting?,” CrowdStrike, April 17, 2022, <https://www.crowdstrike.com/cybersecurity-101/threat-hunting/>.

⁷⁹ Mathias Fuchs and Josh Lemon, *SANS 2022 Threat Hunting Survey* (Rockville, MD: SANS Institute, July 2022).

⁸⁰ Tim Maurer, *Cyber Mercenaries: The State, Hackers, and Power* (Cambridge, UK: Cambridge University Press, 2018), <https://doi.org/10.1017/9781316422724>.

in more detail. In addition, we explore the security-deception trade-off, which hinders deceptive effectiveness in traditional honeypot deployments.

A. CASE STUDY: Q-SHIPS

1. Background: The U-Boat Campaign

On October 20, 1914, the German U-boat *U-17* fired a shot across the bow of SS *Glitra*, a British merchant ship sailing off the coast of Norway.⁸¹ In accordance with naval tradition, *U-17* ordered the crew to escape before scuttling the 866-ton ship.⁸² SS *Glitra* was the first trade vessel to be sunk by a U-boat during the First World War. When *U-20* torpedoed the British commercial liner RMS *Lusitania* in May 1915, the crew and passengers were not so lucky: nearly 1,200 men, women, and children lost their lives.⁸³ The incident caused massive global outrage, reducing Germany's moral standing and energizing the war movement in the United States.⁸⁴ Nonetheless, U-boats continued to attack commercial ships. Acting on promises from the Imperial German Navy that the U-boat campaign could "starve Britain into submission," the German Chancellor even resumed "unrestricted submarine warfare" in January 1917.⁸⁵ Over the course of the war, U-boats would attack thousands of allied and neutral ships, sending more than 11 million gross tons of shipping to the bottom of the ocean.⁸⁶

With the tragedy of RMS *Lusitania* arose an opportunity. After the sinking of the commercial liner, the Kaiser directed a moderation of the U-boat campaign to avoid further

⁸¹ Lawrence Sondhaus, *German Submarine Warfare in World War I: The Onset of Total War at Sea* (Lanham, MD: Rowman & Littlefield, 2017), 5–6.

⁸² J. Ashley Roach, "The Law of Naval Warfare at the Turn of Two Centuries," *The American Journal of International Law* 94, no. 1 (2000): 73–74, <https://doi.org/10.2307/2555231>.

⁸³ Diana Preston, *Lusitania: An Epic Tragedy* (New York, NY: Bloomsbury Publishing USA, 2002).

⁸⁴ Thomas A. Bailey, "The Sinking of the Lusitania," *The American Historical Review* 41, no. 1 (1935): 73, <https://doi.org/10.2307/1839355>.

⁸⁵ Dirk Steffen, "The Holtzendorff Memorandum of 22 December 1916 and Germany's Declaration of Unrestricted U-Boat Warfare," *The Journal of Military History* 68, no. 1 (2004): 215–24.

⁸⁶ R. H. Gibson and Maurice Prendergast, *The German Submarine War 1914–1918* (Cornwall, UK: Periscope Publishing Ltd., 2002), 380–82.

angering the United States and other neutral countries.⁸⁷ Remaining in accordance with “prize rules,” U-boats would surface near trading vessels to give the passengers a chance to evacuate before scuttling the ship, as they had with the SS *Glitra*.⁸⁸ While the U-boats were a deadly and silent threat to any ship when submerged, they were far more vulnerable when surfaced. Admiral Karl Dönitz, who captained a U-boat during the First World War and rose to high command during the Second World War, wrote that surfaced U-boats were “slow” and “low in the water with a restricted field of vision,” leaving them ill-suited for a direct confrontation.⁸⁹ U-boats also lacked armor, such that a single shot could “sink or at least prevent the ship from diving.”⁹⁰ The British, who were desperate to find solutions to the German raids, looked for new ways to exploit this weakness.

2. Q-ships

Naval officers had already proposed the idea of decoys to the British Admiralty in the autumn of 1914.⁹¹ Winston Churchill, First Lord of the Admiralty at the time, instructed the Commander-in-Chief to proceed with arming “small to moderate sized steamers” to “trap the German submarine” and “sink her with gunfire.”⁹² The Royal Navy had a working prototype in the form of SS *Vittoria*, refitted from commercial trawler.⁹³ SS *Vittoria* and other early decoy ships were not impactful, as shipping losses inflicted by the U-boats rose to 1.4 million gross tons by the end of 1915.⁹⁴ An updated concept saw the British deploy a trawler, USS *Taraniki*, paired with the submarine C-24. On June 23, 1915,

⁸⁷ Gerd Hardach, *The First World War, 1914–1918* (Oakland, CA: University of California Press, 1981), 41.

⁸⁸ Andreas Michelsen, *The Submarine Warfare, 1914–1918* (Washington, DC: U.S. Government Publishing Office, 1926), 51.

⁸⁹ Karl Dönitz, *Memoirs Ten Years and Twenty Days* (Chicago, IL: Frontline Books, 2012), 10.

⁹⁰ Richard Compton-Hall, *Submarines at War 1914–1918* (Cornwall, UK: Periscope Publishing Ltd., 2004), 91.

⁹¹ Edward Keble Chatterton, *Q-Ships and Their Story* (Boston, MA: Sidgwick and Jackson, Limited, 1922), 7–8.

⁹² Tony Bridgland, *Sea Killers in Disguise: The Story of the Q-Ships and Decoy Ships in the First World War* (Annapolis, MD: Naval Institute Press, 1999), x.

⁹³ Chatterton, *Q-Ships and Their Story*, 7.

⁹⁴ Gibson and Prendergast, *The German Submarine War 1914–1918*, 380.

the USS *Taraniki* spotted a German submarine and promptly informed C-24 via underwater cable. C-24 positioned itself to attack and sank U-40 with a single torpedo, killing all but three members of the crew.

Decoy vessel tactics gained momentum in the summer of 1915. In the aftermath of the German sinking of RMS *Lusitania*, the British strove to better coordinate the deception campaign.⁹⁵ The Royal Navy launched new decoy ships in Queenstown, Ireland, where they became known as Q-ships. Vice-Admiral Gordon Campbell, who commanded a Q-ship during the war, detailed improvements to their “disguises,” including false cargo and neutral colors according to their itinerary.⁹⁶ Vice-Admiral Andreas Michelsen, Commander of the German Submarine Forces starting in June 1917, reported that “even a visitor on board ... could not recognize [their] true character, and a careful examination through the periscope would fail to reveal anything suspicious.”⁹⁷ Q-ship crews were dressed as civilians (even as women) and trained to simulate commercial crews in every aspect, including the nomination of a “panic crew,” which would abandon the ship in a false demonstration of vulnerability.⁹⁸

Efforts to improve the deceptive effectiveness of Q-ships were soon rewarded. On July 26, 1915, the HMS *Prince Charles* spotted a U-boat to the north of Scotland.⁹⁹ The Q-ship waited for the submarine to engage, at which time it stopped its engines and lowered the boats to simulate an evacuation. When U-36 exposed its broadside in its approach, the crew of HMS *Prince* uncovered their guns and fired on the U-boat. U-36 was the first German submarine to be sunk by a Q-ship, but two more U-boats were lost to the decoy vessels by the end of 1915. As German knowledge of the Q-ship stratagem spread, the

⁹⁵ Bridgland, *Sea Killers in Disguise*, 6.

⁹⁶ Gordon Campbell, *My Mystery Ships* (Garden City, New York: Doubleday, Doran & Company, 1929), 43–47.

⁹⁷ Michelsen, *The Submarine Warfare, 1914–1918*, 70.

⁹⁸ Mary T. Hall, “False Colors and Dummy Ships: The Use of Ruse in Naval Warfare,” *Naval War College Review* 42, no. 3 (1989): 60, <https://www.jstor.org/stable/44636917>.

⁹⁹ Bridgland, *Sea Killers in Disguise*, 10.

Royal Navy adapted by increasing the armor and armament of the decoy ships.¹⁰⁰ However, one of the last Q-ship engagements featured the opposite approach: HMS *Stockforce* was a small (360-ton) steamer commanded by Lieutenant Harold Auten.¹⁰¹ German U-boats found it more difficult to fire effectively on smaller ships and would be forced to surface.¹⁰² In July 2018, HMS *Stockforce* sank *UB-80* in a dramatic engagement that won Auten a Victoria Cross.¹⁰³

Q-ships had a significant impact on the U-boat campaign during the First World War. Over the next years, their numbers soared as the Royal Navy commissioned up to 180 vessels.¹⁰⁴ Actively hunting U-boats, they saved countless commercial vessels and the lives of their crews simply by drawing U-boats away from vulnerable targets. Estimates on the total number of U-boats sunk by Q-ships range from 11 to a “couple of dozen,” with up to 80 U-boats damaged and sent for repairs.¹⁰⁵ Beyond the tangible cost of replacing and repairing U-boats, decoy ships may have inflicted a moral effect on the German submarine crews, who suffered a “great wariness” when approaching any vessel.¹⁰⁶ Such a change in U-boat tactics is indicative of the impact that Q-ships had in the war.¹⁰⁷

The used of Q-ships eventually declined, as any military technology naturally does. Critics have claimed that the U-boat campaign may have become more ruthless because of the Q-ship stratagem.¹⁰⁸ Raiding submarines eventually became reluctant to surface, thus removing any opportunity for the crews of civilian ships to evacuate. Yet, President

¹⁰⁰ Campbell, *My Mystery Ships*, 308.

¹⁰¹ Bridgland, *Sea Killers in Disguise*, 135.

¹⁰² Michelsen, *The Submarine Warfare, 1914–1918*, 70.

¹⁰³ Harold Auten, “Q” *Boat Adventures: The Exploits of the Famous Mystery Ships* (London: H. Jenkins, 1919).

¹⁰⁴ Chatterton, *Q-Ships and Their Story*, 136.

¹⁰⁵ Bridgland provides context to the varying estimates here: Bridgland, *Sea Killers in Disguise*, 147; Chatterton reports “over eighty” damaged U-boats here: Chatterton, *Q-Ships and Their Story*, 137.

¹⁰⁶ Gibson and Prendergast, *The German Submarine War 1914–1918*, 156.

¹⁰⁷ Richard W. Smith, “The Q-Ship—Cause and Effect,” *U.S. Naval Institute Proceedings* 79, no. 3 (1953), <https://www.usni.org/magazines/proceedings/1953/may/q-ship-cause-and-effect>.

¹⁰⁸ Campbell, *My Mystery Ships*, 304.

Woodrow cited “unrestricted submarine warfare,” in the U.S. declaration of war in April 1917.¹⁰⁹ Q-ships and other anti-submarine measures contributed to the reduced effectiveness of the U-boat campaign under “prize rules,” which Germans concluded would be insufficient to starve Britain.¹¹⁰ Thus, the Q-ship (along with other anti-submarine measures) may have indirectly contributed to the United States joining the war. Other critics have wisely noted that the decline in the success rate of U-boats was surely due to a combination of factors, including mines and naval convoys, and cannot be solely attributed to the Q-ship.¹¹¹ Some scholars have also debated the legality of deception, possibly creating a perception of moral equivalency between U-boat raids and Q-ships.¹¹² The nuances between perfidy and ruse may be subtle, but at least the British ordered Q-ships to change their colors just before opening fire, as international law required.¹¹³

3. Cyber Q-ships?

Q-ships illustrate some potential effects of deceptive cyber defense. Decoy vessels lured the stealthy U-boats to a position of vulnerability when surfaced, allowing for the *identification* of threats. In the same way, an effective defensive cyber deception may compel the adversary to an engagement on the defender’s terms. By offering a target of opportunity, Q-ships and deceptive cyber defense *misdirect* the adversary, drawing them away from their intended targets. The ultimate purpose of Q-ships was to destroy or disable U-boats, and they did so through deceptive defense, in *retaliation* for an initial engagement by the U-boat. Current cyber response practice is analogous to launching an anti-submarine ship to search for a U-boat that had sunk a defenseless ship a few months earlier. With Q-

¹⁰⁹ Woodrow Wilson, “Joint Address to Congress Leading to a Declaration of War Against Germany (1917),” National Archives, April 2, 1917, <https://www.archives.gov/milestone-documents/address-to-congress-declaration-of-war-against-germany>.

¹¹⁰ Steffen, “The Holtzendorff Memorandum of 22 December 1916 and Germany’s Declaration of Unrestricted U-Boat Warfare.”

¹¹¹ Chatterton, *Q-Ships and Their Story*, 270.

¹¹² Matthew G. Morris, “‘Hiding Amongst a Crowd’ and the Illegality of Deceptive Lighting,” *Naval Law Review* 54 (2007): 256, <https://apps.dtic.mil/sti/citations/ADA476997>.

¹¹³ Chatterton, *Q-Ships and Their Story*, 8.

ships, as with deceptive cyber defense, retaliation can occur immediately after an adversarial attack.

Lessons from the employment of Q-ships can be applied to the operationalization of deceptive cyber defense. No analogy is perfect, beginning with the violence and other hardships that the crews of Q-ships needed to contend with. Nonetheless, the bravery and genius of the men who served on these decoy vessels demonstrates that effective deception results from creativity, attention to detail, and operational security. As the deceptive quality of Q-ships improved in 1915, so did the stratagem's results. To achieve analogous effects in the cyber domain, defenders must optimize every aspect of deceptive cyber defense. The Royal Navy regularly modified the characteristics of the decoy ships throughout the Great War, extending the useful lifetime of the stratagem. Only once the U-boats fundamentally shifted their tactics did the direct threat of Q-ships abate, and by then, other means of anti-submarine warfare were available to compensate.

B. EFFECT DESCRIPTIONS

1. Identification

Identification provides a general understanding of the cyber threat environment and the specific detection and characterization of adversarial attacks. The former can inform a network's internal defensive measures. As discussed in Chapter II, researchers have long employed deception to characterize emerging threats. In a recent example, scholars utilized honeypots to study the "log4j" attacks.¹¹⁴ Defenders can also leverage the information gained through deceptive cyber defense. For instance, malicious internet addresses can be detected on an external honeypot and blocked from the network potentially faster than they can be identified and indexed by a threat intelligence provider. In addition, these malicious addresses may represent a more relevant asset than a generic blacklist.

¹¹⁴ Shreyas Srinivasa, Jens Myrup Pedersen, and Emmanouil Vasilomanolakis, "Deceptive Directories and 'Vulnerable' Logs: A Honeypot Study of the LDAP and Log4j Attack Landscape," in *2022 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, 2022, 442–47, <https://doi.org/10.1109/EuroSPW55150.2022.00052>.

Deceptive cyber defense supports the detection of attacks by luring the adversary to an engagement on the defender's terms. For instance, a defender can configure a deceptive asset in a way which negates any possibility of legitimate communication. Thus, an actor's mere engagement with the asset constitutes an indicator of compromise. In this sense, deceptive cyber defense serves as an enabler to threat-hunting operations. With additional traffic and attack activity from the adversary, the defender may even observe multiple stages of the attacker's infrastructure (i.e., the initial compromise may be conducted via a botnet while more advanced exploitation through a more permanent channel), facilitating attribution.¹¹⁵ Similarly, deception may also yield insight into the types of malicious activities adversaries are undertaking and their goals.

2. Misdirection

In the context of deceptive cyber defense, misdirection primarily aims to keep the adversary away from one's own valued resources. Many cyber threat actors strike targets of opportunity: they look for assets that are vulnerable to exploits that are available to them. Consequently, a simple deception that mirrors a known vulnerability may be enough to occupy the adversary long enough for the defender to counter the attack. Delay tactics could extend the duration of the effect. For instance, the LaBrea Tarpit responds to traffic destined for unused addresses, purposely delaying incoming connections.¹¹⁶

A cognitive-level misdirection effect can be achieved when the defender leaves false or misleading information on a deceptive enclave of a network. Such a stratagem may not tactically improve network defense but could strategically alter the balance of information in favor of the defender. The defender could follow the trail of misinformation to gain information about the adversary or simply benefit from the effect of this misinformation on the decision processes of the adversary. In certain applications, the

¹¹⁵ David D. Clark and Susan Landau, "The Problem Isn't Attribution: It's Multi-Stage Attacks," in *Proceedings of the Re-Architecting the Internet Workshop* (Conference on emerging Networking EXperiments and Technologies, Philadelphia, PA: Association for Computing Machinery, 2010), 1–6.

¹¹⁶ Tom Liston, "LaBrea – The Tarpit," accessed April 19, 2023, <https://labrea.sourceforge.io/Intro-History.html>.

defender could simply ask the target of the deception to provide information, as the Dutch police did in the case of the AlphaBay decoy.¹¹⁷

3. Retaliation

In retaliation, the roles are flipped: the defender wants to take the offensive to disrupt or dismantle a threat actor, as the U.S. *National Cybersecurity Strategy* suggests. Normal (non-deceptive), targeted offensive cyber operations are much more difficult to conduct than opportunistic attacks that cybercriminals and other cyber threat actors routinely conduct. Rather than match a network to a given exploit as the opportunistic attacker would, a discriminate attacker (e.g., the U.S. government when seeking to disrupt a specific cyber threat actor) must identify a vulnerability on a given network and match it to an exploit to gain access. Thus, without the valuable context potentially provided by defensive operations, offensive planners begin with a blank slate and must scan the adversarial network perimeter to gain access. This deficit should not be understated. Erica Borghard and Shawn Lonergan explain that the task of gaining access to an adversarial network is complex and resource intensive.¹¹⁸ Thus, the information that can be gained on an adversary through deceptive cyber defense can provide vital input to plan retaliatory operations. For example, forensic examination of adversarial network traffic on a deceptive enclave can reveal internet addresses corresponding to multiple stages of a threat actor's network infrastructure as well operating systems and other software that they employ.

Deceptive cyber defense can also enable retaliation by bridging offensive and defensive cyber operations. Offensive and defensive cyber operations are often planned in isolation, through different structures, processes, and people. By collaborating closely, defensive and offensive planners can take advantage of opportunities created by deceptive cyber defense. Within the deceptive enclave, offensive teams can plant canaries which would generate alerts when files are open, backdoors to establish a persistent presence within the adversarial network, or a combination of malware to achieve the desired effect.

¹¹⁷ Barniuk, Chris, "AlphaBay and Hansa Dark Web Markets Shut Down," BBC News, July 20, 2017, <https://www.bbc.com/news/technology-40670010>.

¹¹⁸ Erica D. Borghard and Shawn W. Lonergan, "Cyber Operations as Imperfect Tools of Escalation," *Strategic Studies Quarterly* 13, no. 3 (2019): 125, <https://www.jstor.org/stable/26760131>.

Deceptive cyber defense thus presents a possibility to bypass the complex task of gaining access to an adversarial network. Further, the continuous thread between the initial attack and retaliation provides some confidence that the retaliation will strike the intended target. Closing the gap in time between a malicious action (i.e., a network breach) and the consequences (a swift retaliation enabled by deceptive cyber defense) may have some deterrent effect against future attacks.

4. Summary

The effects of deceptive cyber defense can be illustrated through the concept of time and frameworks such as the cyber kill chain.¹¹⁹ Lockheed Martin's version of the cyber kill chain includes seven stages (Figure 1).¹²⁰ *Identification* exposes and thus enables the interruption of the adversary's cyber kill chain. *Misdirection* focuses the adversary's efforts away from vital network resources and delays the adversary's cyber kill chain to enable a defensive response. *Retaliation* leverages the information provided by deception to speed up one's own cyber kill chain relative to a "blank slate" attack. Thus, deceptive cyber defense aims to optimize the defender's cyber kill chain while hindering the attacker's, potentially contributing to a favorable position for the defender in the overarching information balance.




Figure 1. Lockheed Martin's cyber kill chain¹²¹

¹¹⁹ For examples of how deception can be illustrated through the cyber kill chain, see Li Zhang and Vrizlynn LL Thing, "Three Decades of Deception Techniques in Active Cyber Defense-Retrospect and Outlook," *Computers & Security* 106 (2021): 102288.

¹²⁰ Michael Muckin and Scott C. Fitch, *A Threat-Driven Approach to Cyber Security* (Bethesda, MD: Lockheed Martin Corporation, 2014).

¹²¹ Adapted from Muckin and Fitch.

C. NAVIGATING THE SECURITY-DECEPTION TRADE-OFF

Effective deception underpins the sustained adversarial engagement which leads to the effects we previously described. Ultimately, deception is the result of a target's misperception of an objective reality.¹²² Deceivers seek to alter a target's perception by exploiting cognitive biases and otherwise manipulating the target's senses (e.g., camouflage).¹²³ These artifices will then compete against natural manifestations of reality; a target will weigh all available evidence to establish its perception of the environment. Thus, as Czechoslovak intelligence officer Ladislav Bittman wrote in his 1972 account of Soviet political warfare, a successful deception campaign will likely possess a kernel of truth.¹²⁴ If the deception is too close to the truth, however, one risks drawing attention to that which one intends to conceal.

In the age of modern surveillance, how could one conceal the invasion of a bordering nation with a force of 150,000 troops on the march? Throughout January and February 2022, Russia announced large-scale military exercises with the stated purpose of “suppressing and repelling external aggression while conducting a defensive operation” and simultaneously participating in diplomatic talks.¹²⁵ U.S. leaders struggled to convince their allies, including Ukraine, of the imminent invasion, even armed with overwhelming evidence.¹²⁶ Russia crafted a deception focused on the purpose of troops staged on Ukraine's border, a more credible claim than to deny their existence. However, this acknowledgment also invited scrutiny on their troop movements and lent credibility to the U.S. claim that Russia's intent was to invade. This conundrum illustrates the security-

¹²² Robert Jervis, *Perception and Misperception in International Politics: New Edition* (Princeton, NJ: Princeton University Press, 2017), <https://doi.org/10.2307/j.ctvc77bx3>.

¹²³ Guy Hartcup, *Camouflage: The History of Concealment and Deception in War* (Barnsley, South Yorkshire, UK: Pen and Sword Military, 2008).

¹²⁴ Ladislav Bittman, *The Deception Game: Czechoslovak Intelligence in Soviet Political Warfare* (Syracuse, NY: Syracuse University Research Corp., 1972), 168.

¹²⁵ Nathan Hodge, Tara John, and Anna Chernova, “Ukraine: Russia and Belarus Hold Joint Military Exercises,” CNN, February 10, 2022, <https://www.cnn.com/2022/02/10/europe/ukraine-russia-news-thursday-military-exercises-intl/index.html>.

¹²⁶ Shane Harris et al., “Road to War: U.S. Struggled to Convince Allies, and Zelensky, of Risk of Invasion,” *Washington Post*, August 16, 2023, <https://www.washingtonpost.com/national-security/interactive/2022/ukraine-road-to-war/>.

deception trade-off: *ceteris paribus*, the closer a deception is to reality, the more effective it will be, but the greater a security risk it will represent.

The security-deception trade-off undermines deception technology deployments. In deceptive cyber defense, one typically deploys deception within the DMZ, a low security area outside of the internal network perimeter. DMZs will contain public services such as mail and web servers, which the attacker may find less appealing than the more confidential resources on an organization's internal network. Further, defenders usually configure the honeypot on a separate machine residing on an isolated segment of the DMZ, minimizing both the risk of compromise of other resources in the DMZ and the appeal of the deception. While this architecture may present a lesser form of deception, it is more secure than the alternative, where one would allow the adversary to access a honeypot within the secure network perimeter. Are defenders doomed to implement deception that is weak and secure, or effective but insecure? A study of deceptive techniques employed by cyber threat actors reveals a way to potentially circumvent the security-deception trade-off.

In 2019 and 2020, the Israeli government and technology companies in Israel were subject to a series of cyber attacks.¹²⁷ Attackers employed tools that were associated with Iranian groups, left traces in the Farsi language, and references to Iran in folder structures. As Israel and Iran were no strangers to cyber conflict, there was little reason to doubt the obvious verdict. Only this was not the clear-cut case it seemed to be: researchers have more recently determined that the attacks were far more likely undertaken by a Chinese group.¹²⁸ Using a mix of publicly available and state-developed malware that inadvertently became their signature, the same group had wreaked havoc on several other targets in the Middle East and in Asia to support China's political interests. Analysts from Mandiant suggested the hackers made these mistakes because they were not overly concerned with

¹²⁷ Patrick Howell O'Neill, "Chinese Hackers Disguised Themselves as Iran to Target Israel," *MIT Technology Review*, August 10, 2021, <https://www.technologyreview.com/2021/08/10/1031622/chinese-hackers-false-flag-iran-israel-fireeye/>.

¹²⁸ "UNC215: Spotlight on a Chinese Espionage Campaign in Israel," FireEye, accessed August 29, 2021, <https://www.fireeye.com/blog/threat-research/2021/08/unc215-chinese-espionage-campaign-in-israel.html>.

attribution. Yet, even with these limited efforts, attribution was only possible after extensive analysis that would connect the dots between these attacks.

Defenders have the potential to compensate for the security-deception trade-off by using the same techniques that Chinese hackers used in this false-flag operation. Misleading files and folder structures can give a legitimate appearance to the honeypot, a natural if common step used in honeypots. Truly embracing a hacker's mindset yields more subtle layers to the deception, however: defenders can leverage spoofing, a technique by which a threat actor will impersonate a legitimate device by modifying the source address in their network traffic. While defenders can protect a network against spoofing by implementing additional forms of access control, attackers have limited means to detect address modification. Hence, a honeypot outside of the network perimeter could appear to be a target of choice, if one redirected pre-designated adversarial attack traffic (e.g., based on destination IP address and port) through the edge of the network while modifying network addresses accordingly.

D. CONCLUSION

The *National Cybersecurity Strategy* warns us of a new era of “deepening digital dependencies” and an escalation of the impacts of malicious cyber activity, from a mere nuisance years ago to “damaging attacks against critical infrastructure, to ransomware attacks and cyber-enabled influence campaigns.” Iterating the words “active” or “proactive” a total of 11 times in the 34-page document, the White House is clearly indicating that it wants to regain the initiative. Deceptive cyber defense, like the Q-ships of the past century, can enable an active posture to detect threats and misdirect the adversary (Table 1). Most importantly, deception can allow the United States to regain the initiative in cyber conflicts by bridging offensive and defensive cyber operations. However, the deception-security trade-off and other challenges currently hinder deceptive cyber defense. Thus, defenders must develop new ways to enhance deceptive effectiveness. We present such a concept in the following chapter.

Table 1. Effects of deceptive cyber defense

Effect	Description	Q-Ship Analogy
Identification	Gain understanding of threat environment and detect ongoing attacks.	Q-ships lured U-boats to the surface.
Misdirection	Draw the adversary away from own network assets and intent.	Q-ships were warships camouflaged as trade vessels, intending to draw adversaries away from their intended targets.
Retaliation	Bridge defensive and offensive operations to bypass the adversarial network perimeter.	Q-ships struck U-boats once they were in a vulnerable position.

IV. DRACO: SECURE AND COMPELLING DECEPTION

A. INTRODUCTION

Deceptive Resistance to Adversary Cyber Operations (or DRACO) is a concept born from the desire to improve the operationalization of deceptive cyber defense. As detailed in previous chapters, deception technology shows promise but has not reached its full potential. Deception may represent a significant integration burden, fail to deliver an effective deception, and can represent a security risk. Through the application of a hacker's mindset, DRACO offers secure and compelling engagement with the adversary. In this chapter, we introduce the DRACO concept and illustrate the contrast to typical honeypot deployments through operational scenarios.

B. DECEPTIVE RESISTANCE TO ADVERSARY CYBER OPERATIONS

1. Architecture

DRACO addresses key challenges associated with deception technology by introducing an architecture (Figure 2) which presents an effective deception to cyber threat actors while reducing the integration burden for the end user. Use of network address translation or “spoofing” allows deceptive resources to remain separate from the end user's secure network, while providing responses that appear authentic to cyber threat actors. When designated attack traffic (e.g., an unauthorized connection attempt) is received at the edge of the network, the router forwards it to a network address translation (NAT) service, which modifies the destination address to match a deception network. When the deception network responds, the NAT alters the source address to correspond with the intended target of the attack. This “deception as a service” approach enables hosting the deceptive network component on the Cloud for easy deployment, separation from operational network infrastructure, and enhanced availability.




Figure 2. DRACO architecture

a. Main Components

- (1) **Deception Network** The deception network is the component which interacts with cyber threat actors and is typically managed by the DRACO service provider as separate infrastructure (“deception as a service”). This network may consist of a deception server, analytics/database server, a network firewall, and a router. The role of each sub-component is as follows:
 - Deception Server. This server responds to the cyber threat actor on one or more desired networking protocols and logs relevant network events.
 - Analytics/Database Server. This server receives logs from the deception server and stores them in the database. It also generates dashboards to provide visual analytics to the end user. This server may be inside or outside of the deception network (e.g., firewalled from the deception server).

- Network Firewall. This firewall manages connections to the deception network from the internet. The administration of the deception network is allowed through the firewall.
 - Router. This device routes traffic between the components of the deception network and terminates the virtual private network (VPN) connection to the operational network.
- (2) **Operational Network** The end user's network will forward designated attack traffic to the deception network via a NAT service which will change the destination address on incoming traffic and change the source address on the outgoing traffic. The role of each sub-component is as follows:
- Border Router. This device will be configured to allow designated attack traffic (e.g., matching a certain address range and/or protocol) to NAT service. No attack traffic is ever routed to the operational firewall.
 - NAT. This device will be configured to change the destination address of incoming traffic (from the threat actor to the deception server) to reflect the address of the deception server and to change the source address of outgoing traffic (from the deception server to the threat actor) to reflect the target server.

2. Insider Threat Variant

In its basic configuration, DRACO provides the appearance of an opening in the network perimeter to generate engagement with an external adversary. It is worth noting that not all cyber threat actors are external; malicious insiders also represent a significant attack vector. A report from IBM found that on average, a network compromise initiated by an insider cost \$4.8M to resolve.¹²⁹ Insiders have also been responsible for a series of high profile leaks of classified information, including Edward Snowden, Chelsea Manning,

¹²⁹ IBM, "Cost of a Data Breach."

and most recently Jack Teixeira.¹³⁰ Modification of the original architecture allows DRACO–Insider Threat Variant (DRACO-ITV) to redirect attacks or suspicious traffic that originate from within the network (Figure 3). From any segment of the intranet, we configure the local router to redirect unauthorized traffic to the cloud-hosted deception network via a NAT and achieve the same deceptive effect on the adversary. Since DRACO routes the response back to the attacker through the same network path on which it was received, it is difficult for the attacker to detect this redirection without multiple vantage points within the network.




Figure 3. DRACO–Insider Threat Variant (DRACO-ITV)

¹³⁰ Michael Crowley, "Airman Charged in Leak of Classified Documents," *The New York Times*, April 14, 2023, <https://www.nytimes.com/2023/04/14/us/politics/jack-teixeira-classified-documents-leak.html>.

C. ALTERNATIVE HONEYPOT ARCHITECTURES

Honeypots are usually categorized by application (research or production) or by the degree of adversarial interaction (low, medium, and high).¹³¹ However, these categories oversimplify management and integration concerns. In practice, the network architecture plays a role in the maintenance of the honeypot and in itself contributes to the adversarial interaction levels. Thus, we consider the placement of honeypots according to typical network zones (Figure 4) to focus on practical honeypot deployment trade-offs.




Figure 4. Typical network security zones with corresponding options for honeypot deployment¹³²

1. Operational Scenarios

To assess DRACO’s potential effectiveness, we qualitatively evaluate the outcomes to operational scenarios with each honeypot deployment location and with DRACO. We consider three evaluation criteria:

¹³¹ Mokube and Adams, “Honeypots.”

¹³² Adapted from Tracy Miles et al., *Guidelines on Securing Public Web Servers*, NIST SP 800-44, version 2 (Gaithersburg, MD: National Institute of Technology, 2007), 8–2.

- Ease of integration and maintenance. How much effort, including technical and procedural considerations, is required to deploy, integrate, and maintain the deceptive solution?
- Deceptive effectiveness. Based on the network architecture, how likely are threat actors to engage the deceptive endpoint? Is the deception more likely to simply delay an attack or to generate a more sustained engagement?
- Security risk. What potential vulnerabilities are generated by the integration of the deceptive endpoint, considering the likely abilities and resources of the cyber threat actors?

Our operational scenarios consist of modeled adversaries. In Canada, the Communications Security Establishment (CSE) categorizes cyber threat actors by affiliation, from hacktivists and cybercriminals to nation-state actors and everything in between (Figure 5).¹³³ Each threat actor possesses a distinctive motivation and degree of sophistication beyond these broad categories. While defenders must protect the network against every threat, we select two adversaries: a cybercriminal with moderate sophistication and a nation-state advanced persistent threat (APT). The former may represent the most common threat to a network, while the latter may represent the most dangerous threat (especially to military networks). These example adversarial categories also broadly represent other threat actor categories featuring a similar degree of sophistication but with different motives, whether terrorist groups, hacktivists, or thrill seekers.

¹³³ Canadian Centre for Cyber Security, “An Introduction to the Cyber Threat Environment,” October 28, 2022, <https://www.cyber.gc.ca/en/guidance/introduction-cyber-threat-environment>.




Figure 5. Cyber threat actors and their primary motivation¹³⁴

In our scenario, the cybercriminal indiscriminately scans the address space, opportunistically seeking known vulnerabilities to exploit and breach the network perimeter. Upon gaining access, this adversary intends to, for example, deploy ransomware and exfiltrate data for resale on the dark web. In contrast, the nation-state APT specifically targets the operational network, utilizing zero-day exploits to infiltrate the network perimeter with the goal of compromising deep-rooted targets while maintaining cover through covert communication channels. Owing to its nation-state affiliation, the APT benefits from signals intelligence and advanced tradecraft, enabling it to assess an endpoint’s activity. On the other hand, the cybercriminal relies on scanning endpoints for indications of active use, fingerprinting to identify known honeypot distributions, and leveraging commercial intelligence providers for further insights.

2. External Honeypot

Border routers and network access switches form an external, or internet-facing, zone for hosting a honeypot. Edge devices may carry out basic IP filtering, but their main responsibility is to carry traffic between the firewall (or firewalls if an organization operates many different networks) and the ISP. The external network does not typically

¹³⁴ Source: Canadian Centre for Cyber Security.

serve network applications and most traffic will be encrypted and considered untrusted. Thus, an external honeypot offers a limited security risk and streamlined integration due to its isolation from the internal network and the lack of plaintext traffic, but its deceptive effectiveness may be limited due to the lack of ongoing user activity and limited resemblance to genuine network resources.

External honeypots may initially succeed in deceiving a threat actor with moderate sophistication. As many such adversaries are indiscriminate, they tend to systematically scan and compromise vulnerable targets. However, the lack of activity (open ports) in the subnet and on the endpoint will likely limit the degree to which the cybercriminal will exploit the honeypot. Based on its intelligence-driven knowledge of possible targets on the network, a nation-state APT may not even attempt to scan the external honeypot, preferring to focus on the secure network perimeter. If it were to scan this zone of the network, the APT is likely to detect the deception since genuine services are rarely hosted in the external network. Thus, external honeypots may allow the defender to learn of and possibly delay the cybercriminal actor's attack but are less likely to detect or otherwise affect a nation-state APT.

3. DMZ Honeypot

Another network placement option is the DMZ, usually located between two logical firewalls that control access to the external (or internet-facing) zone and the intranet. Services in the DMZ (e.g., web servers and mail servers) are accessible from the internet. The DMZ allows untrusted (i.e., unauthenticated) traffic, such as HTTP requests to a web server. DMZ honeypots are possibly the most common honeypot deployment location, striking a balance between risk and deceptive effectiveness relative to external honeypots and intranet honeypots.¹³⁵ Between the external and the internal zones, the DMZ's location offers some degree of protection to the internal network while presenting a range of valuable targets to the attacker. However, a more significant integration effort is necessary

¹³⁵ Javier Franco et al., "A Survey of Honeypots and Honeynets for Internet of Things, Industrial Internet of Things, and Cyber-Physical Systems," *IEEE Communications Surveys & Tutorials* 23, no. 4 (2021): 2351–83.

for the defender to ensure that the honeypot mirrors legitimate resources in the environment without generating any conflicts or security risks, as traffic must be allowed past at least one firewall to access the honeypot.

Due to the proximity of the DMZ honeypot to active network resources, both cybercriminal actors and the nation-state APT are more likely to engage with it than the external honeypot. Ongoing activity in this network zone, such as web traffic and email, signals both attackers to the value of this network as a target. Consequently, the cybercriminal actor is more likely to invest time and resources to exploit the DMZ honeypot, leading to a more sustained engagement. While the APT may ultimately discern the honeypot from a legitimate network resource, the honeypot's more strategic location renders it more difficult for the APT to do so: ongoing network activity in the DMZ creates a more realistic environment for the honeypot. Consequently, even against the APT, the DMZ may slow the APT's attack and increase the chances of detection.

4. Intranet Honeypot

Intranets generally host a network's most sensitive resources and are not typically accessible from the internet, with the notable exception of remote users who have connected through a virtual private network gateway. Intranets may be further subdivided into zones of varying levels of security. All communications within the intranet may also be authenticated and encrypted, pursuant to a zero-trust architecture. Nonetheless, the intranet presents far more opportunities for a sophisticated attacker to identify vulnerabilities pertaining to sensitive network data and operations than the DMZ or the external zones. Thus, intranet-hosted honeypots could also offer a vital foothold to the adversary, presenting a greater security risk than placement in the intranet or DMZ. Finally, such honeypots may require significant procedural effort as change management is often more strictly governed within the internal network.

Intranet honeypots are likely to be even more effective at luring and detecting both the cybercriminal actor and the nation-state APT than DMZ or external honeypots. The cybercriminal actor may engage more persistently with the intranet honeypot due to the enticing nature of high-value targets located within this network zone. Unless the network

has already been compromised, the APT is less likely to possess intelligence regarding the design of the intranet and must rely on careful scanning and lateral movement. When informed by commercial threat intelligence or signals intelligence (for government networks), the defender can present a truly compelling deception that reflects the typical tactics of an APT. Nevertheless, the heightened security risk associated with intranet honeypots should be carefully considered, as a compromise could lead to significant damage, for instance in cases where the attacker gains access to sensitive information or disrupts critical operations. The following section shows how DRACO can mitigate this risk while presenting a compelling deception.

5. DRACO

DRACO offers a fundamentally different approach to deception by simulating an active endpoint in any network zone (external, DMZ, or intranet), potentially achieving deceptive effectiveness benefits similar to each corresponding honeypot location due to the appearance of active use on the endpoint. This deceptive approach may prove particularly potent to the discriminate actor. When DRACO simulates an active endpoint of particular importance to the APT or cybercriminal, these actors are more likely to engage with the deceptive endpoint because it aligns with their interests. For instance, an attacker may be targeting the internet address of an organization's web server, which the attacker can verify is an active endpoint. If DRACO simulates this resource and provides responses that appear authentic, the attacker may be tempted to respond. However, a sophisticated actor such as the nation-state APT may detect the redirection, especially if it benefits from multiple vantage points within the network. If this were the case, the APT could possibly learn that DRACO is not routing the APT's attacks to their intended target. Thus, we evaluate DRACO's deceptive effectiveness to be slightly lower than the intranet honeypot. On the other hand, since DRACO redirects the engagement with the adversary to an external deception network, the proposed concept offers much better security. In addition, DRACO's ease of integration and maintenance is potentially much better than any legacy honeypot location option. Rather than sustain and deconflict the deceptive service with ongoing activity within one of the typical network zones (intranet, DMZ, or external), the defender simply redirects designated attack traffic to the cloud-hosted deception network.

D. CONCLUSION

Our evaluation of different honeypot deployment locations and the DRACO concept shows varying levels of performance that ultimately favor DRACO (Table 2). External honeypots offer high ease of integration and maintenance and a very low security risk due to their isolation from the DMZ and internal network. However, their deceptive effectiveness is also low, due to lack of activity within the internet-facing zone. DMZ honeypots strike a balance, providing medium ease of integration and maintenance, deceptive effectiveness, and a low security risk when considering the separation of the DMZ from the most sensitive network resources. Intranet honeypots, while offering very high deceptive effectiveness, come with a low ease of integration and maintenance. Intranet honeypots also offer a potential foothold to cyber threat actors, representing a medium security risk. DRACO, on the other hand, presents a compelling solution, achieving high ease of integration and maintenance, high deceptive effectiveness, and a very low security risk.

Table 2. Conceptual evaluation of DRACO and typical honeypot location options

Criterion	External Honeypot	DMZ Honeypot	Intranet Honeypot	DRACO
Ease of integration and maintenance	High	Medium	Low	High
Deceptive Effectiveness	Low	Medium	Very High	High
Security Risk	Very Low	Low	Medium	Very Low

Based on this evaluation, DRACO represents a promising option to generate secure and effective engagement with the adversary. As we have discussed in the previous chapter, this engagement can lead to powerful effects on the balance of information in cyber conflict. Further, the “deception as a service” architecture has the potential to minimize the integration and maintenance costs of the solution. The following chapter takes steps toward the quantitative validation of DRACO’s capabilities through experimentation on live networks.

THIS PAGE INTENTIONALLY LEFT BLANK

V. EXPERIMENTATION

DRACO intends to provide a secure and compelling deception by simulating an authentic endpoint while remaining outside of the operator's production network. Optimally, the validation of this concept would involve comparative analysis of attack patterns on an authentic endpoint and one which is simulated by DRACO, while controlling all other variables. If similar attack patterns on both endpoints are observed, then it is reasonable to conclude that DRACO effectively simulated the authentic endpoint. However, as this construct invites the exploitation of an authentic endpoint, no organization would submit to such an experiment. Also, while a narrow attack path and complex deception server (i.e., a traditional honeypot) configuration may appear more attractive to an adversary, we evaluate DRACO using a basic honeypot server configuration without additional features, such as an internal file structure.

Validation that DRACO securely simulates an authentic endpoint is shown via experiments on live networks. The first experiment evaluates DRACO's ability to simulate an active endpoint by comparing attack patterns on a simulated endpoint and on typical research honeypots. In the second experiment, we employ DRACO to study attack patterns on an active and an inactive address within a subnet. The methodology section provides an overview of aspects common to all experiments and is followed by sections which present the specific methodology and results for each experiment. Finally, we state the limitations of our experimentation and derive some conclusions.

A. METHODOLOGY

1. Production Network and Collaboration with the U.S. Naval Postgraduate School

DRACO must be attached to an authentic network to simulate an authentic endpoint. Thus, we conducted these experiments in collaboration with the U.S. Naval Postgraduate School. The research university operates an unclassified network for nearly

20,000 students, faculty, and staff.¹³⁶ This network likely has significant value to criminal and state actors, given the nature of research that is carried out at a U.S. Department of Defense institution. For security reasons, only the high-level architecture of the school's network entry point is presented (Figure 6). The detailed specifications and addresses of network components could be used to plan an attack against the network.




Figure 6. High-level architecture of the U.S. Naval Postgraduate School's network entry point

Before we could conduct the experiment, we needed approval from the Configuration Control Board (CCB). We submitted a short proposal detailing the concept, including a risk assessment, to the CCB. The CCB agreed to the proposal, recognizing that the only significant risk would be reputational in nature: the redirection of traffic from the school to the external deception network would not alter the school's security posture, but threat intelligence providers could interpret the ongoing traffic being routed from the school back to known threat actors as an indicator of compromise. We mitigated this risk by designing short experiments and employing a limited scope of network addresses for experimentation. The CCB authorized both experiments and the information technology department subsequently engineered the operational network component of DRACO to specification.

¹³⁶ Office of Institutional Research, Reporting and Analysis (IRRA), "Naval Postgraduate School Fact Sheet," April 21, 2022, <https://nps.edu/documents/10180/0/NPS+Fact+Sheet-2022+%281%29.pdf>.

2. DRACO: Deception Network Architecture

The deception network manages deception and management traffic. We hosted each instance of these networks on the U.S.-West region of Amazon Web Services (AWS) since the U.S. Naval Postgraduate School is also located on the U.S. West Coast. There are two variations of the deception network: with or without a connection to an operational network. When experimenting with DRACO, the deception network was connected to the U.S. Naval Postgraduate School's network using a virtual private network (Figure 7a). The tunnel supported the AWS default encryption settings.¹³⁷

This variation of the deception network required two subnets: one to host the deception server; a second to integrate a stateless firewall and internet gateway for management traffic. Thus, all attack traffic was routed through the virtual private network gateway, while management traffic was routed through the internet gateway. When the deception network was configured without a connection to the U.S. Naval Postgraduate School, it consisted of a single subnet with an internet gateway (Figure 7b). In this case, the internet gateway routed the attack and management traffic to the intended destination address. Both variations of the deception network operated with a network access list, which we configured to allow attack traffic from any address and to allow administrative traffic from the management terminal.

¹³⁷ "Tunnel Options for Your Site-to-Site VPN Connection – AWS Site-to-Site VPN," accessed February 27, 2023, <https://docs.aws.amazon.com/vpn/latest/s2svpn/VPNTunnels.html>.




Figure 7. Deception network architecture, a) with DRACO and b) without

Within the deception network lies the deception server, which responds to attacks and collects all associated data. The server is based on the open-source threat deception platform T-Pot.¹³⁸ Each deception server is an AWS Elastic Compute Cloud (EC2) instance with 256 GB of storage, 16 GB of memory, and four virtual CPUs. T-Pot offers 23 containerized honeypots and many passive security tools.¹³⁹ We experimented with Cowrie, a honeypot on the Secure Shell (SSH) network communications protocol, and SNARE, a web application honeypot over Hypertext Transfer Protocol (HTTP).^{140 141} Further details on the configuration of these honeypots is provided in sections B and C.

We selected the SSH protocol on the basis of its broad use and value as an attack path. As opposed to more specialized honeypots included with T-Pot (e.g., those pertaining to industrial control systems), SSH is a potential attack vector to almost any

¹³⁸ Telekom Security, “T-Pot – The All in One Honeypot Platform,” accessed December 9, 2021, <https://github.com/telekom-security/tpotce>.

¹³⁹ Telekom Security, “T-Pot – Technical Concept,” February 27, 2023, <https://github.com/telekom-security/tpotce#technical-concept>.

¹⁴⁰ “Cowrie SSH and Telnet Honeypot,” Cowrie, accessed February 26, 2023, <https://www.cowrie.org/>.

¹⁴¹ “SNARE,” Python (2015; repr., MushMush, February 27, 2023), <https://github.com/mushorg/snare>.

UNIX-based operating system. An alluring target to attackers, SSH enables the remote management of the endpoint. Cowrie simulates a Linux (Debian 5.0) shell and accepts any login/password combination after a random number of attempts (between two and five attempts). A typical output of the Cowrie log for a given session may include the following messages:

```
New connection: 127.0.0.1

login attempt [test/asdf] failed

login attempt [test/cool] failed

login attempt [test/password] succeeded

CMD: ls

CMD: exit

Connection lost after 13 seconds
```

We also selected SNARE, the web application honeypot, as it operates over HTTP (port 80) and many endpoints we wished to simulate may be active over HTTPS (port 443). Attackers may be tempted to exploit HTTP if they detect HTTPS activity on the same endpoint. However, HTTP is a much less enticing target than SSH because its adoption has crashed worldwide after being restricted by many browsers.¹⁴² SNARE logs the details of all HTTP requests (e.g., method, headers, etc.). The default configuration of this honeypot presents a login page for an installation of ownCloud, an open-source file sharing platform, to attackers.¹⁴³ SNARE has a companion service, TANNER, which classifies these requests.¹⁴⁴

¹⁴² “HTTPS Encryption on the Web,” Google Transparency Report (Google), accessed March 1, 2023, <https://transparencyreport.google.com/https/overview?hl=en>.

¹⁴³ OwnCloud, “Share Files and Folders, Easy and Secure,” accessed March 2, 2023, <https://owncloud.com/>.

¹⁴⁴ “TANNER,” Python (MushMush, February 27, 2023), <https://github.com/mushorg/tanner>.

Each deceptive service generates logs, which `Logstash` ingests and stores on `Elasticsearch`. `T-Pot` includes `Elasticvue`, a browser-based graphical user interface to query the database.¹⁴⁵ We measured attack patterns by centering on events in the database that originated from `Cowrie` and `TANNER`. The database associated each event to many characteristics including a timestamp, source address, and approximate geolocation (imported from Maxmind’s `GeoLite` database).¹⁴⁶ Hence, we were also able to query the database for events matching a specific source address (`src_ip.keyword:"127.0.0.1"`) or approximate geolocation. By default, `T-Pot` rotates logs every 30 days.

`T-Pot` pre-configures `Kibana` to display a dashboard for each service. These dashboards are useful for data monitoring. In addition, `T-Pot` includes a script (`dps.sh`) which provides the uptime for each containerized service. Since `Cowrie` listens on port 22 for SSH attacks, management access via SSH is bound to port 64295. `Kibana` and `Elasticvue` are bound to port 64297. Finally, we saved the deception network configuration as a launch template on AWS and the deception server as an Amazon Machine Image (AMI), ensuring that each instance of the deception network was identical except for network addresses and that we could rapidly launch them.

B. EXPERIMENT 1: SIMULATING AN ENDPOINT WITH DRACO

The aim of this experiment was to validate `DRACO`’s ability to draw attacks intended for an authentic endpoint. We inferred this ability by comparing the volume of attacks and unique source addresses between a control group and an experimental server which employed `DRACO`. In addition, we formulated conjectures to reveal insights regarding the differences between attack patterns on production honeypots versus research honeypots, based on the assumption that `DRACO` is effective in redirecting attacks that are intended for an authentic endpoint.

¹⁴⁵ “Elasticvue Is a Free and Open-Source Elasticsearch GUI for the Browser,” `Elasticvue`, accessed February 27, 2023, <https://elasticvue.com>.

¹⁴⁶ “GeoIP and GeoLite,” Maxmind, accessed February 27, 2023, <https://dev.maxmind.com/geoip?lang=en>.

1. Design

We designed the control group (Figure 8) to minimize the interest from potential attackers. The three control servers were cloud instances in the U.S. West region with no active use: only the deceptive services could be reached on ports 80 (HTTP/SNARE) and 22 (SSH/Cowrie) by an attacker as the network access list restricted the administrative ports (64297/64295) to the address of the management terminal. Control servers had no ongoing affiliation to a known entity (no banners or other misleading content).




Figure 8. Control Group network architecture

When employing cloud providers, however, one will likely inherit certain “baggage” from previous tenants of the addresses. Prior use (open ports and/or affiliation to a known entity) could draw attackers to continue visiting this address later, even on different ports (vertical scanning). Shodan, a threat intelligence service, provides an address history service (open ports).¹⁴⁷ While the reliability of this service is unknown, it

¹⁴⁷ “Shodan Search Engine,” accessed February 26, 2023, <https://www.shodan.io/dashboard>.

can provide an indication of prior activity. According to Shodan, Control Servers 1 and 3 had no known prior use, while Control Server 2 had known use on port 81 in June 2022. Thus, the control servers may have represented targets of varying interest to potential attackers, but this interest was limited by that lack of ongoing activity or open ports.

By contrast, the DRACO server (Figure 9) was designed to be a compelling target to potential attackers. The server on this network was based on the same image as the control servers and was also located in the U.S. West region. However, the deception network with DRACO could only receive attacks through its connection with the U.S. Naval Postgraduate School. The university employs a virtual private network to encrypt communications with remote users. We selected the gateway to the virtual private network, operating on port 443, as the simulated endpoint for DRACO. Thus, designated attack traffic (SSH connections on port 22 and HTTP requests on port 80) were forwarded to the deception network for the duration of the experiment without affecting regular operations of the gateway.




Figure 9. Architecture of the deception network, DRACO

The gateway has a robust pattern of use. It sees a higher volume of traffic than any other endpoint that is accessible from the internet on the school's on-premises network. The IP address of the virtual private network gateway has not changed in at least eight years according to the administrators, presenting a long and uninterrupted history of use to potential attackers. Further, the endpoint sees approximately 100 GB of traffic per day. If a potential attacker browses to the gateway address on port 443, he will see the authentic login page with the school's logo. The gateway's address is geolocated to Monterey, CA, by GeoLite and registered to California State University. By compromising the gateway, attackers could potentially intercept private communications and gain a foothold within the network. Thus, this endpoint presents many characteristics which should be desirable to cyber threat actors.

The characteristics of the control group and the DRACO server (Table 2) suggest that if the proposed concept effectively draws attacks that were intended for the simulated endpoint, the DRACO server should see significantly more attacks than the control group across both deceptive services (HTTP and SSH). Specifically, the DRACO server has more active use, a more robust history of active use, and is affiliated with a high-profile organization: the U.S. Naval Postgraduate School.

Table 3. Characteristics of the addresses assigned to the DRACO server and the control servers

		DRACO	Control Server 1	Control Server 2	Control Server 3
Geolocation		Monterey, CA	Boardman, OR	Boardman, OR	Boardman, OR
Current Use	Open Ports	22, 80, and 443	22 and 80	22 and 80	22 and 80
	Volume of legitimate traffic	~50 GB per day	None	None	None
	Affiliation	U.S. Naval Postgraduate School, California State University	None	None	None
Prior use (1 year)	Open Ports	443	None	81	None
	Affiliation	U.S. Naval Postgraduate School, California State University	Unknown	Unknown	Unknown

2. Analysis and Results

Over a period of 21 days between January 25 and February 15, 2023, the DRACO server logged more attacks and unique source addresses than any control server. In the following analyses, we assess the significance of these results.

a. Graphical Comparison

On the HTTP protocol, DRACO totals 1,432 unique source addresses, while the control group mean is 993.00 unique source addresses—an increase of 44.21%. Daily observations reflect the consistency of this increase (Figure 10). Further, a box plot shows

no significant outliers (Figure 11). On average, the DRACO server shows a mean of 65.09 unique source addresses while the control group sees 45.14.




Figure 10. Line plot, daily HTTP unique source addresses




Figure 11. Box plot, daily HTTP unique source addresses

DRACO totals 5,703 HTTP requests, or 3.98 requests per unique source address, compared to the control group mean of 4,032.67 requests, or 4.06 requests per unique source address. DRACO logs a daily average of 259.23 HTTP requests in comparison to 183.30 for the control group (Figure 12), although this increase is not consistent. An outlier is present in the distribution of daily HTTP requests, under a box-and-whisker plot (Figure 13). In this case, a single source address was responsible for 846 HTTP requests on

DRACO over a 22-hour period between February 14 and 15, 2023. Even without considering this outlier, DRACO still shows a higher total (4,857) and daily average (220.77) of HTTP requests.




Figure 12. Line plot, daily HTTP requests




Figure 13. Box plot, daily HTTP requests

On the SSH protocol, DRACO totals 5,016 unique source addresses, while the control group mean is 2,744.67 unique source addresses—an increase of 83%. We observe much fewer unique source addresses on Control 1 (1,613) relative to Control 2 and Control 3 (3,357 and 3,264, respectively). On average, DRACO shows a mean of 228.00 daily

source addresses while the control group sees 124.76 (Figure 14). There are two outliers in daily SSH source addresses on DRACO (Figure 15). Not only were more unique source addresses drawn to DRACO, but these attackers are also more engaged: DRACO totals 121,477 new connections, or 24.22 new connections per unique source address, while the control server mean is 20,605.67 new connections, or 2.50 new connections per unique source address.




Figure 14. Line plot, daily SSH unique source addresses




Figure 15. Box plot, daily SSH unique source addresses

Daily SSH connections per server (Figure 16) illustrate the consistency of the increase in new connections on DRACO. On average, DRACO shows a mean of 5521.68 SSH connections while the control group sees 936.62. Several outliers are present in the distribution of SSH connections (Figure 17). A single source address was responsible for 21,199 SSH connections on DRACO over a 16-hour period between February 7 and 8. The attacker continued to attempt new login/password combinations even after successful authentication. This attacker demonstrated far more engagement than major attackers on Control 1 (3445 connections over a seven-hour period on January 26, 2023) and Control 3 (2,305 connections over a four-hour period on February 4, 2023). Even if this outlier is disregarded, the DRACO server still shows more total and daily average SSH connections.




Figure 16. Line plot, daily SSH connections




Figure 17. Box plot, daily SSH connections

Our graphical analysis shows a noticeable and consistent increase in attacks and unique source addresses on the DRACO server over the controls. Furthermore, differences among controls are consistently lower than differences between control and DRACO observations, as seen for SSH unique source addresses in Figure 14 and HTTP unique source addresses in Figure 10.

b. Statistical Analysis

Our results suggest that DRACO attracted far more attacks than the control servers. However, for rigor we take a statistical approach to assess significance of this difference.

We first assess time dependency of the distributions. If the data are dependent on the time series (day of the week), then that dependency would need to factor into further assessments. Figures 18–21 display the autocorrelation results on each set of control observations as well as for DRACO. As seen, in only select cases (HTTP address for Control 2, HTTP requests for DRACO, and SSH addresses for Control 1) show a measurement outside of the non-significant (blue) region. The markers indicate the lag interval considered (i.e., the first checking of the autocorrelation from one day to another, the second for every two days, etc.). By and large, these results indicate a lack of significant autocorrelation, (i.e., a lack of time dependence).




Figure 18. ACF plot, daily HTTP unique source addresses




Figure 19. ACF plot, daily HTTP requests




Figure 20. ACF plot, daily SSH unique source addresses




Figure 21. ACF plot, daily SSH connections

Based on the results just described, from this point onward we assume time independence of the data. Given the lack of time dependence, it is possible to organize data by the frequency of attacks, which allows us to compare the overall distributions of control and DRACO data based on those numbers.

We have two primary goals in this assessment. First is to cross compare the raw attack frequencies from the Control and DRACO data sets. Since there is time independence, we look at the data points from the respective Control and DRACO data sets as indistinct measurements of attacks/day without ordering. Based on the assumption that the Control samples come from the same population, we can pool the Control data, allowing us to cross compare a single Control data set of size 66 with a single DRACO data set of size 22.

The second goal is comparison of the differences between the Control and DRACO data sets, namely a difference in difference (DID) comparison. This allows us to take a closer look at the gaps between Controls 1, 2, and 3 and DRACO and assess the relative significance of those gaps. Thus, we create a data set of size 1,452 from the absolute value of differences among Control observations and a second data set of size 1,452 from the absolute value of differences between Control and DRACO observations. These two comparisons (raw Control versus raw DRACO observations and DID Control versus Control/DRACO observations) for our analysis basis.

We first test normality of the individual distributions with the Shapiro-Wilk test, which guides selection of the correct tests on the distribution of the samples.¹⁴⁸ This test, which research suggests is the most powerful normality test, is performed on both the Control and DRACO data sets individually, as well as the pooled Control data set.¹⁴⁹ Additionally, we perform the normality test on the Control differences and Control/DRACO differences to set the stage for the DID test. The Shapiro-Wilk results (Table 4)

¹⁴⁸ Samuel Sanford Shapiro and Martin B Wilk, “An Analysis of Variance Test for Normality (Complete Samples),” *Biometrika* 52, no. 3/4 (1965): 591–611.

¹⁴⁹ Nornadiiah Mohd Razali, Yap Bee Wah, and others, “Power Comparisons of Shapiro-Wilk, Kolmogorov-Smirnov, Lilliefors and Anderson-Darling Tests,” *Journal of Statistical Modeling and Analytics* 2, no. 1 (2011): 21–33.

indicate that the data are not normally distributed for SSH connections and HTTP requests (at the $\alpha = 0.05$ significance level), but there are not sufficient grounds to assume non-normality for unique source addresses for either SSH or HTTP.

Table 4. Results of Shapiro-Wilk test for normality on daily statistics

Sample	Test Statistic	p-value	Conclusion
Daily SSH New Connections			
DRACO	0.49	< 0.0001	Data are not normally distributed
Pooled Control	0.83	< 0.0001	Data are not normally distributed
DRACO Differences	0.87	< 0.0001	Data are not normally distributed
Control Differences	0.50	< 0.0001	Data are not normally distributed
Daily SSH Unique Source Addresses			
DRACO	0.94	0.2371	Insufficient evidence to assume non-normality
Pooled Control	0.97	0.1666	Insufficient evidence to assume non-normality
DRACO Differences	0.99	0.0012	Data are not normally distributed
Control Differences	0.96	< 0.0001	Data are not normally distributed
Daily HTTP Requests			
DRACO	0.57	< 0.0001	Data are not normally distributed
Pooled Control	0.84	< 0.0001	Data are not normally distributed
DRACO Differences	0.59	< 0.0001	Data are not normally distributed
Control Differences	0.86	< 0.0001	Data are not normally distributed
Daily HTTP Unique Source Addresses			
DRACO	0.95	0.3450	Insufficient evidence to assume non-normality
Pooled Control	0.98	0.3450	Insufficient evidence to assume non-normality
DRACO Differences	0.99	< 0.0001	Data are not normally distributed
Control Differences	0.85	< 0.0001	Data are not normally distributed

We assume that the DRACO and Control servers are independent samples. Each server is located on a different network subnet in a very large space of addresses. Given the non-normality of SSH connections and HTTP requests across all categories (Controls, DRACO, Control differences, Control/DRACO differences), we apply non-parametric tests appropriate for independent samples to assess whether the distributions are equivalent. With SSH and HTTP unique source addresses, we select a parametric test for independent samples.

- **Mann-Whitney U test.**¹⁵⁰ This non-parametric test on independent samples is employed when comparing the distributions of raw SSH connections and HTTP requests. In this test, the null hypothesis is that the distributions being compared are equivalent, and the alternative hypothesis is that the difference in distributions is greater than zero.
- **Permutation test.**¹⁵¹ When samples are large and of unequal size, research has shown that the Mann-Whitney test may be less accurate if there are unequal variances between the samples.¹⁵² This non-parametric test on independent samples is employed with $N_{perm} = 100,000$ when comparing the distributions between the DID DRACO and Control samples. In this test, the null hypothesis is that the data are sampled from equivalent distributions.
- **Welch's t-test.**¹⁵³ This parametric test on independent samples with unequal variances is employed when comparing the distributions of raw SSH and HTTP unique source addresses. In this test, the null hypothesis is that the means of the distributions are equal, and the alternative hypothesis is that the difference of means is greater than zero.

The results shown in Table 5 indicate a significant difference between Control results and DRACO results, whether in the comparison of raw attack observations or in a DID analysis of the gaps between the distributions.

¹⁵⁰ Henry B. Mann and Donald R. Whitney, "On a Test of Whether One of Two Random Variables Is Stochastically Larger than the Other," *The Annals of Mathematical Statistics* 18, no. 1 (1947): 50–60.

¹⁵¹ Ronald Aylmer Fisher, "The Design of Experiments.," *British Medical Journal* 1, no. 3923 (1936): 554.

¹⁵² Donald W. Zimmerman, "A Warning about the Large-Sample Wilcoxon-Mann-Whitney Test," *Understanding Statistics* 2, no. 4 (2003): 267–80.

¹⁵³ Bernard L. Welch, "The Generalization of 'STUDENT'S' Problem When Several Different Population Variances Are Involved," *Biometrika* 34, no. 1–2 (1947): 28–35.

Table 5. Results of distribution tests on daily statistics

Test	Sample	Test Statistic	p-value	Conclusion
Daily SSH Connections				
Mann-Whitney	DRACO vs Pooled Control	1401.00	< 0.0001	Distributions are significantly different
Permutation	DRACO Differences vs Control Differences	3740.75	0.0002	Distributions are significantly different
Daily SSH Unique Source Addresses				
Welch's T-Test	DRACO vs Pooled Control	9.68	< 0.0001	Means are significantly different
Permutation	DRACO Differences vs Control Differences	43.17	0.0002	Distributions are significantly different
Daily HTTP Requests				
Mann-Whitney	DRACO vs Pooled Control	965.50	0.0106	Distributions are significantly different
Permutation	DRACO Differences vs Control Differences	43.49	0.0002	Distributions are significantly different
Daily HTTP Unique Source Addresses				
Welch's T-Test	DRACO vs Pooled Control	12.03	< 0.0001	Means are significantly different
Permutation	DRACO Differences vs Control Differences	13.01	0.0002	Distributions are significantly different

c. Conjectures

The previous results provide an indication of DRACO's ability to redirect attacks intended for an authentic endpoint, which in turn suggests that DRACO is effective at simulating authentic endpoints. Based on the assumption that DRACO is effective, we further analyze the data from Experiment 1 to assess the differences in attack patterns between an active endpoint (i.e., the U.S. Naval Postgraduate School virtual private network gateway) and inactive endpoints (i.e., the control group).

Geolocation of source IP addresses should be taken with a grain of salt given the ease with which cyber threat actors can move infrastructure across borders. For instance, Cisco Talos reported in June 2022 that a cyber threat actor was leveraging cloud infrastructure in Europe and North America to spread malware.¹⁵⁴ Notably, the list of top countries by proportion of source addresses which attacked DRACO (Figure 22) was missing an obvious entry: China only accounts for 1% of addresses. By contrast, China-based addresses accounted for 2% of the total on Control 1, and 6% of the total on Control 2 and Control 3. This limited sample may suggest that Chinese actors take additional

¹⁵⁴ Chetan Raghuprasad and Vanja Svajcer, "Nanocore, Netwire and AsyncRAT Spreading Campaign Uses Public Cloud Infrastructure," *Cisco Talos Blog* (blog), January 12, 2022, <https://blog.talosintelligence.com/nanocore-netwire-and-asyncrat-spreading/>.

precautions when attacking U.S. government infrastructure, although concrete attribution is outside of the scope of this work.




Figure 22. Top countries by proportion of estimated geolocation of source addresses on the DRACO server

Could attackers targeting the U.S. Naval Postgraduate School also be more deliberate in planning their attacks outside of normal working hours? Since we noted in the previous section that there had been single addresses responsible for a large volume of attacks, we focus on unique source addresses rather than new connections. Data from Experiment 1 show no clear indication of a difference in attack patterns throughout the day, or between the control group mean and the DRACO server (Figure 23).




Figure 23. Proportion of unique addresses by hour of the day for each server

3. Conclusion

Experiment 1 provides an initial indication that DRACO may effectively redirect attacks intended for an authentic endpoint. As predicted, we observe significantly more attacks and unique source addresses on the DRACO server relative to the control group across both protocols. This conclusion depends on the assumption that if the authentic U.S. Naval Postgraduate School virtual private network gateway were exposed to SSH attacks, it would likely observe more attacks and unique source addresses than a control group which consists of deception servers with no active use or affiliation to an entity and a variable history of prior use. If the results of this experiment can be replicated while controlling more variables, more evidence could be gathered to evaluate the assumption (and this conclusion).

C. EXPERIMENT 2: COMPARISON OF ATTACK PATTERNS ON ACTIVE AND INACTIVE ENDPOINTS

The aim of this experiment is to analyze the differences in attack patterns between active and inactive addresses on a network. If statistically significant differences are identified, they will inform cybersecurity research conducted on inactive networks (also known as dark networks or network telescopes). Experiment 2 rests on the assumption that DRACO is effective at simulating inactive and active addresses on a network.

1. Design

We configure two DRACO deception networks, one simulating an active endpoint and another simulating an inactive endpoint (Figure 24). In this case, an active endpoint is one which hosts at least one service for legitimate users of the network (i.e., an attacker will receive a response on more ports than just the port(s) corresponding to the deception). Conversely, an inactive endpoint is one which will only respond on the port(s) corresponding to the deception. DRACO simulates a different active endpoint than in Experiment 1 to avoid an influence from prior use of the endpoint as a deception. We select the address of the central authentication service, which authenticates users to access various public-facing resources. The volume of traffic on this server is approximately 1.5 GB per day, much lower than the virtual private network gateway we selected in Experiment 1. Nonetheless, there are many years of prior active use on port 443 and port 80. Due to active use on port 80, we configure the deception server only for the SSH protocol and not the HTTP protocol.




Figure 24. Architecture of the active and inactive simulated endpoints

2. Analysis and Results

Over a period of 21 days between February 23 and March 15, 2023, the simulated active endpoint logged more connection attempts, unique source addresses, and commands issued than the simulated inactive endpoint. We analyze total statistics and the correlation between daily attack patterns.

a. Total Statistics

The active endpoint totals 74,456 connection attempts, 3,154 unique source addresses, and 103,832 commands issued, while the inactive endpoint totals 71,962 new connections, 3,027 unique source addresses, and 64,801 commands (Figure 25), with a single observation from each group and variable.




Figure 25. Total new connections, unique source addresses, and commands on the active and inactive simulated endpoints

While the data for connection attempts and unique source addresses are similar (a difference of 3.47% and 4.20% in favor of the active endpoint), the active endpoint logs 60.23% more commands on the simulated SSH session. This considerable gap may indicate that attackers are more inclined to exploit an endpoint which is active on multiple ports.

b. Daily Statistics

Daily statistics are consistent with the totals in terms of the relative differences in daily averages for new connections, unique source addresses, and commands issued (Figures 26–28). On average, the active endpoint sees 3,545.52 connections, 208.09 unique source addresses, and 4,944.38 commands issued on the simulated SSH session. On average, the inactive endpoint sees 3,426.76 new connections, 197.10 unique source addresses, and 3,085.76 commands issued on the simulated SSH session. Thus, the active endpoint averages 3.47% more new connections, 4.20% more unique source addresses, and 60.23% more commands issued per day.




Figure 26. Line plot, daily SSH connections




Figure 27. Line plot, daily SSH unique source addresses




Figure 28. Daily commands issued by attacks on the simulated SSH shell

We turn to the analysis of attacks by hour of the day. SSH connections over the duration of the experiment are binned by the hour of the day (in the local time zone) in which they occur. While it is possible that some attackers may plan their attacks based on the silent hours when an organization's network operations center is unmanned, the time zones of the attackers may also be an important factor. Graphical analysis of these patterns on the active and inactive endpoints (Figure 29) suggests no clear pattern or difference is discernible.




Figure 29. Proportion of unique source addresses by time of the day (local time)

Finally, we examine the estimated geolocation of source addresses with Maxmind's GeoLite database. Analysis of geolocation data should be taken with a grain of salt due to the ease with which attackers can leverage network infrastructure beyond their borders. This metric only relates to the final egress point of the attack and does not permit the attribution of an attack to a specific country. The top ten countries by estimated geolocation of source addresses is identical across both simulated endpoints, with only a few minor differences in the ordered list, as shown in Figure 30 and Table 6.




Figure 30. Top countries by proportion of estimated geolocation of source addresses

Table 6. Sources IP address by estimated geolocation

Active Endpoint		Inactive Endpoint	
Country	Source Addresses	Country	Source Addresses
United States	782	United States	739
South Korea	397	South Korea	376
Taiwan	158	Taiwan	148
Japan	146	Japan	137
India	122	India	127
China	111	Singapore	116
Russia	107	China	107
Brazil	99	Germany	90
Singapore	98	Russia	88
Germany	94	Brazil	81

c. Statistical Analysis

To select the appropriate method to measure the statistical significance of the results, we first determine if the data follow a normal distribution with the Shapiro-Wilk

test. Results suggest that there are no grounds to assume non-normality in daily connections (active and inactive endpoints) and commands (active endpoint only) at the specified significance level ($\alpha = 0.05$). However, we can reject the null hypothesis of a normal distribution for daily source addresses and commands on the inactive endpoint (Table 7).

Table 7. Results Shapiro-Wilk test results for non-normality

Sample	p-value	Conclusion
Active Connections	0.070	Insufficient evidence to assume non-normality
Inactive Connections	0.231	Insufficient evidence to assume non-normality
Active Source Addresses	0.015	Data are not distributed normally
Inactive Source Addresses	0.012	Data are not distributed normally
Active Commands	0.413	Insufficient evidence to assume non-normality
Inactive Commands	0.018	Data are not distributed normally

We also assume that the active and inactive samples for a given variable are dependent. Contrary to the previous experiment where all samples originated from different address ranges, both servers in this experiment reside on the same subnet (with mask /16) on the organization's network. We consider that the presence of an active endpoint could affect scanning and attack patterns against neighboring addresses.¹⁵⁵ Owing to the dependence of the samples and the results of normality testing, we select the following tests:

- **Wilcoxon Signed-Rank Test.**¹⁵⁶ This non-parametric test is employed to compare distributions for variables where we reject the assumption of normality for at least one of the two samples (daily source addresses and commands). In this test, the null hypothesis is that the two samples come

¹⁵⁵ Richter and Berger, "Scanning the Scanners: Sensing the Internet from a Massively Distributed Network Telescope."

¹⁵⁶ Frank Wilcoxon, "Individual Comparisons by Ranking Methods," *Biometrics Bull.* 1 (1945): 80–83.

from the same distribution, and the alternative hypothesis is that the difference between the samples is greater than zero.

- **Paired t-Test.**¹⁵⁷ This parametric test is employed to compare the distributions for variables where we assume normality for both samples (daily connections). In this test, the null hypothesis is that the means of the samples are equal, and the alternative hypothesis is that the difference between the samples is greater than zero.
- **Chi-Square Test of Independence.**¹⁵⁸ This non-parametric test compares the frequency of a variable measured at a nominal level. The null hypothesis is that there is no association between the two variables. We employ this test to compare the frequencies of top countries by estimated geolocation of source addresses and attack frequencies by hour of the day.

Distribution tests suggest that we can assume that the distributions of daily connections and commands are significantly different ($p = 0.0001$ with the Paired t-test for connections and $p < 0.0001$ with the Wilcoxon tests for commands). For daily unique source addresses, we have insufficient grounds to assume that the means of the inactive and active samples are different ($p = 0.0895$ with the Wilcoxon test). A significant increase in the daily unique source addresses supports the notion that attackers are more likely to engage with active endpoints. Further, the robust increase in daily commands implies that these attackers are significantly more engaged with active terminals.

On the frequency of attacks by hour of the day, the chi-square test of independence indicates that we can reject the null hypothesis ($p = 0.0005$), stating that there is no relationship between the nominal variables. This result confirms the intuition from the graphical analysis that there is no clear trend in hourly patterns that could inform network defenses.

¹⁵⁷ Henry Hsu and Peter A. Lachenbruch, "Paired t Test," in *Encyclopedia of Biostatistics* (Wiley Online Library, 2005).

¹⁵⁸ Mary L. McHugh, "The Chi-Square Test of Independence," *Biochemia Medica* 23, no. 2 (June 15, 2013): 143–49, <https://doi.org/10.11613/BM.2013.018>.

Chi-square results on estimated geolocation of source addresses returns a result ($p = 0.8206$). Thus, we can assume that the geolocation is different on active servers and inactive servers. While the top ten countries were identical for the active and inactive servers, this result suggests that the difference in their respective proportions is an additional differentiator when comparing attack patterns on active and inactive endpoints.

3. Summary

Experiment 2 demonstrates DRACO's potential for cybersecurity research and reinforces the need to study attack patterns on active endpoints. Our results indicate that there are similarities between the attack patterns on inactive endpoints (i.e., those with no ports open other than the deceptive services) and active endpoints (i.e., those with some legitimate activity on more ports than the deceptive services). Notably, the differences between the number of daily connections, daily source addresses and estimated geolocation of source addresses are not statistically significant. However, attackers were far more engaged in their attempts to compromise the active endpoint, as measured by the number of commands issued on the simulated SSH session. In addition, the difference in distribution of attacks by hour of the day is also statistically significant. Future research with greater sample sizes across different networks can further clarify the differences between attack patterns on active and inactive endpoints. Nonetheless, researchers who study inactive networks (also known as dark networks or network telescopes) may wish to consider these potential differences when mapping their findings to cybersecurity considerations on active networks.

D. LIMITATIONS

Limitations to DRACO's initial validation may be illustrated by contrasting the ideal experiment with the context of this study, as previously mentioned. Although the control servers in experiment 1 were configured identically and located in the same geographical region, they may present different patterns of prior use. This could explain the discrepancy between control servers, though the results remain statistically significant. With regard to DRACO's application as a cybersecurity research tool in the second experiment, the main limitation stems from the small sample size. We experimented with

a single active endpoint and a single inactive endpoint. In this case, the limitation is not conceptual but rather practical: with limited resources and time to conduct the experiment, we were simply not able to generate a larger sample size. Nonetheless, the significance of the results, specifically in commands issued on the active and inactive addresses, strongly suggest a difference in attacker behavior.

E. CONCLUSION

This research provides an early indicator that DRACO can simulate authentic endpoints for applications in research and network defense. We compare attack patterns between the DRACO-enabled deception network and a control group, with results showing that DRACO attracts more new connections and unique source addresses than the control servers, in the tested case. Our findings hold across HTTP and the SSH deception services, and even if outliers are removed that may have favored the DRACO server. Furthermore, in our second experiment, we employ DRACO to compare attack patterns between active and inactive endpoints. Results show similar total and daily new connections and unique source addresses on the inactive and active endpoints, but attackers issue far more commands on the active endpoint's simulated SSH session.

More research is necessary to test the full efficacy and potential of the DRACO deception network to assess threats to live networks. Future steps include comparison with an Intranet honeypot control or multiple instantiations to increase the sample size of DRACO deception servers and control servers. Nonetheless, our work provides an initial indication that DRACO may 1) effectively redirect attacks intended for an authentic endpoint and 2) be useful for targeted research regarding specific threats to the production network.

THIS PAGE INTENTIONALLY LEFT BLANK

VI. CONCLUSION

This thesis began with a brief discussion of military deception, a practice with a long and successful track record in physical operations. It is unsurprising that our findings reveal that the value of deception applies also to virtual operations.

A. SUMMARY

In a qualitative analysis, we identified the potential effects of this technology. The success of Q-ships against U-boats during the First World War demonstrates the power of deceptive defense when combatting a concealed and opportunistic enemy. Specifically, we identified three potential effects of deceptive cyber defense: *identification*, *misdirection*, and *retaliation*.

Through *identification*, deceptive cyber defense may improve one's general understanding of the threat environment and enable the specific detection of threats against a network. Deceptive cyber defense can also *misdirect* the adversary, drawing the adversary's focus away from vital network resources and potentially delaying the adversary's offensive campaign. *Retaliation*, the union of defensive and offensive operations, is greatly enabled by deceptive cyber defense. When deception is effective, an adversarial engagement may reveal information about the adversary that is vital to offensive planners, in addition to presenting opportunities to bypass the adversarial network perimeter.

Although deceptive cyber defense shows promise, current approaches offer limited deceptive effectiveness. Typical honeypot deployments are hindered by the security-deception trade-off, presenting the unacceptable options of deception that is weak but secure, or effective but insecure. For example, internal honeypots may present an effective deception due to the target-rich environment within the network zone, yet they represent a potentially significant security risk. Even with configuration control processes and security precautions in place, intranet honeypots offer a potential foothold to the adversary deep within a network. Deception technology can also represent a significant integration burden for network defenders.

To balance these extremes and suggest a viable path towards realistic and secure cyber deception we presented DRACO—a unique honeypot placement and integration approach. DRACO leverages self-spoofing and redirection so that attack traffic can be handled away from the operational network while the attack receives network traffic that appears to have originated from their intended target. In contrast to the honeypot deployment options just mentioned, DRACO offers compelling deception while minimizing security risks and reducing the integration burden for the defender. We conceptually evaluated DRACO and typical honeypot placement options against two cyber actor profiles and concluded that the proposed concept may present deceptive benefits derived by internal honeypot deployments while generating a lower security risk and integration burden.

We quantitatively evaluated DRACO through experimentation on the U.S. Naval Postgraduate School network. In a first experiment, we tested DRACO’s ability to redirect attacks that are intended for an authentic endpoint. To do so, we compared attack patterns on an experimental server with DRACO, which received attack traffic directed against the U.S. Naval Postgraduate School, with those on control servers, with no ongoing activity or association to an organization. Over the three-week experiment, we found a significant difference in daily attack patterns based on hypothesis testing against raw and difference-in-differences data. In a second experiment, we employed DRACO to study the differences in attack patterns on active and inactive addresses within an operational network. Notably, we found a significant difference in the number of commands issued by attacks on a simulated shell and inferred that attacks were more engaged against active endpoints. Our work provides an initial indication that DRACO may: 1) effectively redirect attacks intended for an authentic endpoint and 2) be useful for targeted research regarding specific threats to the production network.

B. FUTURE WORK

Deceptive cyber defense includes a wide range of techniques and associated effects. While our work aims to optimize the honeypot, we only briefly touched on other means of deception, such as delaying tactics and network concealment. From a strategic standpoint,

future research should continue to build on other means of deceptive cyber defense as well.¹⁵⁹ Within the research domain of honeypots, we reported that there is a lack of literature to inform operationalization, and this remains the case (as of Spring 2023).¹⁶⁰

Our experiments have also been limited in terms of scope and scale. Future experiments with larger sample sizes across different organizations and types of networks could inform the generalizability of our conclusions. In addition, DRACO is fundamentally a network architecture, so we tested it against different network architectures. The next logical step is to experiment with different deceptive protocols and levels of realism within the simulated network environment. Such research could employ DRACO to compare the effects of elaborate folder and file structures, false traffic generation, or other deception enhancements. In turn, these results would inform where defenders should invest their limited time and resources.

¹⁵⁹ Rowe and Rrushi, *Introduction to Cyberdeception*.

¹⁶⁰ Franco et al., “A Survey of Honeypots and Honeynets for Internet of Things, Industrial Internet of Things, and Cyber-Physical Systems.”

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

- Amado, Rafael. “WannaCry: An Analysis of Competing Hypotheses.” *Digital Shadows*, May 18, 2017. <https://www.digitalshadows.com/blog-and-research/wannacry-an-analysis-of-competing-hypotheses/>.
- Antonioli, Daniele, Anand Agrawal, and Nils Ole Tippenhauer. “Towards High-Interaction Virtual ICS Honey pots-in-a-Box.” In *Proceedings of the 2nd ACM Workshop on Cyber-Physical Systems Security and Privacy*, 13–22, 2016.
- Arquilla, John. *Bitskrieg: The New Challenge of Cyberwarfare*. 1st edition. Cambridge, UK and Medford, MA: Polity, 2021.
- . “Cyberwar Is Already Upon Us.” *Foreign Policy* (blog), February 27, 2012. <https://foreignpolicy.com/2012/02/27/cyberwar-is-already-upon-us/>.
- . “The Strategic Implications of Information Dominance.” *Strategic Review* 22 (1994): 22–34.
- Arquilla, John, and David Ronfeldt. “Cyberwar Is Coming!” *Comparative Strategy* 12, no. 2 (1993): 141–65.
- Auten, Harold. “*Q*” *Boat Adventures: The Exploits of the Famous Mystery Ships*. London: H. Jenkins, 1919.
- AWS. “Tunnel Options for Your Site-to-Site VPN Connection – AWS Site-to-Site VPN.” Accessed February 27, 2023. <https://docs.aws.amazon.com/vpn/latest/s2svpn/VPNTunnels.html>.
- Bailey, Thomas A. “The Sinking of the Lusitania.” *The American Historical Review* 41, no. 1 (1935): 54–73. <https://doi.org/10.2307/1839355>.
- Barniuk, Chris. “AlphaBay and Hansa Dark Web Markets Shut Down.” BBC News, July 20, 2017. <https://www.bbc.com/news/technology-40670010>.
- Belqruch, Amine, and Abdelilah Maach. “SCADA Security Using SSH Honey pot.” In *Proceedings of the 2nd International Conference on Networking, Information Systems & Security*, 1–5. Morocco: Association for Computing Machinery, 2019.
- Benson, Karyn, Alberto Dainotti, Kc Claffy, Alex C Snoeren, and Michael Kallitsis. “Leveraging Internet Background Radiation for Opportunistic Network Analysis.” In *Proceedings of the 2015 Internet Measurement Conference*, 423–36. Tokyo: Association for Computing Machinery, 2015.
- Bittman, Ladislav. *The Deception Game: Czechoslovak Intelligence in Soviet Political Warfare*. Syracuse, NY: Syracuse University Research Corp., 1972.

- Borghard, Erica D., and Shawn W. Loneragan. "Cyber Operations as Imperfect Tools of Escalation." *Strategic Studies Quarterly* 13, no. 3 (2019): 122–45. <https://www.jstor.org/stable/26760131>.
- Bridgland, Tony. *Sea Killers in Disguise: The Story of the Q-Ships and Decoy Ships in the First World War*. Annapolis, MD: Naval Institute Press, 1999.
- Campbell, Gordon. *My Mystery Ships*. Garden City, NY: Doubleday, Doran & Company, 1929.
- Canadian Centre for Cyber Security. "An Introduction to the Cyber Threat Environment," October 28, 2022. <https://www.cyber.gc.ca/en/guidance/introduction-cyber-threat-environment>.
- Chatterton, Edward Keble. *Q-Ships and Their Story*. Boston: Sidgwick and Jackson, Limited, 1922.
- Cheswick, Bill. "An Evening with Berferd in Which a Cracker Is Lured, Endured, and Studied," 1991. <https://www.cheswick.com/ches/papers/berferd.pdf>.
- Clark, David D., and Susan Landau. "The Problem Isn't Attribution: It's Multi-Stage Attacks." In *Proceedings of the Re-Architecting the Internet Workshop*, 1–6. Philadelphia, PA: Association for Computing Machinery, 2010.
- . "Untangling Attribution." In *Proceedings of A Workshop on Deterring Cyberattacks: Informing Strategies and Developing Options for U.S. Policy*, 25–40. Washington, D.C.: The National Academies Press, 2011.
- Clausewitz, Carl von (Edited and translated by Michael Eliot Howard Peter Paret). *On War*. Princeton, NJ: Princeton University Press, 1989.
- Compton-Hall, Richard. *Submarines at War 1914–1918*. Cornwall, UK: Periscope Publishing Ltd., 2004.
- Cowrie. "Cowrie SSH and Telnet Honeypot." Accessed February 26, 2023. <https://www.cowrie.org/>.
- CrowdStrike. "Honeypots in Cybersecurity Explained," March 9, 2022. <https://www.crowdstrike.com/cybersecurity-101/honeypots-in-cybersecurity-explained/>.
- . "What Is Cyber Threat Hunting?," April 17, 2022. <https://www.crowdstrike.com/cybersecurity-101/threat-hunting/>.
- Crowley, Michael. "Airman Charged in Leak of Classified Documents." *The New York Times*, April 14, 2023. <https://www.nytimes.com/2023/04/14/us/politics/jack-teixeira-classified-documents-leak.html>.

- Cyber National Mission Force Public Affairs. “Before the Invasion: Hunt Forward Operations in Ukraine.” U.S. Cyber Command, November 28, 2022. <https://www.cybercom.mil/Media/News/Article/3229136/before-the-invasion-hunt-forward-operations-in-ukraine/>.
- Demarest, Colin. “Pentagon Seeks \$11.2 Billion for Cyber in FY23 Budget Request.” C4ISRNet, March 28, 2022. <https://www.c4isrnet.com/cyber/2022/03/28/pentagon-seeks-112-billion-for-cyber-in-fy23-budget-request/>.
- Denning, Peter J. “The Science of Computing: The ARPANET after Twenty Years.” *American Scientist* 77, no. 6 (1989): 530–34. <https://www.jstor.org/stable/27856002>.
- Dönitz, Karl. *Memoirs Ten Years and Twenty Days*. Chicago, IL: Frontline Books, 2012.
- Elasticvue. “Elasticvue Is a Free and Open-Source Elasticsearch GUI for the Browser.” Accessed February 27, 2023. <https://elasticvue.com>.
- Fan, Wenjun, Zhihui Du, Max Smith-Creasey, and David Fernández. “HoneyDOC: An Efficient Honeypot Architecture Enabling All-Round Design.” *IEEE Journal on Selected Areas in Communications* 37, no. 3 (2019): 683–97. <https://doi.org/10.1109/JSAC.2019.2894307>.
- FireEye. “UNC215: Spotlight on a Chinese Espionage Campaign in Israel.” Accessed August 29, 2021. <https://www.fireeye.com/blog/threat-research/2021/08/unc215-chinese-espionage-campaign-in-israel.html>.
- Fischerkeller, Michael P., and Richard J. Harknett. “Cyber Persistence Theory, Intelligence Contests and Strategic Competition.” Institute for Defense Analyses, June 1, 2020. <https://apps.dtic.mil/sti/pdfs/AD1118679.pdf>.
- Fisher, Ronald Aylmer. “The Design of Experiments.” *British Medical Journal* 1, no. 3923 (1936): 554.
- Franco, Javier, Ahmet Aris, Berk Canberk, and A. Selcuk Uluagac. “A Survey of Honeypots and Honeynets for Internet of Things, Industrial Internet of Things, and Cyber-Physical Systems.” *IEEE Communications Surveys Tutorials* 23, no. 4 (2021): 2351–83. <https://doi.org/10.1109/COMST.2021.3106669>.
- Fuchs, Mathias, and Josh Lemon. *SANS 2022 Threat Hunting Survey*. SANS Institute Whitepaper. Rockville, MD: SANS Institute, July 2022. <https://www.sans.org/white-papers/sans-2022-threat-hunting-survey/>
- Gartzke, Erik. “The Myth of Cyberwar: Bringing War in Cyberspace Back Down to Earth.” *International Security* 38, no. 2 (October 2013): 41–73. https://doi.org/10.1162/ISEC_a_00136.

- Gartzke, Erik, and Jon R. Lindsay. "Weaving Tangled Webs: Offense, Defense, and Deception in Cyberspace." *Security Studies* 24, no. 2 (April 3, 2015): 316–48. <https://doi.org/10.1080/09636412.2015.1038188>.
- Gibson, R. H., and Maurice Prendergast. *The German Submarine War 1914–1918*. Cornwall, UK: Periscope Publishing Ltd., 2002.
- GitHub. "SNARE." Python. 2015. Reprint, MushMush, February 27, 2023. <https://github.com/mushorg/snare>.
- . "TANNER." Python. MushMush, February 27, 2023. <https://github.com/mushorg/tanner>.
- Google. "HTTPS Encryption on the Web." Google Transparency Report. Accessed March 1, 2023. <https://transparencyreport.google.com/https/overview?hl=en>.
- Greengard, Samuel. *The Internet of Things, Revised and Updated Edition*. Cambridge, Massachusetts: MIT Press, 2021.
- Guarnizo, Juan David, Amit Tambe, Suman Sankar Bhunia, Martín Ochoa, Nils Ole Tippenhauer, Asaf Shabtai, and Yuval Elovici. "Siphon: Towards Scalable High-Interaction Physical Honeypots." In *Proceedings of the 3rd ACM Workshop on Cyber-Physical System Security*, 57–68. Abu Dhabi, 2017.
- Hall, Mary T. "False Colors and Dummy Ships: The Use of Ruse in Naval Warfare." *Naval War College Review* 42, no. 3 (1989): 52–62. <https://www.jstor.org/stable/44636917>.
- Handel, Michael I. "Sun Tzu and Clausewitz: The Art of War and on War Compared:" Fort Belvoir, VA: Defense Technical Information Center, January 1, 1991. <https://doi.org/10.21236/ADA239084>.
- Hardach, Gerd. *The First World War, 1914–1918*. Oakland, CA: University of California Press, 1981.
- Harris, Shane, Karen DeYoung, Isabelle Khurshudyan, Ashley Parker, and Liz Sly. "Road to War: U.S. Struggled to Convince Allies, and Zelensky, of Risk of Invasion." *Washington Post*, August 16, 2023. <https://www.washingtonpost.com/national-security/interactive/2022/ukraine-road-to-war/>.
- Hartcup, Guy. *Camouflage: The History of Concealment and Deception in War*. Barnsley, South Yorkshire, UK: Pen and Sword Military, 2008.
- Hesketh, Roger. *Fortitude: The D-Day Deception Campaign*. New York, NY: Abrams, 2002.

- Hodge, Nathan, Tara John, and Anna Chernova. “Ukraine: Russia and Belarus Hold Joint Military Exercises.” CNN, February 10, 2022. <https://www.cnn.com/2022/02/10/europe/ukraine-russia-news-thursday-military-exercises-intl/index.html>.
- Hollis, David. “Cyberwar Case Study: Georgia 2008.” *Small Wars Journal*, January 6, 2011, 10. <https://smallwarsjournal.com/blog/journal/docs-temp/639-hollis.pdf>.
- Holt, Thaddeus. *The Deceivers: Allied Military Deception in the Second World War*. New York: Scribner, 2014.
- Hsu, Henry, and Peter A Lachenbruch. “Paired t Test.” In *Encyclopedia of Biostatistics*. Wiley Online Library, 2005. <https://onlinelibrary.wiley.com/doi/abs/10.1002/0470011815.b2a15112>
- IBM. “Cost of a Data Breach.” 2022. <https://www.ibm.com/reports/data-breach>.
- International Data Corporation. “IoT Growth Demands Rethink of Long-Term Storage Strategies, Says IDC.” IDC. Accessed August 10, 2021. <https://www.idc.com/getdoc.jsp?containerId=prAP46737220>.
- Jajko, Walter. “Deception: Appeal for Acceptance; Discourse on Doctrine; Preface to Planning.” *Comparative Strategy* 21, no. 5 (October 1, 2002): 351–63. <https://doi.org/10.1080/01495930290043092>.
- Jervis, Robert. *Perception and Misperception in International Politics: New Edition*. Princeton, NJ: Princeton University Press, 2017. <https://doi.org/10.2307/j.ctvc77bx3>.
- Kelly, Christopher, Nikolaos Pitropakis, Alexios Mylonas, Sean McKeown, and William J. Buchanan. “A Comparative Analysis of Honeypots on Different Cloud Platforms.” *Sensors* 21, no. 7 (January 2021): 2433. <https://doi.org/10.3390/s21072433>.
- Koniaris, Ioannis, Georgios Papadimitriou, and Petros Nicopolitidis. “Analysis and Visualization of SSH Attacks Using Honeypots.” In *Eurocon 2013*, 65–72. Zagreb, Croatia: IEEE, 2013.
- Latimer, Jon. *Deception in War: The Art of the Bluff, the Value of Deceit, and the Most Thrilling Episodes of Cunning in Military History, from the Trojan Horse to the Gulf War*. First Edition. New York, NY: Abrams, 2003.
- Leiner, Barry M., Vinton G. Cerf, David D. Clark, Robert E. Kahn, Leonard Kleinrock, Daniel C. Lynch, Jon Postel, Larry G. Roberts, and Stephen Wolff. “A Brief History of the Internet.” *ACM SIGCOMM Computer Communication Review* 39, no. 5 (October 7, 2009): 22–31. <https://doi.org/10.1145/1629607.1629613>.

- Liebowitz, David, Surya Nepal, Kristen Moore, Cody J. Christopher, Salil S. Kanhere, David Nguyen, Roelien C. Timmer, Michael Longland, and Keerth Rathakumar. "Deception for Cyber Defence: Challenges and Opportunities." In *2021 Third IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications (TPS-ISA)*, 173–82, 2021. <https://doi.org/10.1109/TPSISA52974.2021.00020>.
- Lin, Herbert. "Attribution of Malicious Cyber Incidents: From Soup to Nuts." SSRN Scholarly Paper. Rochester, NY: Social Science Research Network, September 2, 2016.
- Lindsay, Jon R. "Cyber Conflict vs. Cyber Command: Hidden Dangers in the American Military Solution to a Large-Scale Intelligence Problem." *Intelligence and National Security* 36, no. 2 (February 23, 2021): 260–78. <https://doi.org/10.1080/02684527.2020.1840746>.
- Liston, Tom. "LaBrea – The Tarpit." Accessed April 19, 2023. <https://labrea.sourceforge.io/Intro-History.html>.
- Lopez, C. Todd. "DOD Releases Path to Cyber Security Through Zero Trust Architecture." U.S. Department of Defense. November 28, 2022. <https://www.defense.gov/News/News-Stories/Article/Article/3229211/dod-releases-path-to-cyber-security-through-zero-trust-architecture/>.
- Mandiant. "Highly Evasive Attacker Leverages SolarWinds Supply Chain to Compromise Multiple Global Victims with SUNBURST Backdoor," December 13, 2020. <https://www.mandiant.com/resources/blog/evasive-attacker-leverages-solarwinds-supply-chain-compromises-with-sunburst-backdoor>.
- Mann, Henry B, and Donald R Whitney. "On a Test of Whether One of Two Random Variables Is Stochastically Larger than the Other." *The Annals of Mathematical Statistics* 18, no. 1 (1947): 50–60. <https://projecteuclid.org/journals/annals-of-mathematical-statistics/volume-18/issue-1/On-a-Test-of-Whether-one-of-Two-Random-Variables/10.1214/aoms/1177730491.full>.
- Maschmeyer, Lennart, and Nadiya Kostyuk. "There Is No Cyber 'Shock and Awe': Plausible Threats in the Ukrainian Conflict." War on the Rocks, February 8, 2022. <https://warontherocks.com/2022/02/there-is-no-cyber-shock-and-awe-plausible-threats-in-the-ukrainian-conflict/>.
- Maurer, Tim. *Cyber Mercenaries: The State, Hackers, and Power*. Cambridge: Cambridge University Press, 2018. <https://doi.org/10.1017/9781316422724>.
- Maxmind. "GeoIP and GeoLite." Accessed February 27, 2023. <https://dev.maxmind.com/geoip?lang=en>.

- McHugh, Mary L. “The Chi-Square Test of Independence.” *Biochemia Medica* 23, no. 2 (June 15, 2013): 143–49. <https://doi.org/10.11613/BM.2013.018>.
- Michelsen, Andreas. *The Submarine Warfare, 1914–1918*. Washington, D.C.: U.S. Government Publishing Office, 1926.
- Microsoft Digital Security Unit. “Special Report: An Overview of Russia’s Cyberattack Activity in Ukraine,” April 27, 2022. Add URL.
- Mokube, Iyatiti, and Michele Adams. “Honeypots: Concepts, Approaches, and Challenges.” In *Proceedings of the 45th Annual Southeast Regional Conference, ACM-SE 45*, 321–26. New York, NY: Association for Computing Machinery, 2007. <https://doi.org/10.1145/1233341.1233399>.
- Morris, Matthew G. “‘Hiding Amongst a Crowd’ and the Illegality of Deceptive Lighting.” *Naval Law Review* 54 (2007). <https://apps.dtic.mil/sti/citations/ADA476997>.
- Muckin, Michael, and Scott C. Fitch. *A Threat-Driven Approach to Cyber Security*. Bethesda, MD: Lockheed Martin Corp., 2014.
- NIST. “Honeypot.” Computer Security Resource Center. Accessed February 10, 2022. <https://csrc.nist.gov/glossary/term/honeypot>.
- . *Risk Management Framework*. NIST Special Publication (SP) 800–53. Gaithersburg, MD: National Institute of Standards and Technology, November 30, 2016.
- Office of Institutional Research, Reporting and Analysis (IRRA). “Naval Postgraduate School Fact Sheet,” April 21, 2022. <https://nps.edu/documents/10180/0/NPS+Fact+Sheet-2022+%281%29.pdf>.
- O’Neill, Patrick Howell. “Chinese Hackers Disguised Themselves as Iran to Target Israel.” *MIT Technology Review*, August 10, 2021. <https://www.technologyreview.com/2021/08/10/1031622/chinese-hackers-false-flag-iran-israel-fireeye/>.
- OwnCloud. “Share Files and Folders, Easy and Secure.” Accessed March 2, 2023. <https://owncloud.com/>.
- Pour, Morteza Safaei, Joseph Khoury, and Elias Bou-Harb. “HoneyComb: A Darknet-Centric Proactive Deception Technique For Curating IoT Malware Forensic Artifacts.” In *NOMS 2022–2022 IEEE/IFIP Network Operations and Management Symposium*, 1–9, 2022. <https://doi.org/10.1109/NOMS54207.2022.9789827>.

- Preston, Diana. *Lusitania: An Epic Tragedy*. New York, NY: Bloomsbury Publishing USA, 2002.
- Raghuprasad, Chetan, and Vanja Svajcer. “Nanocore, Netwire and AsyncRAT Spreading Campaign Uses Public Cloud Infrastructure.” *Cisco Talos Blog* (blog), January 12, 2022. <https://blog.talosintelligence.com/nanocore-netwire-and-asyncrat-spreading/>.
- Razali, Nornadiah Mohd, and Yap Bee Wah. “Power Comparisons of Shapiro-Wilk, Kolmogorov-Smirnov, Lilliefors and Anderson-Darling Tests.” *Journal of Statistical Modeling and Analytics* 2, no. 1 (2011): 21–33.
- Reti, Daniel, David Klaaßen, Simon Duque Anton, and Hans Dieter Schotten. “Secure (S)Hell: Introducing an SSH Deception Proxy Framework.” *ArXiv:2104.03666 [Cs]*, April 8, 2021. <http://arxiv.org/abs/2104.03666>.
- Richter, Philipp, and Arthur Berger. “Scanning the Scanners: Sensing the Internet from a Massively Distributed Network Telescope.” In *Proceedings of the Internet Measurement Conference*, 144–57. New York, NY: Association for Computing Machinery, 2019.
- Rid, Thomas. *Cyber War Will Not Take Place*. New York, NY: Oxford University Press, 2013.
- Roach, J. Ashley. “The Law of Naval Warfare at the Turn of Two Centuries.” *The American Journal of International Law* 94, no. 1 (2000): 64–77. <https://doi.org/10.2307/2555231>.
- Rovner, Joshua. “Cyber War as an Intelligence Contest.” War on the Rocks, September 16, 2019. <https://warontherocks.com/2019/09/cyber-war-as-an-intelligence-contest/>.
- Rowe, Neil C., and Julian Rrushi. *Introduction to Cyberdeception*. Cham: Springer International Publishing, 2016. <https://doi.org/10.1007/978-3-319-41187-3>.
- Shapiro, Samuel Sanford, and Martin B Wilk. “An Analysis of Variance Test for Normality (Complete Samples).” *Biometrika* 52, no. 3/4 (1965): 591–611.
- Shodan. “Shodan Search Engine.” Accessed February 26, 2023. <https://www.shodan.io/dashboard>.
- Smith, Richard W. “The Q-Ship—Cause and Effect.” *U.S. Naval Institute Proceedings* 79, no. 3 (1953). <https://www.usni.org/magazines/proceedings/1953/may/q-ship-cause-and-effect>.
- Sondhaus, Lawrence. *German Submarine Warfare in World War I: The Onset of Total War at Sea*. Lanham, MD: Rowman & Littlefield, 2017.

- Soro, Francesca, Idilio Drago, Martino Trevisan, Marco Mellia, João Ceron, and José J. Santanna. “Are Darknets All The Same? On Darknet Visibility for Security Monitoring.” In *2019 IEEE International Symposium on Local and Metropolitan Area Networks (LANMAN)*, 1–6, 2019. <https://doi.org/10.1109/LANMAN.2019.8847113>.
- Spitzner, L. *Honeypots: Tracking Hackers*. Boston, MA: Addison-Wesley, 2003. <https://books.google.com/books?id=oNFQAAAAMAAJ>.
- Srinivasa, Shreyas, Jens Myrup Pedersen, and Emmanouil Vasilomanolakis. “Deceptive Directories and ‘Vulnerable’ Logs: A Honeypot Study of the LDAP and Log4j Attack Landscape.” In *2022 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, 442–47, 2022. <https://doi.org/10.1109/EuroSPW55150.2022.00052>.
- Steffen, Dirk. “The Holtzendorff Memorandum of 22 December 1916 and Germany’s Declaration of Unrestricted U-Boat Warfare.” *The Journal of Military History* 68, no. 1 (2004): 215–24.
- Sun Tzu (Translated by Samuel B. Griffith). *The Art of War*. Clarendon Press, 1963.
- Taylor, Peter. *The Thirty-Six Stratagems: A Modern Interpretation of a Strategy Classic*. Oxford, UK: Infinite Ideas, 2013.
- Telekom Security. “T-Pot – Technical Concept,” February 27, 2023. <https://github.com/telekom-security/tpotce#technical-concept>.
- . “T-Pot – The All in One Honeypot Platform.” Accessed December 9, 2021. <https://github.com/telekom-security/tpotce>.
- Temple-Raston, Dina. “A ‘Worst Nightmare’ Cyberattack: The Untold Story Of The SolarWinds Hack.” *NPR*, April 16, 2021. <https://www.npr.org/2021/04/16/985439655/a-worst-nightmare-cyberattack-the-untold-story-of-the-solarwinds-hack>.
- Thom, Jay, Yash Shah, and Shamik Sengupta. “Correlation of Cyber Threat Intelligence Data Across Global Honeypots.” In *2021 IEEE 11th Annual Computing and Communication Workshop and Conference (CCWC)*, 0766–72, 2021. <https://doi.org/10.1109/CCWC51732.2021.9376038>.
- Tracy, Miles, Wayne Jansen, Karen Scarfone, and Theodore Winograd. *Guidelines on Securing Public Web Servers*. NIST SP 800-44, Version 2. Gaithersburg, MD: National Institute of Technology, 2007.
- Valeriano, Brandon, Benjamin M. Jensen, and Ryan C. Maness. *Cyber Coercion: The Evolving Character of Cyber Power and Strategy*. New York, NY: Oxford University Press, 2018.

- Washington Post. “The Real Story of How the Internet Became so Vulnerable.” Accessed July 7, 2021. <http://www.washingtonpost.com/sf/business/2015/05/30/net-of-insecurity-part-1/>.
- Welch, Bernard L. “The Generalization of ‘STUDENT’S’ Problem When Several Different Population Variances Are Involved.” *Biometrika* 34, no. 1–2 (1947): 28–35.
- Whaley, Barton. *Stratagem: Deception and Surprise in War*. Norwood, MA: Artech House, 2007. <https://us.artechhouse.com/Stratagem-Deception-and-Surprise-in-War-P1754.aspx>.
- . “The One Percent Solution: Costs and Benefits of Military Deception.” In *Information Strategy and Warfare*. London, UK: Routledge, 2007.
- . “Toward a General Theory of Deception.” *Journal of Strategic Studies* 5, no. 1 (March 1982): 178–92. <https://doi.org/10.1080/01402398208437106>.
- The White House. “Imposing Costs for Harmful Foreign Activities by the Russian Government,” April 15, 2021. <https://www.whitehouse.gov/briefing-room/statements-releases/2021/04/15/fact-sheet-imposing-costs-for-harmful-foreign-activities-by-the-russian-government/>.
- The White House. *National Cybersecurity Strategy*. Washington, DC: White House, 2013. <https://www.whitehouse.gov/wp-content/uploads/2023/03/National-Cybersecurity-Strategy-2023.pdf>.
- Wilcoxon, Frank. “Individual Comparisons by Ranking Methods.” *Biometrics Bull.* 1 (1945): 80–83.
- Wilson, Woodrow. “Joint Address to Congress Leading to a Declaration of War Against Germany (1917).” National Archives, April 2, 1917. <https://www.archives.gov/milestone-documents/address-to-congress-declaration-of-war-against-germany>.
- Wirtz, James J. “The Cyber Pearl Harbor.” *Intelligence and National Security* 32, no. 6 (2017): 758–67.
- Yener, Devran, Tansel Özyer, and Emin Kuğu. “Analysis and Comparison of Honeypot Activities on Two ISP Networks.” In *2021 2nd International Informatics and Software Engineering Conference (IISEC)*, 1–6. Add place of publication: IEEE, 2021.
- Zhang, Li, and Vrizlynn LL Thing. “Three Decades of Deception Techniques in Active Cyber Defense-Retrospect and Outlook.” *Computers & Security* 106 (2021): 102288.

Zimmerman, Donald W. “A Warning about the Large-Sample Wilcoxon-Mann-Whitney Test.” *Understanding Statistics* 2, no. 4 (2003): 267–80.

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



DUDLEY KNOX LIBRARY

NAVAL POSTGRADUATE SCHOOL

WWW.NPS.EDU

WHERE SCIENCE MEETS THE ART OF WARFARE