

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing 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. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 02-03-2023		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 1-Sep-2013 - 30-Sep-2021	
4. TITLE AND SUBTITLE Final Report: Adversarial and Uncertain Reasoning for Adaptive Cyber Defense: Building the Scientific Foundation			5a. CONTRACT NUMBER W911NF-13-1-0421		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 611103		
6. AUTHORS			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES George Mason University 4400 University Drive, MSN 4C6 Fairfax, VA 22030 -4422			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 63820-NC-MUR.80		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			Sushil Jajodia
					19b. TELEPHONE NUMBER 703-993-1653

RPPR Final Report
as of 13-Mar-2023

Agency Code: 21XD

Proposal Number: 63820NCMUR
INVESTIGATOR(S):

Agreement Number: W911NF-13-1-0421

Name: Sushil Jajodia
Email: jajodia@gmu.edu
Phone Number: 7039931653
Principal: Y

Organization: **George Mason University**

Address: 4400 University Drive, MSN 4C6, Fairfax, VA 220304422

Country: USA

DUNS Number: 077817450

EIN: 540836354

Report Date: 31-Dec-2021

Date Received: 02-Mar-2023

Final Report for Period Beginning 01-Sep-2013 and Ending 30-Sep-2021

Title: Adversarial and Uncertain Reasoning for Adaptive Cyber Defense: Building the Scientific Foundation

Begin Performance Period: 01-Sep-2013

End Performance Period: 30-Sep-2021

Report Term: 0-Other

Submitted By: Sushil Jajodia

Email: jajodia@gmu.edu

Phone: (703) 993-1653

Distribution Statement: 1-Approved for public release; distribution is unlimited.

STEM Degrees: 37

STEM Participants: 37

Major Goals: Please see the attached FINAL REPORT.

Accomplishments: Please see the attached FINAL REPORT.

Training Opportunities: Please see the attached FINAL REPORT.

Results Dissemination: Project research results were disseminated through publications in peer reviewed journals and conferences, and presentations at professional meeting and to the government.
Please see the attached FINAL REPORT for details.

Honors and Awards: Please see the attached FINAL REPORT for details.

Protocol Activity Status:

Technology Transfer: Project research results were disseminated through publications in peer reviewed journals and conferences, and presentations at professional meeting and to the government.
Please see the attached FINAL REPORT for details.

ARTICLES:

RPPR Final Report as of 13-Mar-2023

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published

Journal: ACM Transactions on Intelligent Systems and Technology

Publication Identifier Type: DOI

Publication Identifier: 10.1145/2882969

Volume: 8

Issue: 1

First Page #:

Date Submitted: 9/24/18 12:00AM

Date Published: 6/29/16 10:16PM

Publication Location:

Article Title: Dynamic Scheduling of Cybersecurity Analysts for Minimizing Risk Using Reinforcement Learning

Authors: Rajesh Ganesan, Sushil Jajodia, Ankit Shah, Hasan Cam

Keywords: Cybersecurity, Reinforcement Learning, Risk

Abstract: This paper presents a reinforcement learning-based stochastic dynamic programming optimization model that incorporates the estimates of future alert rates and responds by dynamically scheduling cybersecurity analysts to minimize risk and maintain the risk under a pre-determined upper bound. The paper tests the dynamic optimization model and compares the results to an integer programming model that optimizes the static staffing needs based on a daily-average alert generation rate with no estimation of future alert rates (static workforce model).

Distribution Statement: 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support: Y

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published

Journal: ACM Transactions on Intelligent Systems and Technology

Publication Identifier Type: Other

Publication Identifier:

Volume:

Issue:

First Page #:

Date Submitted: 6/29/16 12:00AM

Date Published: 6/29/16 6:28PM

Publication Location:

Article Title: 2.Dynamic Scheduling of Cybersecurity Analysts for Minimizing Risk Using Reinforcement Learning

Authors: Rajesh Ganesan, Sushil Jajodia, Ankit Shah, Hasan Cam

Keywords: Cybersecurity, Reinforcement Learning, Risk

Abstract: The paper presents a reinforcement learning-based stochastic dynamic programming optimization model that incorporates the estimates of future alert rates and responds by dynamically scheduling cybersecurity analysts to minimize risk and maintain the risk under a pre-determined upper bound. The paper tests the dynamic optimization model and compares the results to an integer programming model that optimizes the static staffing needs based on a daily-average alert generation rate with no estimation of future alert rates (static workforce model).

Distribution Statement: 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support: Y

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published

Journal: Journal of Computer Security

Publication Identifier Type: DOI

Publication Identifier: 10.3233/JCS-160555

Volume: 24

Issue: 2016

First Page #: 735

Date Submitted: 9/24/18 12:00AM

Date Published: 6/30/16 5:33PM

Publication Location:

Article Title: Using Temporal Probabilistic Logic for Optimal Monitoring of Security Events

Authors: Sushil Jajodia, Noseong Park, Edoardo Serra, V.S. Subrahmanian

Keywords: Security, Probabilistic Logic

Abstract: In this paper, we study the problem of allocating cybersecurity analysts to alerts generated by intrusion detection systems and other security software. In particular, given an enterprise network (or set of enterprise networks) and information about the value of assets stored at a node (e.g. computer, router) in the network, together with probabilities of compromising a neighbor of a compromised vertex, we show that annotated probabilistic temporal (APT) logic programs allow a defender to express knowledge about the network that captures the probabilities that different nodes will be attacked. In addition, certain APT logic computations, in conjunction with a Stackelberg game theoretic formalization, enable us to capture the attacker's maximal probability of success as well as his ability to maximize damage. We show how the defender can come up with optimal allocations of tasks to cybersecurity analysts.

Distribution Statement: 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support: Y

RPPR Final Report as of 13-Mar-2023

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 3-Accepted

Journal: ACM Transactions on Intelligent Systems and Technology

Publication Identifier Type:

Publication Identifier:

Volume:

Issue:

First Page #:

Date Submitted: 6/30/16 12:00AM

Date Published: 6/30/16 1:40PM

Publication Location:

Article Title: 3.Understanding the Relationship between Human Behavior and Susceptibility to Cyber-Attacks: A Data-Driven Approach

Authors: Michael Ovelgonne, Tudor Dumitras, B. Aditya Prakash, V. S. Subrahmanian, Benjamin Wang, Matthew

Keywords: Cyber-Attacks, Behavior

Abstract: While human users are often considered to be the weakest link in security systems, the risks associated with typical day-to-day computing habits are not well understood. Using Symantec's WINE platform, we conduct a detailed study of 1.6 million machines over an 8-month period in order to learn the relationship between user behavior and cyber attacks against their personal computers. We classify users into 4 categories (gamers, professionals, software developers, others plus a fifth category comprising everyone) and identify a total of 7 independent variables to study: (i) number of binaries (executables) on a machine, (ii) fraction of low-prevalence binaries on a machine, (iii) fraction of high-prevalence binaries on a machine, (iv) fraction of unique binaries on a machine, (v) fraction of downloaded binaries on a machine, (vi) fraction of unsigned binaries on a machine and (vii) travel history of the machine based on number of ISPs from whom the machine connected to the Internet.

Distribution Statement: 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support: Y

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published

Journal: IEEE Security & Privacy

Publication Identifier Type: DOI

Publication Identifier: 10.1109/MSP.2018.1331034

Volume: 16

Issue: 1

First Page #: 92

Date Submitted:

Date Published: 1/1/18 5:00AM

Publication Location:

Article Title: Defending from Stealthy Botnets Using Moving Target Defenses

Authors: Massimiliano Albanese, Sushil Jajodia, Sridhar Venkatesan

Keywords: Moving target defense, botnets

Abstract: Stealthy botnets pose significant threats due their ability to evade traditional defenses and persist in the target system for extended periods of time. Detecting and mitigating these threats is a multifaceted problem that calls for novel solutions to address the several interrelated challenges that stealthy bots introduce. We have shown how the problem can be decomposed into relatively simpler problems that can be tackled separately, while keeping the big picture in mind. Essentially, we need to understand where and when to monitor for potentially suspicious activity, how to look at observed traffic to identify potentially compromised machines, and how to ensure that each and every bot has been removed from the network. Moving target defense has proved to be a viable and promising approach in tackling these challenges, enabling us to achieve some interesting results.

Distribution Statement: 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support: Y

RPPR Final Report as of 13-Mar-2023

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published

Journal: Service Oriented Computing and Applications

Publication Identifier Type: DOI

Publication Identifier: 10.1007/s11761-018-0235-3

Volume: 12

Issue: 2

First Page #: 123

Date Submitted: 9/24/18 12:00AM

Date Published: 4/1/18 12:00AM

Publication Location:

Article Title: Adaptive reallocation of cybersecurity analysts to sensors for balancing risk between sensors

Authors: Ankit Shah, Rajesh Ganesan, Sushil Jajodia, Hasan Cam

Keywords: Cybersecurity analysts · Adaptive reallocation · Balance risk between sensors · Reactive scheduling · Goal Programming

Abstract: Cyber Security Operations Center (CSOC) is a service-oriented system. Analysts work in shifts, and the goal at the end of each shift is to ensure that all alerts from each sensor (client) are analyzed. The goal is often not met because the CSOC is faced with adverse conditions such as variations in alert generation rates or in the time taken to thoroughly analyze new alerts. Current practice at many CSOCs is to pre-assign analysts to sensors based on their expertise, and the alerts from the sensors are triaged, queued, and presented to analysts. Under adverse conditions, some sensors have more number of unanalyzed alerts (backlogs) than others, which results in a major security gap for the clients if left unattended. The paper presents a technical solution to achieve the objectives and answers two important research questions: (i) detection of triggers, which determines when-to reallocate, and (ii) how to optimally reallocate analysts to sensors.

Distribution Statement: 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support: Y

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published

Journal: ACM Transactions on Privacy and Security

Publication Identifier Type: DOI

Publication Identifier: 10.1145/3196884

Volume: 21

Issue: 4

First Page #: 1

Date Submitted: 9/24/18 12:00AM

Date Published: 6/1/18 4:00AM

Publication Location:

Article Title: VULCON

Authors: Katheryn A. Farris, Ankit Shah, George Cybenko, Rajesh Ganesan, Sushil Jajodia

Keywords: Cyber-security analysts, vulnerability triage and management

Abstract: Vulnerability remediation is a critical task in operational software and network security management. In this article, an effective vulnerability management strategy, called VULCON (VULnerability CONTROL), is developed and evaluated. The strategy is based on two fundamental performance metrics: (1) time-to-vulnerability remediation (TVR) and (2) total vulnerability exposure (TVE). VULCON takes as input real vulnerability scan reports, metadata about the discovered vulnerabilities, asset criticality, and personnel resources. VULCON uses a mixed-integer multiobjective optimization algorithm to prioritize vulnerabilities for patching, such that the above performance metrics are optimized subject to the given resource constraints. VULCON has been tested on multiple months of real scan data from a cyber-security operations center (CSOC). Results indicate an overall TVE reduction of 8.97% when VULCON optimizes a realistic security analyst workforce's effort.

Distribution Statement: 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support: Y

CONFERENCE PAPERS:

Publication Type: Conference Paper or Presentation

Publication Status: 1-Published

Conference Name: Proceedings of USENIX Security 2015.

Date Received:

Conference Date:

Date Published:

Conference Location:

Paper Title: Towards Discovering and Understanding Task Hijacking in Android

Authors:

Acknowledged Federal Support:

RPPR Final Report
as of 13-Mar-2023

DISSERTATIONS:

Publication Type: Thesis or Dissertation

Institution: George Mason University

Date Received: 29-Sep-2017

Completion Date: 8/31/17 4:00AM

Title: A QUANTITATIVE FRAMEWORK FOR CYBER MOVING TARGET DEFENSES

Authors: Warren Cnnell

Acknowledged Federal Support: **Y**

Publication Type: Thesis or Dissertation

Institution: Dartmouth College

Date Received: 29-Sep-2017

Completion Date: 8/31/17 4:40PM

Title: The Quantification and Analysis of Cyber-Security Operations Center Vulnerability Data

Authors: KATHERYN FARRIS

Acknowledged Federal Support: **Y**

Publication Type: Thesis or Dissertation

Institution: George Mason University

Date Received: 29-Sep-2017

Completion Date: 8/31/17 4:00AM

Title: ADAPTIVE CYBER DEFENSES TO MITIGATE BOTNET-BORNE THREATS

Authors: Sridhar Venkatesan

Acknowledged Federal Support: **Y**

Partners

,

I certify that the information in the report is complete and accurate:

Signature: Sushil Jajodia

Signature Date: 3/2/23 12:53PM

Final Report

ADVERSARIAL AND UNCERTAIN REASONING FOR ADAPTIVE CYBER DEFENSE: BUILDING THE SCIENTIFIC FOUNDATION

ARO Grant No. W911NF-13-1-0421

Submitted by

Sushil Jajodia
Center for Secure Information Systems
George Mason University
Fairfax, VA 22030-4422
jajodia@gmu.edu

Principal Investigators

Sushil Jajodia, Massimiliano Albanese
Center for Secure Information Systems, George Mason University, Fairfax, VA 22030-4422

George Cybenko
Thayer School of Engineering, Dartmouth College, Hanover, NH 03755

Peng Liu, Minghui Zhu,
College of Information Sciences and Technology, Pennsylvania State University, University
Park, PA 16802

Michael Wellman, Satinder Singh Baveja, Demosthenis Teneketzis,
Department of Electrical Engineering and Computer Science, University of Michigan,
Ann Arbor, MI 48109-2122

TABLE OF CONTENTS

Executive Summary	5
1. Introduction.....	6
2. Scientific Progress in Year 1	7
Manipulating a system's attack surface.....	7
Modeling defensive situational awareness and adversaries.....	9
Control-Theoretic Reasoning and Game-Theoretic Reasoning	10
A control theory framework for comparing cost-effectiveness of different MTD techniques.....	11
Reinforcement learning algorithms for adaptive cyber defense against Heartbleed attack	12
A Novel Probabilistic Approach to Detect and Shield Heartbleed Attack	13
3. Scientific Progress in Year 2	14
Quantification of MTDs and Situational Awareness.....	15
Optimal Scheduling of Cybersecurity Analysts for Minimizing Risk.....	18
Dynamic Scheduling of Cybersecurity Analysts for Minimizing Risk Using Reinforcement Learning	18
Control-Theoretic Reasoning Techniques	18
Game-Theoretic Reasoning Techniques.....	19
Manipulating the Attack Surface of a System	19
A Practical Approach for Adaptive Data Structure Layout Randomization	22
A Novel Defense Mechanism Providing Effectiveness Guarantees Against Heartbleed.....	23
An Innovative Symbolic N-Variant System for Practical N-Variant System MTD Against Zero-Day Attacks.....	24
4. Scientific Progress in Year 3	25
MTD Quantification.....	25
A Moving Target Defense Approach to Mitigate DDoS Attacks against Proxy-Based Architectures..	27
Moving Target Defense against DDoS Attacks: An Empirical Game-Theoretic Analysis.....	27
Efficient Computation of Centrality Measures	27
Cyber Resiliency	28
Vulnerability Remediation.....	28
5. Scientific Progress in Year 4	28
Game-Theoretic Learning for Adaptive Cyber Defense.....	28

Adaptive ASLR (Address Space Layout Randomization).....	29
Cooperative Attack Surface “Divide and Conquer” for Memory Protection	30
Applying Deep Learning to Predict How the Attacker Can Bypass CFI (Control-Flow Integrity)	31
A POMDP Approach to the Dynamic Defense of Large-Scale Cyber Networks.....	32
Dynamic Games with Asymmetric Information	33
Stochastic dynamic games with delayed sharing information structure.....	33
A Stackelberg Game Model for Botnet Traffic Exfiltration.....	33
Deep Learning	34
6. Scientific Progress in Year 5	34
Quantification and Evaluation of Adaptive Cyber Defenses	35
Quantification and Mitigation of Advanced Persistent Threats	36
Defeating Data Structure Manipulation Attacks	36
Bayesian Attack Graphs.....	37
Applying Deep Learning to Accurately Detect ROP Payloads	38
Adaptive Management of System Vulnerabilities	38
Formulating and Analyzing Dynamic Cyber-Security Game Models	39
Fundamental Control- and Game-Theoretic Methods.....	40
7. Scientific Progress in Year 6	40
Game-Theoretic Reinforcement Learning.....	41
Repackaging attacks against mobile apps.....	41
Defending against zero-day buffer over-read attacks.....	42
Cyber Security Operation Centers (CSOCs)	42
Edited Book.....	43
8. Scientific Progress in Year 7	43
Modeling Stealthy Adversaries	44
Maintaining the Level of Operational Effectiveness of a CSOC under All Adverse Conditions	45
Defending Against Multi-Stage Attacks Using Reinforcement Learning-based POMDP.....	45
Deep Learning-based Technique to Detect Previously Unseen ROP Payloads Without Using Any Heuristics.....	46

Reinforcement Learning-based ACD Technique to Ensure Quantitatively Bounded Damages in Defending Heartbleed-Like Attacks	46
9. Scientific Progress in Year 8	47
Attributable Machine Learning	47
Game-Theoretical Analysis	48
Novel Anomaly Detection Approach for Discrete Event Logs.....	48
Deep-learning-based technique to identify silent buffer overflow vulnerabilities in program binaries without source code.....	49
Tackling Imbalanced Data in Cybersecurity with Transfer Learning: A Case with ROP Payload Detection.....	49
10. Technology Transfer.....	50
11. Broader Impact	51
12. Student/Supported Personnel.....	51
Graduate Students	51
Post Doctorates.....	52
Faculty.....	52
Undergraduate Students	53
Undergraduate Degrees Awarded.....	53
Masters Degrees Awarded.....	53
Ph.D.s Awarded.....	53
13. Honors and Awards	54
14. PUBLICATIONS	56
Papers published in peer-reviewed journals.....	56
Papers published in non-peer-reviewed journals.....	75
Presentations at meetings, but not published in Conference Proceedings	75
Peer-Reviewed Conference Proceeding publications	79
Book Chapters.....	97
Books.....	102
Edited Book.....	102
Book chapters	104

EXECUTIVE SUMMARY

This Multidisciplinary University Research Initiative (MURI) project has developed a new class of technologies called Adaptive Cyber Defense (ACD) by building on two active but separate research areas: Adaptation Techniques (AT) and Adversarial Reasoning (AR). Our research has yielded a rich repertoire of AT methods for introducing diversity and uncertainty into networks, applications, and hosts. Moreover, we have investigated the criteria for deciding where, when, and how to best employ available AT options. Such management decisions are complex due to the performance and security tradeoffs inherent in AT approaches. To address such challenges, this project has harnessed a broad array of AR techniques to identify effective and stable strategies for deploying AT options in operational systems. AR combines machine learning, behavioral science, operations research, control theory, and game theory to address the goal of computing effective strategies in dynamic, adversarial environments.

Building on previous results from this multi-year effort and on state-of-the-art machine learning and adversarial modeling techniques, we have advanced the state of the art in autonomous adaptive cyber defense by developing the scientific foundations and associated understanding of autonomous cyber operations and specifically how adaptive cyber defenses play a role in such operations, thus paving the way for the development of next-generation autonomous cyber operations. Specifically, we developed several novel ACD mechanisms based on reinforcement learning and deep learning, advanced our understating of how stealthy attackers establish and expand their foothold within target systems, and developed the capability to ensure a predetermined level of operational effectiveness in Cyber Security Operations Centers (CSOCs) under all adverse conditions.

1. INTRODUCTION

Today's cyber defenses are largely static. They are governed by slow deliberative processes involving testing, security patch deployment, and human-in-the-loop monitoring. As a result, adversaries can systematically probe target networks, pre-plan their attacks, and ultimately persist for long times inside compromised networks and hosts.

In response to this situation, researchers in recent years have started to investigate various methods that make networked information systems less homogeneous and less predictable. The basic idea of Adaptation Techniques (AT) is to engineer systems that have homogeneous functionalities but randomized manifestations. Homogeneous functionality allows authorized use of networks and services in predictable, standardized ways while randomized manifestations make it difficult for attackers to engineer exploits remotely.

Examples of Adaptation Techniques include concepts such as Moving Target Defenses (MTD) as well as artificial diversity and bio-inspired defenses to the extent that they involve system adaption for security and resiliency purposes. Unfortunately, the majority of AT research has been focused on developing specific new techniques as opposed to understanding their overall operational costs, when they are most useful, and what their possible inter-relationships might be. Moreover, the AT approaches assume stationary and stochastic, but non-adversarial, environments. Situations with intelligent peer adversaries operating in and changing a networked environment produce dynamic behaviors that violate these assumptions, potentially defeating these adaptations.

This Multidisciplinary University Research Initiative (MURI) project has developed a new class of technologies called Adaptive Cyber Defense (ACD) by building on two active but separate research areas: Adaptation Techniques (AT) and Adversarial Reasoning (AR).

Our research has yielded a rich repertoire of AT methods for introducing diversity and uncertainty into networks, applications, and hosts. Moreover, we have investigated the criteria for deciding where, when, and how to best employ available AT options. Such management decisions are complex due to the performance and security tradeoffs inherent in AT approaches. To address such challenges, this project has harnessed a broad array of Adversarial Reasoning (AR) techniques to identify effective and stable strategies for deploying AT options in operational systems. AR combines machine learning, behavioral science, operations research, control theory, and game theory to address the goal of computing effective strategies in dynamic, adversarial environments.

These techniques force adversaries to continually re-assess, re-engineer and re-launch their cyber-attacks. By integrating game-theoretic and control-theoretic analyses for tradeoff analysis, ACD presents adversaries with optimized and dynamically changing attack surfaces and system configurations, thereby significantly increasing the attacker's workloads and decreasing their probabilities of success.

Building on previous results from this multi-year effort and on state-of-the-art machine learning and adversarial modeling techniques, we have advanced the state of the art in autonomous adaptive cyber defense by developing the scientific foundations and associated understanding of autonomous cyber operations and specifically how adaptive cyber defenses play a role in such operations, thus paving the way for the development of next-generation autonomous cyber operations. Specifically, we developed several novel ACD mechanisms based on reinforcement learning and deep learning, advanced our understating of how stealthy attackers establish and expand their foothold within target systems, and developed the capability to ensure a predetermined level of operational effectiveness in Cyber Security Operations Centers (CSOCs) under all adverse conditions.

This coherent and focused research effort have yielded: (a) scientific and engineering principles that enable effective Adaptive Cyber Defense, and (b) prototypes and demonstrations of technologies embodying these principles in several real-world scenarios. Details about specific achievements are provided in the subsections below.

2. SCIENTIFIC PROGRESS IN YEAR 1

The primary objective of the effort in Year 1 has been twofold. On one side, we focused on formalizing and quantifying several aspects of cyber defenses. To this aim, we have proposed an initial formalization of a system's attack surface, as well as metrics and approaches to assess the cost-effectiveness of different adaptive techniques in reshaping a system's attack surface. On the other side, we began exploring game-theoretic and control-theoretic aspects of adaptive cyber defenses. This effort resulted in a number of preliminary, yet significant, scientific accomplishments, which were disseminated nationally and internationally through a number of presentations and publications in peer-reviewed venues. First, we defined a framework for modeling, evaluating, and countering an attacker's reconnaissance effort. The proposed defensive strategy consists in controlling how the attacker perceives the attack surface of the target system. Second, we started to develop the quantification of moving target defenses needed for control and game models. Analysis of such defenses has led to the conclusion that moving target defenses are the most effective defenses against well designed botnets. Third, we developed a general (supervisory) control-theoretic formulation of a dynamic cyber-security problem, from the perspective of a defender aiming to maximize its worst-case performance. Fourth, we designed and implemented a simulation model of the basic adaptive cyber-defense scenario, and commenced game-theoretic analysis of this model. Fifth, we developed a control theory-based framework for comparing cost-effectiveness of different MTD techniques. Finally, we developed reinforcement learning algorithms for adaptive cyber defense against Heartbleed attack and a novel probabilistic approach to detect and shield Heartbleed attacks.

The following sections describe each of our major accomplishments in more details.

MANIPULATING A SYSTEM'S ATTACK SURFACE

Today's approach to cyber defense is governed by slow and deliberative processes such as security patch deployment, testing, episodic penetration exercises, and human-in-the-loop monitoring of security events. Adversaries can greatly benefit from this situation, and can continuously and systematically probe target networks with the confidence that those networks will change slowly if at all. In fact, cyber-attacks are typically preceded by a reconnaissance phase in which adversaries aim at collecting valuable information about the target system, including network topology, service dependencies, and unpatched vulnerabilities. As most system configurations are static – hosts, networks, software, and services do not reconfigure, adapt, or regenerate except in deterministic ways to support maintenance and uptime requirements – it is only a matter of time for attackers to acquire accurate knowledge about the target system. This will eventually enable them to engineer reliable exploits and pre-plan their attacks. In order to address this important problem, significant work has been done in the area Adaptive Cyber Defense (ACD), which includes concepts such as Moving Target Defense (MTD), as well as artificial diversity and bio-inspired defenses, to the extent they involve system adaptation for security and resiliency purposes.

Essentially, a number of techniques have been proposed to dynamically change a system's attack surface by periodically reconfiguring some aspects of the system.

A system's attack surface has been defined as the "subset of the system's resources (methods, channels, and data) that can be potentially used by an attacker to launch an attack". Intuitively, dynamically reconfiguring a system is expected to introduce uncertainty for the attacker and increase the cost of the reconnaissance effort. However, one of the major drawbacks of current approaches is that they force the defender to periodically reconfigure the system, and this may introduce a costly overhead for legitimate users, as well as the potential for denial of service conditions. Additionally, most of the existing techniques are purely proactive in nature or do not adequately consider the attacker's behavior.

During Year 1, we aimed at advancing the state of the art in adaptive cyber defense by looking at the problem from a control perspective and proposing a graph-based approach to manipulate the attacker's perception of a system's attack surface. To achieve this objective, we formalized the notion of *system view* as well as the notion of distance between views. We refer to the attacker's view of the system as the external view and to the defender's view as the internal view. A system's attack surface can then be thought of as the subset of the internal view that would be exposed to potential attackers when no deceptive strategy is adopted. Starting from these definitions, we developed a principled yet practical approach to manipulate responses to attacker's probes so as to induce an external view of the system that satisfies certain desirable properties. In particular, we proposed efficient algorithmic solutions to different classes of problems, namely

- (i) inducing an external view that is at a minimum distance from the internal view while minimizing the cost for the defender;
- (ii) inducing an external view that maximizes the distance from the internal one, given an upper bound on the admissible cost for the defender.




Figure 1 - Control Perspective

Figure 1 shows how standard and ACD-based controllers can be adopted in order to alter the attacker's perception of the attack surface. State of the art solutions aim at dynamically reconfiguring the system in order to create a moving target for the attacker. The controller uses internal and external system variables to monitor the system's state and realize a moving target objective. On the other hand, our approach aims

at transparently altering the attacker's perception. The attacker's knowledge of the system is inferred from the system's observable variables (e.g., ports, packets, web-pages content, DNS records). In order to gather additional information, the attacker probes the system. The proposed ACD controller identifies such probes and controls the information exposed through the system's observable variables by crafting adequate probe responses to bias the attacker's view of the system.

Compared to traditional approaches, our novel approach presents two key advantages: (i) we can introduce uncertainty for the attackers without actually changing the system's configuration, thus minimizing the overhead; (ii) we can deceive the attackers and steer them away from critical resources rather than simply introducing uncertainty and forcing them to use a random strategy. Experiments conducted on a prototypal implementation of the proposed algorithms confirm that our approach is efficient and effective in steering the attackers away from critical resources.

This overall framework is presented in a paper accepted for publication at CNS 2014. In another paper, recently published at SECURITY 2014, we presented an approach to artificially and arbitrarily delay intrusions, which fits within the scope of this framework. In this work, we formalized the attacker's capability of penetrating a system, and altered his perception of the attack surface by virtually modifying the configuration through the addition of *distraction clusters*.

MODELING DEFENSIVE SITUATIONAL AWARENESS AND ADVERSARIES

Cyber situational awareness is an important ingredient when seeking to apply adaptive cyber defenses or moving target technologies dynamically in an operational environment. Additionally, models of cyber adversaries' capabilities and intentions are needed to properly construct control and game theoretic models that can lead to effective strategies for moving target deployments. Finally, some quantification of moving target defenses themselves is needed for control and game models. This project is undertaking the quantification of these ingredients.

We have developed a model for defensive postures that captures the reality of modern enterprise intrusion technologies as actually deployed in commercial, government and military networks. Sensor reports are combined into a disjunction normal form - a disjunction of conjunctions - that represents alert rules. The resulting alert logic is an abstraction of what situations will result in alerts and what combination of sensors will fail to alert. The latter then can be analyzed further to define the malware's coyness score with respect to specific defensive posture.

Analysis of such defenses has led to the conclusion that moving target defenses are the most effective defenses against well designed botnets. Moreover, by randomizing some of the alert logic rules, it is possible to create a moving target defense based on detection logics. This is especially relevant given that current cyber analysts' workloads prohibit investigating all alerts equally so randomization might allow a defender to create an environment in which the attacker cannot reliably estimate the defenses in place.

Additionally, we have identified and formalized several realistic cyber attack-defend games that arise in moving target defenses. Those models are currently being studied to reveal promising strategies and scaled to larger and more complex system configurations.

In addition, we have developed an information market type approach to quantifying moving target defenses in terms of their effectiveness and cost-to-deploy. Initial results of that market will be available in 4th quarter 2014.

We will complete the information market design and actually run the market several times to ascertain the accuracy and efficacy of that approach. Given that there are some early-stage simulation approaches to quantifying moving target defenses, we hope to be able to compare the two approaches.

Using the progress made on modeling defensive situational awareness and adversaries, we hope to have a generic control framework and near-optimal control computation ready by the end of the second project year.

CONTROL-THEORETIC REASONING AND GAME-THEORETIC REASONING

During this first year of the project, we began exploring both game-theoretic and control-theoretic reasoning. In both cases we started with a relatively simple adaptive cyber-defense scenario, collaboratively developed by team members from all participating institutions. The scenario comprises one attacker, and one defender. The attacker may probe or attack any of the defender's hosts, probabilistically succeeding in compromise. The defender pursues an adaptive strategy (e.g., a moving target defense), which results in resetting the machine state to erase the attacker's progress, and regain control (if it was compromised). Our premise is that thoroughly understanding the control and strategic implications of this simple scenario will provide insights for addressing more complex situations, and drive the development of techniques for reasoning about extended scenarios.

During this first year, we developed a general (supervisory) control-theoretic formulation of a dynamic cyber-security problem, from the perspective of the defender. In this formulation, the defender aims to maximize its worst-case performance, considering that the attacker may behave arbitrarily. The model we formalized captures several key features of cyber-security in this project: (1) progressive attacks, where attacker's actions have cumulative effect, and (2) incomplete observability of the state of the system (e.g., the degree of attacker success in compromising hosts on the network).

The first paper on this approach (appeared in GameSec, November 2014) accomplishes several goals, including the following.

1. A formalization of the supervisory control approach to dynamic cyber-defense.
2. A dynamic-programming algorithm for deriving optimal defender policies.
3. Computational analysis of simulations (using a discrete-event system modeling package) illustrating the structure of defender policy under a range of parametric environments.

The model studied so far is extensible in several directions, which we plan to pursue in follow-on work.

With respect to game-theoretic reasoning, we designed and implemented a simulation model of the basic adaptive cyber-defense scenario, and commenced game-theoretic analysis of this model. We take an empirical game-theoretic (EGT) approach, which uses results of extensive simulation to construct a formal game model. The model we started with can be viewed as an extension of the FlipIt game, with multiple hosts, asymmetric stealth, and progressive attack.

The first paper on this approach (appeared in GameSec, November 2014) reports preliminary results of the EGT study. Contributions include:

1. Simulation model exposing relevant environment parameters; initial set of heuristic strategies for attacker and defender.
2. Analysis of game over 87x58 strategies, in six environment settings.
3. Identified three qualitatively distinct equilibria, including "Maximal Defense", which is proved to be in equilibrium regardless of parameter settings and strategy sets.

One of the key insights from this study is that deterrence hinges pivotally on the defender's ability to perfectly observe attacker probes. We plan to relax this assumption, as well as explore other extensions in the coming year.

The year 1 effort establishes a solid foundation for both tracks, which we plan to build on in the coming year. In addition to extending the current models, we expect to explore more detailed domain-specific problem instances. We are also planning to devote significantly more effort on reinforcement learning techniques (which in a sense lie between game- and control-theoretic approaches).

A CONTROL THEORY FRAMEWORK FOR COMPARING COST-EFFECTIVENESS OF DIFFERENT MTD TECHNIQUES

Quantitatively comparing the cost-effectiveness of different MTD techniques is a critical part of the scientific foundation of Adaptive Cyber Defense. When there is no ongoing mission, the existing MTD comparison techniques can be broken down into four categories: (1) Attack-based experiments (on a test-bed); (2) Probability models; (3) Simulation-based (cost-effectiveness) evaluation; (4) Expert opinion based MTD comparison. When there are ongoing missions in progress (i.e., real-time MTD), these four categories of MTD comparison methods, however, are very limited in quantitatively comparing the cost-effectiveness of different MTD techniques. The main limitations include (a) inability to handle the mission dynamics and semantics; (b) assuming prior distributions (about the attack); and (c) viewing a host as a black-box.

When a mission is ongoing, the literature suggests three possible ways to compare different MTD options/strategies. The first method is to use security monitoring techniques to "observe" the impact of the cyber-attack (if any) on the mission. However, this method is only cost-effective to certain extent: the performance overhead could be very significant. The second method is to use standby systems to simultaneously deploy more than one MTD option. However, this method could be too costly, and inconsistency could emerge between the primary system and the standby system. The third method is high-fidelity simulation. However, this method may exhibit poor scalability.

During Year 1, we developed several key components of a control theory inspired framework for comparing the cost-effectiveness of different MTD techniques. In particular, our main research contributions are as follows:

- The conceptual framework is developed. It is a totally new MTD framework. It is inspired by control theory, especially RHC (Receding Horizon Control).

- We developed a novel three-layer model to evaluate and compare the cost-effectiveness of different MTD options. The first layer captures low-level mission asset information (e.g., heap buffers, files, network messages) from individual programs (i.e., service processes). For each program, a state machine is created to present how attacks cause damage (i.e., effects on mission assets) to the program when it runs to conduct a mission task. The second layer inter-connects the state machines from the first layer, and models damage propagation between a set of programs that work together to conduct mission tasks. In this layer, a state machine is constructed to demonstrate how attacks damage a mission system that runs a number of programs. The third layer explicitly shows the cost-effectiveness of different MTD options in terms of the aggregated effects on mission assets. The state machine in the third layer shows how different MTD options result in different state machines in the second layer. Each state machine in the second layer is an aggregation of a subset of state machines in the first layer.
- We have been developing an innovative simulation technology to implement the three-layer model in a software system prototype. To address the black-box limitation of the existing simulation-based evaluation method, we propose a new concept called “execution path based simulation”. This new concept leads to white-box (high-fidelity) simulation of the effects of cyber-attacks on the mission assets owned or accessed by a program that runs to conduct a mission task. To address the fidelity challenge, we propose to leverage the high execution-path coverage (over 90%) achieved by the state-of-the-art symbolic execution technology. To address the scalability challenge, we have been exploring the use of cutting-edge “big data” technology, including highly scalable parallel data analytics running on a cluster.
- We have partially built the system prototype. A cluster that consists of 4 workstations is built. A million of execution paths have been extracted through symbolic execution of programs in a particular mission scenario. Basic execution path based simulation capabilities have been developed.
- We have done preliminary evaluation of the system prototype. For example, we have compared the execution speed of the simulator with that of the real mission system.

REINFORCEMENT LEARNING ALGORITHMS FOR ADAPTIVE CYBER DEFENSE AGAINST HEARTBLEED ATTACK

In adaptive cyber defense, a critical task of the defender is to identify optimal defenses which can minimize damages caused by attacks. Note that the damages are jointly determined by both deployed defenses and launched attacks. In order to precisely predict the effectiveness of defenses, the defender needs to access the information of the attacker, e.g., his/her targets, resources, intentions, preferences and launched attacks. However, the information is difficult for the defender to gather in adversarial environments. As a result, the defender may spend substantial time and effort to protect an asset which may not be the target of the attacker.

To address the challenge, we designed and formally analyzed two reinforcement learning algorithms for adaptive cyber defense against Heartbleed attack. Our contributions are summarized as follows:

- We came up with the adaptive reinforcement learning algorithm. At each time instant, the defender, on one hand, reinforces the successful defense in the last two time instants with a certain exploitation

rate, and on the other hand, randomly chooses any feasible defense. By using stochastic stability of Markov chains, we formally proved that, with probability one, the defense converges to the best response to the attacks when the exploration rate of the attacker is diminishing.

- We proposed the robust reinforcement learning algorithm for the case when the exploration rate of the attacker is persistent. At each time instant, the defender, on one hand, reinforces the defense which minimizes the maximal utility in the whole observed history with a certain exploitation rate, and on the other hand, randomly chooses any feasible defense. We formally proved that, with a probability arbitrarily close to one, the defense converges to the min-max strategy despite that the attacker could persistently explore the system.
- We used numerical simulations to verify the performance of the proposed algorithms.
- The results are summarized in a paper which was published in the first ACM Workshop on Moving Target Defense.

A NOVEL PROBABILISTIC APPROACH TO DETECT AND SHIELD HEARTBLEED ATTACK

Heap buffer overflow bugs, where reads or writes go beyond heap buffer boundaries, can lead to crashes, erroneous execution, and security vulnerabilities. In this work, we propose HeapTherapy, a comprehensive defense against both heap buffer over-read and over-write. HeapTherapy applies several light-weight detection and shielding methods, including a novel technique for probabilistic over-read detection (detecting Heartbleed attack). A key insight behind our approach is the leveraging of calling context information to locate vulnerable buffers and generate temporary patches at user side. The temporary patches are used to prevent recurring buffer overflows exploiting from the same vulnerability and enable the normal execution of the program. Our experiment results show that HeapTherapy can detect and defend against the Heartbleed attack with little throughput slowdown. And the performance overhead on SPEC CPU2006 is 6%.

To the best of our knowledge, this is the first probabilistic detection technique ever developed to deal with Heartbleed attack. It works very effectively even if the defender has zero knowledge about the vulnerable heap buffer which the adversary will exploit. In our evaluation experiments, we run Nginx, a widely-deployed web server program, and we simulate benign HTTPS requests using ApacheBench. Then, we use the same Heartbleed attack scripts used by real-world hackers to send over-read requests (we set the length to be 64KB) to Nginx continuously. The results show that even if we use a (page guarding) probability as low as 0.01, the number of attacker's successful over-reads before the first guard page is triggered is only on average less than 8, which means that the attacker can only read Nginx's memory 8 times (64KB per over-read) and then all the subsequent Heartbleed requests will be completely shielded (i.e., blocked) by HeapTherapy. Given that real-world Heartbleed attacks usually require more than 100,000 over-reads before achieving the attacker's goal, HeapTherapy should be able to defeat almost all of the real-world Heartbleed attacks. Compared to 100,000 over-reads, 8 over-reads provides almost no help to the adversary.

A paper which summarizes this work was been submitted to NDSS'15.

3. SCIENTIFIC PROGRESS IN YEAR 2

The primary objective of the effort in Year 2 has been on quantification of MTD techniques and on methods for deploying available techniques in optimal ways. We also investigated new and more effective adaptive techniques to increase complexity for the attacker. This effort resulted in a number of preliminary, yet significant, scientific accomplishments, which were disseminated nationally and internationally through a number of presentations and publications in peer-reviewed venues.

First, we identified and compared techniques for quantifying Moving Target Defenses (MTD) so that these quantifications can be used by the various control and game theoretic optimization algorithms being developed in this project. That work includes the organization and hosting of an ARO-funded workshop on the topic of MTD Quantification, held at George Mason University on August 31 and September 1, 2015. As part of the quantification aspect of the ACD effort, we have been developing approaches for modeling and analyzing attacker and defender behaviors within a network with the goals of populating realistic dynamic models of network vulnerabilities and normal behaviors on a network for cyber situational awareness to drive real-time control and game theoretic decisions.

Second, we looked at how the demand for cybersecurity analysts outpaces supply, which calls for optimal management of the analyst resource. Therefore, a key component of the cybersecurity defense system is the optimal scheduling of its analysts. We developed two models – a static model and a reinforcement learning-based stochastic dynamic programming optimization model – for dynamically scheduling cybersecurity analysts to minimize risk and maintain risk under a pre-determined upper bound. This work is important because it focuses on the requirements of human cybersecurity analysts and ways to improve their efficiency. This work is not only of theoretical but practical interest as well.

Third, we continued our development of game-theoretic and control-theoretic reasoning methods for adaptive cyber-defense. The goal of this effort is not to develop new defense primitives, but rather to produce techniques for deploying available primitives in an optimal way against a fixed or worst-case adversary (control theory approach), or against a rational adversary (game theory approach). In both cases we built on and significantly extended models and methods developed during the first year. The models employed are somewhat different, but both start from a scenario comprising one attacker and one defender. The attacker may probe or attack any of the defender's hosts, probabilistically succeeding in compromising them. The defender pursues an adaptive strategy (e.g., a moving target defense), which results in resetting the machine state to erase the attacker's progress, and regain control (if it was compromised).

Fourth, we continued to investigate techniques to increase the complexity for the attacker by manipulating a system's attack surface.

Fifth, we developed the first adaptive DSLR (Data Structure Layout Randomization) technique to protect enterprise software systems from memory content leakage attacks and DKOM (Direct Kernel Object Manipulation) attacks through a rootkit

Sixth, we developed a novel defense mechanism that provides mathematically provable effectiveness guarantees against the Heartbleed attack through mathematics (i.e., stochastic online optimization with bandit feedback) and software-system co-design.

Last we developed an innovative symbolic N-variant system to do more practical N-variant defense against zero-day attacks.

The following sections describe each of our major accomplishments in more details.

QUANTIFICATION OF MTDs AND SITUATIONAL AWARENESS

During Year 1, the early stage of MTD quantification research was identified as a serious challenge to the field in general and this ACD project in particular. During Year 2, progress was made on two fronts to address this issue.

A survey was developed and presented to experts in the MTD field with the goal of eliciting quantifications of MTD techniques in a variety of dimensions. Using the classification scheme developed by the MIT Lincoln Labs MTD research group, the survey was conducted over the summer of 2014 and results were analyzed during Year 2 of this project. One depiction of those results is presented in Figure 1.




Figure 2: Effectiveness and Costs of Various MTD Techniques

Figure 1 depicts a variety of MTD techniques plotted using two dimensions: cost (to deploy and operate) and effectiveness. The various MTD techniques are color-coded based on what class the technique falls into (according to a taxonomy developed previously by MIT Lincoln Labs). That work has identified several techniques as relatively ineffective for cyber defense against advanced adversary types. In particular, techniques that introduce uncertainty into network addressing and discovery appear to be easy to circumvent for skilled attackers and offer limited defenses only in certain early stages of the Cyber Kill Chain.




Figure 3: Applicability of ACD MTD Techniques to Attacks

Figure 2 further depicts how the MTD techniques being developed in the ACD project address various attack classes (as defined in the MIT Lincoln Labs taxonomy). This figure suggests that the project's techniques have good coverage collectively across attack classes (namely that each column, representing an attack class, is being covered by multiple techniques and the individual techniques are relatively broad in scope (namely that each row is addressing multiple attack classes)).

Additional work has been done on modeling the time-varying vulnerability of machines within an enterprise-scale network using empirical data. Monthly vulnerability scans, using the commercial Nessus scanner, were analyzed to discover what kinds of known vulnerabilities are present on which machines on which monthly scans. Analysis of the high risk and critical vulnerabilities over scans and machines lead to the models of vulnerabilities depicted in Figure 3, Figure 4, and Figure 5. While the specific parameters of the model will vary by organization, this methodology can be applied to such vulnerability scan data to better understand how networks and their individual constituent machines are exposed to attacks. This understanding is necessary to optimize the application of MTD techniques to operational networks because it quantifies the risk-benefit tradeoffs needed for optimization and game theoretic algorithms.




Figure 4: A Model of the Birth Process for Vulnerabilities




Figure 5: A Model of the Death Process for Vulnerabilities




Figure 6: The Distribution of Vulnerability Lifetimes

These analyses are based on empirical data that are incomplete, noisy and drawn over about a dozen months (with gaps that are interpolated) so that standard queuing theory results are not necessarily

applicable (such as Little's Theorem). This work does however provide an empirical approach for analyzing such scan data and producing models suitable for understanding enterprise scale vulnerability exposure and risk accordingly.

OPTIMAL SCHEDULING OF CYBERSECURITY ANALYSTS FOR MINIMIZING RISK

Cybersecurity threats are on the rise with evermore digitization of the information that many day-to-day systems depend upon. The demand for cybersecurity analysts outpaces supply, which calls for optimal management of the analyst resource. Therefore, a key component of the cybersecurity defense system is the optimal scheduling of its analysts. Sensor data is analyzed by automatic processing systems, and alerts are generated. A portion of these alerts is considered to be significant, which requires thorough examination by a cybersecurity analyst. Risk, in this paper, is defined as the percentage of unanalyzed or not thoroughly analyzed alerts among the significant alerts by analysts. We developed a generalized optimization model for scheduling cybersecurity analysts to minimize risk (a.k.a maximize significant alert coverage by analysts) and maintain risk under a pre-determined upper bound. We tested the optimization model and its scalability on a set of given sensors with varying analyst experiences, alert generation rates, system constraints, and system requirements. Results indicate that the optimization model is scalable, and is capable of identifying both the right mix of analyst expertise in an organization and the sensor-to-analyst allocation in order to maintain risk below a given upper bound. Several meta-principles were derived from the optimization model which serve as guiding principles for hiring and scheduling cybersecurity analysts.

DYNAMIC SCHEDULING OF CYBERSECURITY ANALYSTS FOR MINIMIZING RISK USING REINFORCEMENT LEARNING

We next developed a reinforcement learning-based stochastic dynamic programming optimization model for dynamically scheduling cybersecurity analysts to minimize risk and maintain risk under a pre-determined upper bound. We tested the dynamic optimization model and compared the results to an integer programming model that optimizes the staffing needs based on a daily-average alert generation rate with no estimation of future workload (static workforce model). Results indicate that the learning-based optimization model, through a dynamic (on-call) workforce in addition to the static workforce, is capable of reducing risk better than the static model, is scalable and capable of identifying the right mix of analyst expertise in an organization, and is able to determine their dynamic (on-call) schedule and their sensor-to-analyst allocation in order to maintain risk below a given upper bound. Several meta-principles were derived from the optimization model which may serve as guiding principles for hiring and scheduling cybersecurity analysts. Days-off scheduling was performed to determine analyst weekly work-schedules that met the cybersecurity system's workforce constraints and requirements.

CONTROL-THEORETIC REASONING TECHNIQUES

In Year 1, we had developed a general (supervisory) control-theoretic formulation of a dynamic cybersecurity problem, from the perspective of the defender. In this formulation, the defender aims to maximize its worst-case performance, considering that the attacker may behave arbitrarily. In Year 2, we built on that work, focusing on scalable representations and algorithms. We also developed a new model,

based on Bayesian attack graphs. In our control formulation, the defender's goal is to determine an optimal defense strategy given its current knowledge of the attacker's capabilities. The Bayesian attack graph itself models how an attack progresses through the graph according to a probabilistic spreading process. The defender summarizes the current system state into an information state which consists of a probability mass function on the possible system states. The optimal policy takes into account the future evolution of the system via a dynamic program.

The first paper on this approach (to appear in the 2nd ACM Workshop on Moving Target Defense, October 2015) accomplishes several goals, including the original formulation in terms of a partially observable Markov decision process (POMDP), and an algorithm to compute optimal defender countermeasures. The model studied so far is extensible in several directions, which we plan to pursue in follow-on work.

GAME-THEORETIC REASONING TECHNIQUES

With respect to game-theoretic reasoning, we extended our simulation model of the basic adaptive cyber-defense scenario explored in Year 1. We continue to pursue an empirical game-theoretic (EGT) approach, which uses results of extensive simulation to construct a formal game model. Our basic model can be viewed as a major extension of the FlipIt game, with multiple hosts, asymmetric stealth, and progressive attack.

Our paper on this approach (appeared in the 2nd ACM Workshop on Moving Target Defense, October 2015) reports results of the EGT study. Findings include:

1. With perfect probe detection, a maximal defense (aggressive application of moving target methods) is often effective at deterring attacks. When the detection capability degrades, however, the deterrent becomes much less effective and eventually evaporates.
2. Optimal policies typically involve monitoring and estimating the degree of compromise, and responding when it passes a threshold.
3. In addition to maximal defense, three other qualitative patterns are often found in equilibrium: maximal attack (aggressive attack makes defender give up), fight (both sides act vigorously, are unable to maintain objectives reliably), and share (both sides back off, satisfied with partial control). The share equilibria are only possible in environments where contention is relatively low.

These results are illustrative of the capabilities of EGT to identify qualitatively meaningful strategic patterns in complex environments, beyond the reach of analytic game-theoretic methods.

Based on the Year 2 effort, we have solid modeling baselines for both control-theoretic and game-theoretic reasoning tracks. We have also made significant progress (as yet unpublished) on reinforcement learning methods for deriving optimal responses to adversary behavior. Over the next year, these will be integrated into our game-theoretic reasoning techniques, as well as our control approaches.

MANIPULATING THE ATTACK SURFACE OF A SYSTEM

During Year 1, we looked at the problem of manipulating a system's attack surface from a control perspective and we proposed a graph-based approach to manipulate the attacker's perception of a system's attack surface. To achieve this objective, we formalized the notion of *system view* as well as the

notion of distance between views. We refer to the attacker's view of the system as the *external view* and to the defender's view as the *internal view*. A system's attack surface can then be thought of as the subset of the internal view that would be exposed to potential attackers when no deceptive strategy is adopted. Starting from these definitions, we developed a principled yet practical approach to manipulate responses to attacker's probes so as to induce an external view of the system that satisfies certain desirable properties.




Figure 7 - Control Perspective

Figure 6 shows how standard and ACD-based controllers can be adopted in order to alter the attacker's perception of the attack surface. State of the art solutions aim at dynamically reconfiguring the system in order to create a moving target for the attacker. The controller uses internal and external system variables to monitor the system's state and realize a moving target objective. On the other hand, our approach aims at transparently altering the attacker's perception. The attacker's knowledge of the system is inferred from the system's observable variables (e.g., ports, packets, web-pages content, DNS records). In order to gather additional information, the attacker probes the system. The proposed ACD controller identifies such probes and controls the information exposed through the system's observable variables by crafting adequate responses to bias the attacker's view of the system.

Compared to traditional approaches, our novel approach presents two key advantages: (i) we can introduce uncertainty for the attackers without actually changing the system's configuration, thus minimizing the overhead; (ii) we can deceive the attackers and steer them away from critical resources rather than simply introducing uncertainty and forcing them to use a random strategy.

First, building on the work we started in the previous year, we further investigated network diversity as a security property of networks that can guide further diversification of the attack surface. The interest in diversity as a security mechanism has recently been revived in various applications, such as Moving Target Defense (MTD), resisting worms in sensor networks, and improving the robustness of network routing. However, most existing efforts on formally modeling diversity have focused on a single system running diverse software replicas or variants. At a higher abstraction level, as a global property of the entire network, diversity and its impact on security have received limited attention. In our work, we took the first step towards formally modeling network diversity as a security metrics for evaluating the robustness

of networks against potential zero-day attacks. We have demonstrated that intuitive notions of diversity usually lead to misleading results, whereas our formal model of network diversity enables a better understanding of the impact of diversity on security. Specifically, we first devised a biodiversity-inspired metrics based on the effective number of existing distinct network resources. We then proposed two complementary diversity metrics, based on the least and the average attacking efforts, respectively. The most significant contributions during Year 2 include: (i) a new probabilistic model for addressing various limitations of the previous model; (ii) a study on how to instantiate the metrics, and in particular on how to collect inputs about software diversity; and (iii) a number of different realistic use cases, and a set of simulations for analyzing the proposed metrics under these different use cases. In our previous diversity model, modeling the effect of reusing exploits as a conditional probability that a resource may be exploited, given that other instances of the same type have already been exploited, essentially assumes a total order over different instances of the same resource type in any resource graph, which represents a major limitation. Intuitively, what allows an attacker to more likely succeed in exploiting a previously exploited type of resources is the knowledge, skills, or exploit code he/she has obtained. Therefore, instead of directly modeling the casual relationship between reused exploits, we explicitly model such advantages of the attacker as separate events, and model their effect on increasing the likelihood of success in subsequent exploits as conditional probabilities. This work has been submitted to IEEE Transactions on Information Forensics & Security.

Second, we proposed a proactive approach to strategically deploy detectors on selected network nodes, so as to either completely disrupt the communication between bots and command and control nodes (C2), or at least force the attacker to create more bots, therefore increasing the footprint of the botnet and the likelihood of detection. In our approach, we assume that the attacker can identify detector nodes, and can deploy bots in such a way to create detector-free paths through the network. However, since traffic is routed along the shortest path between source and destination, when a detector is deployed on the shortest path between a bot and a C2 site, the attacker will have to deploy additional bots to relay traffic exfiltrated by a bot to the C2 site in such a way to avoid detector nodes. In our approach, we leverage this mechanism to force the attacker to create a more complex botnet by strategically placing detectors on critical nodes. In the context of adaptive cyber defense, detector deployment may be periodically redefined to invalidate the attacker's current knowledge about detector-free paths. Clearly, the problem of optimally placing detectors to monitor a network is intractable. Therefore, we proposed heuristics based on several centrality measures, and this approach enables us to identify the most promising nodes to use as detectors in a time-efficient manner. Simulations results confirm that our approach can effectively increase complexity for the attacker. This work was presented at the 3rd IEEE Conference on Communications and Network Security (September 2015).

Third, within the general framework developed in Year 1, we developed a mechanism to disrupt an attacker's reconnaissance effort and protect critical systems by manipulating outgoing traffic in such a way to resemble traffic generated by a set of systems running different operating systems and services. In fact, in our work, we go beyond simply introducing uncertainty for the attacker, and propose an approach to deceive potential intruders into making incorrect inferences about important system characteristics, including operating systems and active services. Differently from many existing techniques, we do so without changing the actual configuration of the system. In fact, our approach

mainly consists in manipulating outgoing traffic such that, not only important details about operating systems and services are not revealed, but the traffic also resembles traffic generated by hosts and networks with different characteristics. Protocol specifications may leave some degrees of freedom to developers, and the specific choices that a developer makes may reveal the nature of the operating system or even the type of device. All the information required to impersonate a certain operating system or device can be extracted from the SinFP fingerprinting tool's signature database. All the outgoing packets that may reveal relevant information about the OS are modified to reflect the deceptive signature. Similarly, to defeat service fingerprinting, we replace the original service banner with a deceptive banner. Experiments conducted on a prototypal implementation show that the overhead introduced by the proposed approach is negligible, thus rendering this solution completely transparent to legitimate users. At the same time, our approach can effectively deceive the attackers, and steer them away from critical resources we wish to protect. This work will be presented at the 3rd IEEE Conference on Communications and Network Security (September 2015).

A PRACTICAL APPROACH FOR ADAPTIVE DATA STRUCTURE LAYOUT RANDOMIZATION

This work is motivated by two limitations of ASLR (Address Space Layout Randomization). ASLR is one of the most successful techniques of MTD (Moving Target Defense). The basic goal of ASLR is to ensure that the same malicious payload (e.g., buffer overflow payload), which contains code pointers (e.g., a function), can no longer be repeatedly used. Fine-grained ASLR is a set of fine-grained versions of ASLR in which the randomization unit is no longer a whole 64K segment.

However, both ASLR and fine-grained ASLR have two limitations. (a) They are vulnerable to memory content leakage attacks against both code addresses and data field layout inside a critical data structure. Some of the leakage is through side channels. Due to this limitation, the attacker can do data structure manipulation with memory content leakage in various applications. For example, launching such an attack against OpenSSH would enable the attacker to modify data field `pw_uid` inside data structure `passwd` and escalate the remote shell with root privilege. (b) They can be easily circumvented by rootkits launching DKOM attacks. Due to this limitation, a rootkit can launch DKOM attacks to achieve various malicious goals such as to hide one process, to hide a module, and to give the root privilege to certain process.

Besides ASLR, static DSLR techniques have also been developed in the literature. Unfortunately, static DSLR techniques also suffer from these two limitations. In this work, we developed (to the best of our knowledge) the first adaptive DSLR (Data Structure Layout Randomization) technique that can effectively address the two aforementioned limitations. We designed and implemented an adaptive DSLR system prototype, called SALADS, that automatically translates a program to a DSSR (Data Structure Self-Randomizing) program through an extended compiler. At runtime, a DSSR program dynamically randomizes the layout of each security-sensitive data structure by itself autonomously. DSSR programs regularly re-randomize a data structure when it has been accessed several times after last randomization. More importantly, DSSR programs automatically determine the randomizability of instances and randomize each instance independently (i.e., no need to consider how other data structures are randomized).

We have implemented SALADS based on gcc-4.5.0 and generated DSSR user-level applications, OS kernels, and hypervisors. Our experiments show that the DSSR programs can defeat a wide range of attacks with reasonable performance overhead. For applications, our experiments show that DSSR OpenSSL and OpenSSH programs can defeat attacks exploiting CVE-2002-0656 and CVE-2001-0144. For OS kernels, our experiments show that DSSR Linux kernels can defeat 12 real-world rootkits. For hypervisors, our experiments show that DSSR Xen hypervisors can defeat the Blue Pill attack.

Regarding performance overhead, for DSSR Linux kernel and DSSR Xen-3.2.0, we use the Lmbench to evaluate the performance overhead. Specifically, we measure the overhead with the bandwidth and the latency benchmarks. By only randomizing the selected data structures (around 20 data structures are selected), DSSR Linux kernel introduces 6.7% to 28.8% (16.7% on average) runtime overhead, and DSSR Xen-3.2.0 introduces 0.1% to 14.8% (4.5% on average) runtime overhead. For applications, we tested SPECint2000, httpd, OpenSSL and OpenSSH applications, and the overhead (we randomly select 20% of data structures) ranges from 0.2% to 23.5% on average.

A paper which summarizes this work is accepted into ESORICS 2015.

A NOVEL DEFENSE MECHANISM PROVIDING EFFECTIVENESS GUARANTEES AGAINST HEARTBLEED

According to the Community Emergency Response Team of U.S., almost 1/3 of newly published vulnerabilities are related to buffer overflow; e.g., buffer over-write and over-read. The Heartbleed bug is a recently disclosed buffer over-read vulnerability. A number of countermeasures have been proposed to defend against buffer over-read attacks and some of them include safety-type system retrofitting, bounds checking and adding guard pages. However, to the best of our knowledge, all the existing countermeasures are ad hoc and lack of mathematically provable effectiveness. This would significantly lower the trustworthiness of the countermeasures.

In order to overcome this crucial shortcoming, we proposed the co-design of mathematical algorithms and software systems against the Heartbleed attack. In particular, we came up with a defense scheme based on a modified stochastic bandit algorithm. This scheme enables the defender to dynamically and adaptively adjust guard pages in a heap in response to measured utilities. In particular, at the beginning of each defense cycle, the defender first measures the utility which quantifies the effectiveness of the defense action, a set of guard pages, in the last defense cycle. The defender then chooses to deploy the defense action whose empirical utility in the past is highest. Notice that the defense deployment is only based on measured utilities in the history, and these quantities do not require the defender to access the actions and targets of the attacker. So our proposed defense scheme is suitable for handling zero-day attacks. In addition, the proposed defense scheme is robust to the delays and errors of the utility reports. Most importantly, in contrast to all the existing countermeasures, the cost effectiveness of our proposed defense scheme is mathematically provable. (We have successfully developed the mathematical proofs.) In particular, we used the regret as the effectiveness metric, and the regret is defined as the distance between the aggregate utility caused by the optimal defense action and that caused by our proposed scheme. We have successfully developed the mathematical proofs through which we formally showed

that the regret is in the order of $\log N$ where N is the number of defense cycles. It should be noted that the regret in the order of N corresponds to the worst defense scheme.

In order to achieve the provable effectiveness, we found that we need to modify the Linux memory allocator in a specific way, since the current Linux Malloc memory allocator is not mathematically equal to the abstract system model used in our mathematical analysis. Through mathematical equivalence analysis, we have identified the set of properties that an equivalent Linux memory allocator must have. Based on these properties, we have designed and implemented a different memory allocator. Although it is currently not very space efficient, its extra runtime overhead is only 20% to 30%.

A paper which summarizes this work is going to be submitted to NDSS 2016 in August 2015.

AN INNOVATIVE SYMBOLIC N-VARIANT SYSTEM FOR PRACTICAL N-VARIANT SYSTEM MTD AGAINST ZERO-DAY ATTACKS

Based on a 2015 expert survey (around 20 experts participated in this survey) conducted by researchers at Dartmouth College, N-Variant System MTD was ranked in top 7 among over 35 MTD techniques covered by the survey. Among the top 7 MTD techniques, N-Variant System MTD is having the same level of effectiveness (i.e., medium effectiveness) with 5 others, but its costliness was ranked the worst. This survey result indicates that if the costliness of N-Variant System MTD can be well addressed, N-Variant System MTD may become a highly competitive top 3 MTD choice for cyber-security professionals.

Nevertheless, conventional N-Variant Systems are by nature more costly than other peer MTD techniques, since several diversified replicas (of the same system/application) are required to run simultaneously online. Hardware virtualization could be leveraged to let several replicas share the same hardware resources, but virtualization does not change the fact that the total resource demands of an N-Variant System is roughly N times as large as a single system/application.

In this work, we have been developing an innovative symbolic N-Variant system to significantly reduce the cost of N-Variant System MTD. The proposed system is dramatically different from conventional N-Variant Systems. Replicas are no longer required online. In contrast, given a concrete malicious “seed” input observed in the wild, we will firstly do offline symbolic execution against diversified replica programs to explore hopefully all the potential attack-related execution paths. The exploration will end as soon behavior deviation is detected. At this moment, the corresponding symbolic path specifications will be recorded.

In runtime, our symbolic N-Variant system will firstly use a light-weight filter to filter off almost all the benign inputs. Next, the malicious (mutant) inputs that belong to the same “family” as the “seed” input (if there are any) and a few benign inputs will be matched against the symbolic path specifications recorded during the offline analysis. Whenever a match is found, we immediately know the behavior deviation, because the deviation was already observed during the offline symbolic execution. Since our symbolic N-Variant system is essentially a match finder in runtime, it can simultaneously “simulate” multiple real N-Variant systems. In this way, the cost of N-Variant system MTD could be significantly reduced.

In runtime, our symbolic N-Variant system run together with the real applications. The real applications no longer need any real replicas. If the real applications run faster than the simulator, our approach can do intrusion detection, but not malicious input blocking. If our simulator runs faster than the real applications, malicious input blocking could be achieved.

As of August 2015, we have developed a system prototype of the simulator, and we have implemented a set of optimizations on scalability. Our preliminary experiments show that the simulator has good scalability. Nevertheless, the effectiveness and wide applicability of our approach are still under investigation.

4. SCIENTIFIC PROGRESS IN YEAR 3

During Year 3, the MURI team, building on the foundations defined during previous years, focused on developing new MTD techniques and a formal quantification framework to assess the cost and performance of individually or jointly deployed Moving Target Defense techniques. This effort resulted in a number of scientific accomplishments that were disseminated nationally and internationally through presentations and publications in peer-reviewed venues.

MTD QUANTIFICATION

Despite the significant work done in this area and the many different techniques that have been proposed, MTD has not yet gained widespread adoption due to several limitations. Specifically, interactions between multiple techniques have not been studied yet and a unified framework for quantifying and comparing very diverse techniques is still lacking. To overcome these limitations, we developed a framework to model how different MTD techniques can affect the information an attacker needs to exploit a system's vulnerabilities, so as to introduce uncertainty and reduce the likelihood of successful attacks. This framework can be used to compare two sets of MTDs, and to select an optimal set of MTDs that maximize security within a given budget.

As shown for the simple example of Figure 1, the framework consists of four layers: (i) a time-invariant service layer representing the set $\mathcal{S}$ of services to be protected; (ii) a weakness layer representing the set $\mathcal{W}$ of general classes of weaknesses that may be exploited; (iii) a knowledge layer representing the set $\mathcal{K}$ of all possible knowledge blocks required to exploit those weaknesses; and (iv) an MTD layer representing the set $\mathcal{M}$ of available MTD techniques. Therefore, the proposed MTD quantification framework can be formally defined as a 7-tuple $(\mathcal{S}, \mathcal{R}_{SW}, \mathcal{W}, \mathcal{R}_{WK}, \mathcal{K}, \mathcal{R}_{KM}, \mathcal{M})$, where: (i) $\mathcal{R}_{SW} \subseteq \mathcal{S} \times \mathcal{W}$ represents relationships between services and the common weaknesses they are vulnerable to; (ii) $\mathcal{R}_{WK} \subseteq \mathcal{W} \times \mathcal{K}$ represents relationships between weaknesses and the knowledge blocks required for an attacker to exploit them; and (iii) $\mathcal{R}_{KM} \subseteq \mathcal{K} \times \mathcal{M}$ represents relationships between knowledge blocks and the MTD techniques that affect them. The proposed model induces a k -partite graph (with $k = 4$) $G = (\mathcal{S} \cup \mathcal{W} \cup \mathcal{K} \cup \mathcal{M}, \mathcal{R}_{SW} \cup \mathcal{R}_{WK} \cup \mathcal{R}_{KM})$.




Figure 8 - MTD quantification framework: an example

The framework enables to decouple MTDs from the specific vulnerabilities they are designed to address through the knowledge layer, which assumes the role of unifying how we can handle very diverse MTD techniques. Experimental results have shown that our approach is effective in quantifying MTDs and driving the selection of an optimal set of MTDs.




Figure 9 - Analytic model

We also developed quantitative analytic model for assessing the resource availability and performance of MTDs, and a method for determining the reconfiguration rate that maximizes a utility function that incorporates the tradeoffs between the attacker's success probability and response time. Our analytic model is based on queuing theory and can be further divided, as shown in Figure 2, into two parts: a *reconfiguration model* for reconfiguration requests generated by an MTD and a *performance model* for service requests. The reconfiguration model takes as inputs the target rate α at which resources are reconfigured, the average reconfiguration time S , the number of resources c , and the maximum number of resources c^* that can be reconfigured at the same time. This model produces as outputs the availability of resources, the average number of resources available, and the probability distribution p_k of the number of resources being reconfigured. This distribution, along with the number of resources, the maximum number of resources that can be reconfigured at the same time, the average arrival rate λ of service requests, and the average service time T of requests are inputs to the performance model, which produces the distribution P_k of service requests in the system and the average response time of requests.

The reconfiguration and performance models are solved using Continuous Time Markov Chains and are combined iteratively.

A MOVING TARGET DEFENSE APPROACH TO MITIGATE DDoS ATTACKS AGAINST PROXY-BASED ARCHITECTURES

Distributed Denial of Service attacks against high-profile targets have become more frequent in recent years. In response to such massive attacks, several architectures have adopted proxies to introduce layers of indirection between end users and target services and reduce the impact of a DDoS attack by migrating users to new proxies and shuffling clients across proxies so as to isolate malicious clients. However, the reactive nature of these solutions presents weaknesses that we leveraged to develop a new attack – the proxy harvesting attack – which enables malicious clients to collect information about a large number of proxies before launching a DDoS attack. We show that current solutions are vulnerable to this attack, and propose a moving target defense technique consisting in periodically and proactively replacing one or more proxies and remapping clients to proxies. Our primary goal is to disrupt the attacker's reconnaissance effort. Additionally, to mitigate ongoing attacks, we propose a new client-to-proxy assignment strategy to isolate compromised clients, thereby reducing the impact of attacks. We validate our approach both theoretically and through simulation, and show that the proposed solution can effectively limit the number of proxies an attacker can discover and isolate malicious clients.

MOVING TARGET DEFENSE AGAINST DDoS ATTACKS: AN EMPIRICAL GAME-THEORETIC ANALYSIS

Distributed denial-of-service attacks are a rampant problem facing web applications, for which many defense techniques have been proposed, including several moving-target strategies. Proposed moving-target strategies to prevent denial-of-service attacks typically work by relocating targeted servers over time, increasing uncertainty for the attacker while trying not to disrupt normal users or incur excessive costs. Prior work has not shown, however, whether and how a rational defender would choose a moving-target method against an adaptive attacker, and under what conditions. We formulate a denial-of-service scenario as a two-player game, and solve a restricted-strategy version of the game using the methods of empirical game-theoretic analysis. Using agent-based simulation, we evaluate the performance of strategies from prior literature under a variety of attacks and environmental conditions. We find evidence for the strategic stability of various proposed strategies, such as proactive server movement, delayed attack timing, and suspected insider blocking, along with guidelines for when each is likely to be most effective.

EFFICIENT COMPUTATION OF CENTRALITY MEASURES

The identification of important nodes is a ubiquitous problem in the analysis of social and computer networks. Centrality indices (such as degree centrality, closeness centrality, betweenness centrality, PageRank, and others) are used across many domains to accomplish this task. However, the computation of such indices is expensive on large graphs. Moreover, evolving graphs are becoming increasingly important in many applications. It is therefore desirable to develop on-line algorithms that can

approximate centrality measures using memory sublinear in the size of the graph. In particular, we apply recent advances in the streaming and sketching literature to provide a preliminary streaming approximation algorithm for degree centrality utilizing CountSketch and a multi-pass semi-streaming approximation algorithm for closeness centrality leveraging a spanner obtained through iteratively sketching the vertex-edge adjacency matrix. We also discuss possible ways forward for approximating betweenness centrality, as well as spectral measures of centrality. We have provided a preliminary result using sketched low-rank approximations to approximate the output of the HITS algorithm.

CYBER RESILIENCY

Cyber resiliency has become an increasingly relevant and timely research and operational concept in cyber security. While several metrics have been proposed for quantifying cyber resiliency, a considerable gap remains between those metrics and operationally measurable and meaningful concepts that can be empirically determined in an objective manner. This work describes a concrete quantitative notion of cyber resiliency that can be tailored to meet specific needs of organizations that seek to introduce resiliency into their assessment of their cyber security posture.

VULNERABILITY REMEDIATION

Vulnerability remediation is a critical task in operational software and network security management. In this work, an effective vulnerability management strategy, called VULCON (VULnerability CONtrol), has been developed and evaluated. The strategy is based on two fundamental performance metrics: i). Time-to-Vulnerability Remediation (TVR) and; ii). Total Vulnerability Exposure (TVE). VULCON takes as input real vulnerability scan reports, metadata about the discovered vulnerabilities, asset criticality and personnel resources. VULCON uses a mixed integer multi-objective optimization algorithm to prioritize vulnerabilities for patching, such that the above performance metrics are optimized subject to the given resource constraints. VULCON has been tested on multiple months of real scan data from a Cyber-Security Operations Center (CSOC). Results indicate an overall Total Vulnerability Exposure reduction of 8.97% when VULCON optimizes a realistic security analyst workforce's effort. Additionally, we show how VULCON can determine monthly resources required to maintain a target TVE score. As such, VULCON provides valuable operational guidance for improving vulnerability response processes in CSOCs.

5. SCIENTIFIC PROGRESS IN YEAR 4

During Year 4, the MURI team, building on the foundations defined during previous years, focused on developing an array of new game-theoretic, control-theoretic, and graph-based adaptive techniques to address a range of different security threats. This effort resulted in a number of scientific accomplishments that were disseminated nationally and internationally through presentations and publications in peer-reviewed venues.

GAME-THEORETIC LEARNING FOR ADAPTIVE CYBER DEFENSE

Game theory provides a mathematically rigorous framework for multiple strategic players to reason about each other. A number of learning algorithms have been proposed to compute Nash equilibrium, including best-response dynamics, better-response dynamics, fictitious play, regret matching, logit-based

dynamics and replicator dynamics. In security, the defender and the attacker can only access limited information about the game of interest. For example, each player may not know the structure of its own utility function. Additionally, during repeated interactions, each player may not be aware of the actions of the other. This imposes additional challenges for the design of game-theoretic learning algorithms.

To address the challenges, we have developed a game-theoretic learning algorithm, which allows a defender and an attacker to identify in probability the set of action profiles with maximal stochastic potential. More specifically, at each iteration, each player, on one hand, exploits successful actions in recent history via comparing received utility values, and on the other hand, randomly explores any feasible action with a certain exploration probability. The proposed algorithm requires minimal prior information for the defender and the attacker. In particular, each player knows its own past actions but is unaware of the other's deployed actions. Each player is unaware of both utility functions and only able to access utility values given a pair of actions. When the interactions consist of a weakly acyclic game, the convergence to the set of Nash equilibria is guaranteed. The convergence rate is explicitly quantified using tools from Markov chains and resistance trees. The proposed algorithm is validated by the real-world security scenario of dynamic platforms. In particular, a server contains several zero-day security vulnerabilities. A zero-day attack happens once that a software/hardware vulnerability is exploited by the attacker before software the engineers develop any patch to fix the vulnerability. The attacker is equipped with a set of zero-day attack scripts and the defender is equipped with a set of platforms. The defender uses a defensive technique called dynamic platforms, which changes the properties of the server such that it is harder for the attacker to succeed.

A paper, which summarizes the above results, was submitted to *Automatica*, a top journal in the field of automatic control.

ADAPTIVE ASLR (ADDRESS SPACE LAYOUT RANDOMIZATION)

With the broad deployment of W XOR X and Data Execution Prevention (DEP), the effectiveness of simple code injection attacks like worm wanes. In its place rises a more sophisticated attack, Return-Oriented Programming (ROP). An ROP attack hijacks a target machine by gaining control of the call stack and executing carefully chosen machine instruction sequences (i.e., "gadgets"). This attack exploits the fact that the gadget locations in memory are deterministic and an attacker can easily chain them to perform arbitrary operations.

Address space randomization is a fundamentally important MTD technique that obfuscates the locations of gadgets by randomizing the memory layout of a process. As such, it has been shown to be effective for the defense against ROP. However, recent research indicates this technique is vulnerable to Just-In-Time ROP (JIT-ROP) in which an attacker employs memory disclosure attacks to locate gadgets in memory and thus circumvents address randomization.

To counteract JIT-ROP, existing defense mechanisms either restrict memory execution or perform adaptive randomization. To some extent, these solutions are effective in terms of obstructing JIT-ROP. However, they are impractical for real-world applications and systems. From a correctness standpoint, existing adaptive randomization techniques do not guarantee program correctness especially when dealing with binary code. For example, Remix and TASR need to locate code pointers for randomization.

However, a prior study has demonstrated that it is generally impossible to correctly locate all code pointers in memory. From a performance standpoint, existing adaptive randomization techniques like Isomeron and RuntimeASLR perform dynamic code instrumentation which introduces significant overhead at run time.

In this project, we have developed CHAMELEON, a novel practical mechanism to counteract JIT-ROP. Our basic idea is to obfuscate the address of code pages, making code exploit infeasible. In JIT-ROP, an attacker typically exploits memory disclosure vulnerability, and thus obtains the address of code pages and corresponding code instructions. Using this information, the attacker then constructs ROP payload on call stack and performs arbitrary operations. In this work, we randomize the address of code pages through an address space transformation mechanism. In this way, the attacker is no longer able to obtain the actual address of a code page, dump the code instructions and thus construct a valid ROP payload.

Innovation claim. Technically speaking, CHAMELEON draws on address randomization. However, it is fundamentally different from existing address space randomization mechanisms. In fact, existing randomization techniques do not hinder code disclosure but impede malicious ROP payload to perform arbitrary operations. In contrast, our randomization mechanism serves as the first-line defense that counteracts attacks by obstructing the code page disclosure directly. In terms of significance of this work, CHAMELEON is, to the best of our knowledge, the first practical adaptive ASLR technique for real-world applications and systems. CHAMELEON ensures program correctness, and has acceptable performance overhead.

During Year 3 of the MURI project, we have built a preliminary system prototype of CHAMELEON. During this year, we have substantially refined the system design, substantially solidified the evaluation by running new sets of experiments, and significantly improved the paper write-up. We have implemented a prototype of CHAMELEON on Linux operating system and extensively experimented it in different settings. In particular, we have conducted 16 representative test cases, including three web service applications (Nginx 1.6.0, httpd 2.2.27 and lighttpd 1.4.35), one compression program xz 5.1.4 and twelve benchmark programs from SPEC CPU2006. The results show that CHAMELEON is very effective in thwarting Just-in-Time ROP attacks; the memory overhead is less than 4% for all tested programs except for one; and the performance overhead is in general less than 8%.

A paper, which summarizes the above results, has been published on the Proceedings of IEEE DSN 2017, a top conference in dependable and secure systems.

COOPERATIVE ATTACK SURFACE “DIVIDE AND CONQUER” FOR MEMORY PROTECTION

In recent years, more and more large-scale cyberattacks have exploited a zero-day memory corruption vulnerability to compromise a large amount of servers and service program instances. Since the attacker’s goal is not a particular server/program, we call such attacks a greedy memory corruption attack. The Heartbleed attack is a representative greedy attack. In this work, we propose an innovative cooperative attack surface “divide and conquer” approach to defend against greedy memory corruption attacks.

Innovation claim. This approach is, to the best of our knowledge, the first defense that uses CCID (Calling Context ID numbers) to define the attack surface. It is also the first defense that conducts cooperative CCID-defined attack surface “divide and conquer”.

In this work, we aim to effectively defend against greedy zero-day memory corruption attacks. Since such attacks try to compromise a large amount of server program instances, it is reasonable to let a large group of instances of the same program (e.g., Apache web server) running in a server farm to cooperate and do attack surface “divide and conquer”. In this way, no single machine or virtual machine will suffer from significant performance degradation. In particular, we aim to provide three essential properties in the server farm: (P1) under a greedy zero-day heap overflow attack, at least one program instance is able to detect the attack; (P2) detection on a single instance is adequate to automatically generate a CCID-based shield which can be immediately deployed on all the other instances; (P3) the performance penalty incurred by applying the cooperative defense on individual instances is low.

We have designed a system that can provide the 3 security properties. This system defines the attack surface as the set of all possible CCIDs; the system divide the attack surface into multiple subsets and assign each subset of CCIDs to a program instance; each program instance is only responsible for protecting the calling contexts assigned to itself; the system then inserts guard pages at run-time to guard the heap buffers dynamically allocated in each assigned calling context; as soon as there is a guard-page-related segment fault, the zero-day attack is detected; then the system will use a technique we previously proposed to generate a simple CCID-based shield; since the shield is orthogonal to the execution states of any program instance, the shield can be pervasively deployed; as soon as the shield is deployed on all the program instances, the zero-day attack will fail to harvest any new compromises.

During this year, we have proposed the new CCID-based definition of the attack surface; we have designed and developed a preliminary system prototype of the new attack surface “divide and conquer” defense; and we have conducted the preliminary evaluation based on analytical results.

APPLYING DEEP LEARNING TO PREDICT HOW THE ATTACKER CAN BYPASS CFI (CONTROL-FLOW INTEGRITY)

With the broad deployment of W XOR X and Data Execution Prevention (DEP), the effectiveness of simple code injection attacks like worm wanes. In its place rises a more sophisticated attack, Return-Oriented Programming (ROP). An ROP attack hijacks a target machine by gaining control of the call stack and executing carefully chosen machine instruction sequences (i.e., “gadgets”).

Control flow integrity (CFI) is a defense mechanism that can mitigate ROP attacks from arbitrarily controlling program behavior. In particular, it confines all control flow transfers to follow a valid path defined in a static control flow graph (CFG). Theoretically, CFI can prevent any ROP attack. However, theoretically safe CFI is not a practical solution in the real world. As is demonstrated in the literature, fine-grained CFI provides strong security guarantee but at the same time introduces performance overhead significantly high; e.g., 52% in CFI deployment. In real world, many practical CFI implementations have been proposed, such as clang compiler and Microsoft Visual Studio, however many sophisticated ROP attacks can bypass CFI implementations.

In this work, we consider the attacker’s target machine is protected by some practical CFI implementation which is unknown to the attacker. To successfully launch ROP attacks, the attacker must bypass the CFI implementation. It is not difficult for the attacker to generate ROP gadgets. There are several tools which can help the attacker generate ROP gadgets. Then the attacker can easily chain the gadgets based on

his/her purposes. But how to guarantee that the gadget chains can bypass the CFI implementation is not an easy job. To determine whether a gadget chain can bypass the CFI implementation, the attacker usually needs the knowledge of the CFI implementation. If the knowledge is unavailable to the attacker, he/she usually uses binary analysis to recover the CFG and simulates the CFI implementation based on a finite, static CFG. However, recovering CFG from binary files is very time consuming and the recovered CFG is usually imprecise and incomplete.

Innovation claim. This work is, to the best of our knowledge, the first work that applies deep learning to predict how the attacker could bypass CFI. In particular, from the perspective of adversary reasoning, we aim to gain new understanding about (a) whether deep learning could be effectively used to help the attacker address the above challenges when bypassing CFI defenses; (b) whether there are enough incentives for the attacker to apply deep learning in launching ROP attacks.

During this year, we have developed the problem formulation, and conducted a preliminary study towards achieving the aforementioned project goal. In particular, we have used cutting-edge tools to generate massive amount of gadget chains; we have used static analysis to automatically generate the labels for each data item in the training dataset; we have implemented our DNN (Deep Neural Network) in Python using Tensorflow and its high-level library Keras; and we have conducted preliminary evaluation of the trained DNN. Our empirical evaluation results indicate that the trained DNN can accurately and efficiently classify huge amount of gadget chains into two categories: can bypass CFI vs. cannot bypass CFI. The classification accuracy is over 93%, and the trained DNN outperforms the classic SVM approach.

A technical report, which summarizes the above results, has been fully written.

A POMDP APPROACH TO THE DYNAMIC DEFENSE OF LARGE-SCALE CYBER NETWORKS

We investigated the problem of optimally mitigating the progression of an adversary through a network in real-time, decreasing the probability of the attacker reaching its goal(s) while minimizing the negative impact to availability. Our model is based on a dependency graph which models the dependencies between security conditions (attacker capabilities) and exploits. The dependency graph describes how an attacker can use its capabilities to perform exploits and progressively move through the network towards its goal(s). The defender is able to interfere with this progression by blocking some exploits from being carried out. The true progression of the attacker at any given time is not observable by the defender and must be inferred using security alerts (generated by an intrusion detection system) and previously deployed defense actions. The security alerts are noisy, subject to both missed detections and false alarms. The resulting problem of determining the optimal defense actions over time is formulated as a partially observable Markov decision process (POMDP). Due to the scale of the defense problem, obtaining an optimal solution is computationally intractable. As a result, we make use of an existing online solution method, termed the partially observable Monte-Carlo planning (POMCP) algorithm. The algorithm samples future possible scenarios from the current belief in order to select actions, avoiding the state-space explosion problem. The behavior of the computed defense policy is demonstrated on an illustrative example. The scalability of the algorithm is further improved by taking advantage of the specific structure of our model.

DYNAMIC GAMES WITH ASYMMETRIC INFORMATION

We formulate and analyze a general class of stochastic dynamic games with asymmetric information arising in dynamic systems. In such games, multiple strategic agents control the system dynamics and have different information about the system over time. Because of the presence of asymmetric information, each agent needs to form beliefs about other agents' private information. Therefore, the specification of the agents' beliefs along with their strategies is necessary to study the dynamic game. We use Perfect Bayesian equilibrium (PBE) as our solution concept. A PBE consists of a pair of strategy profile and belief system. In a PBE, every agent's strategy should be a best response under the belief system, and the belief system depends on agents' strategy profile when there is signaling among agents. Therefore, the circular dependence between strategy profile and belief system makes it difficult to compute PBE. Using the common information among agents, we introduce a subclass of PBE called common information based perfect Bayesian equilibria (CIB-PBE), and provide a sequential decomposition of the dynamic game. Such decomposition leads to a backward induction algorithm to compute CIB-PBE. We illustrate the sequential decomposition with an example of a multiple access broadcast game. We prove the existence of CIB-PBE for a subclass of dynamic games.

STOCHASTIC DYNAMIC GAMES WITH DELAYED SHARING INFORMATION STRUCTURE

We formulate and analyze dynamic games with d -step ($d \geq 1$) delayed sharing information structure. The resulting game is a dynamic game of asymmetric information with hidden actions, imperfect observations, and controlled and interdependent system dynamics. We adopt common information based perfect Bayesian equilibrium (CIB-PBE) as the solution concept, and provide a sequential decomposition of the dynamic game. Such a decomposition leads to a backward induction algorithm to compute CIB-PBEs. We discuss the features of our approach to the above class of games and address the existence of CIB-PBEs.

A STACKELBERG GAME MODEL FOR BOTNET TRAFFIC EXFILTRATION

Cyber-criminals can distribute malware to control computers on a networked system and leverage these compromised computers (i.e., botnets) to perform their malicious activities inside the network. Botnet-detection mechanisms, based on a detailed analysis of network traffic characteristics, provide a basis for defense against botnet attacks. In this work, we formulate the botnet defense problem as a Stackelberg security game, allocating detection resources to deter botnet attacks considering the strategic response of cyber-criminals. Based on the new game model, we propose a game-theoretic algorithm, ORANI, to compute an optimal detection resource allocation strategy in zero-sum game settings. Our algorithm employs the double-oracle method to deal with an exponential number of players' actions. Furthermore, we provide greedy heuristics to approximately compute an equilibrium of these botnet defense games. Finally, we conduct extensive experiments based on both simulated and real-world network topologies to demonstrate advantages of our game-theoretic solution compared to previously proposed defense policies.

DEEP LEARNING

Machine learning, especially deep learning, is revolutionizing how many engineering problems are being solved. Three critical ingredients are needed to apply deep machine learning to significant real world problems: (i) large data sets; (ii) software to implement deep learning and; (iii) significant computing cycles. This effort addresses the state of each ingredient with a specific focus on: (a) how deep learning can apply to large-scale network analysis and; (b) the computing resources required to make such analyses feasible.

6. SCIENTIFIC PROGRESS IN YEAR 5

During Year 5, the MURI team – building upon the technical accomplishments achieved in previous years – primarily focused on integrating results from different lines of research as well as on expanding work in the most promising research areas. Our ultimate and overarching goal during this 5-year effort has been that of developing the understanding and the analytical tools to enable researchers and practitioners to develop novel and effective ACD techniques and evaluate both new and existing techniques. Our work on quantification and metrics is a core element of this vision, and it received significant attention during Year 5. In particular, we extended our work on quantification in order to characterize the behavior of the most sophisticated threats we are currently facing, including Advanced Persistent Threats.

In an effort to better capture attacker and defender dynamics, we developed graph-based, control-theoretic, and game-theoretic models, which have guided the development of an array of new techniques designed to address a range of different security threats, including Data Structure Manipulation Attacks (DSMA) and Return-Oriented Programming (ROP). We also developed a vulnerability management strategy to dynamically prioritize remediation efforts based on available resources and severity of the exposed vulnerabilities.

Our effort resulted in a number of scientific accomplishments that were disseminated nationally and internationally through 30 publications in peer-reviewed journals and conference proceedings, a book, and several book chapters.

In Year 5, we primarily focused on integrating results achieved in different research thrusts into a framework that will guide researchers and practitioners in the development of novel and effective ACD techniques and in the evaluation of both new and existing techniques. We also expanded our efforts in the most promising research areas. The MTD quantification framework developed in Years 3 and 4 represents a core element of our broader vision. To achieve our ultimate goal, we have further expanded this framework to address limitations and gaps that have been identified at this stage, and we have refined the different approaches to adaptive cyber defense that have been developed so far, including approaches based on game and control theory. Additionally, we have explored new and promising research directions, such as mitigation of Advanced Persistent Threats. In the following subsections, we provide details about specific goals, results, and future plans for each of the key research thrusts.

QUANTIFICATION AND EVALUATION OF ADAPTIVE CYBER DEFENSES

Our framework for quantifying moving target defenses (developed in Years 3 and 4) can accommodate any existing or future MTD, regardless of their specific nature. Our multi-layered approach captures the relationship between MTDs and the knowledge blocks they are designed to protect and the relationship between knowledge blocks and generic classes of weaknesses that can be exploited using that knowledge. This framework can be used to compare two sets of MTDs, and to select an optimal set of MTDs that maximize security within a given budget. Although this work represented a significant step towards effective MTD quantification, several limitations still existed and were addressed as part of our Year 5 effort or are being addressed as part of our ongoing work, as detailed below.

- The original method for computing the probability that an MTD would disrupt a specific block of attacker's knowledge only provided rough estimates. Therefore, we developed a better understanding about an MTD's ability to disrupt specific knowledge blocks and showed how this capability can be effectively quantified.
- To validate the model in more realistic scenarios, we developed an experimental testbed within our computing infrastructure and deployed multiple MTDs to study their performance both in isolation – in order to determine the probability $P_{\{i,j\}}$ that MTD i can disrupt knowledge block j – and in combination – in order to accurately study their interactions and their combined cost and performance.
- Originally, we adopted a very simple notion of cost, and just used an overall budget as an additional constraint. We have explored a more sophisticated notion of cost that takes into account both the cost of implementing an MTD and the impact on availability. Our results show that there is a tradeoff between the frequency of reconfiguration of any MTD and the availability of the resources being reconfigured. Our work provide insights on how to fine tune the frequency of reconfiguration to strike a balance between security requirements and system performance.
- We have modeled and analyzed a formulation of an adaptive attack-defend situation in which attackers seek to exploit the least cost/effort attack path to achieve a goal while the defender has a fixed budget to invest in hardening their systems. This creates a type of “network interdiction” problem or game, whose solution in the limit as the defender's budget increases is shown to be determined by minimal cut sets in the attack graph, although minimal cut sets may not play a role in system hardening initially.
- Rapid on-line estimation of network and system situational awareness suitable for adaptive, dynamic security management typically depends on streaming data collected from various sensors and logs deployed across the system being defended. We have been extending and applying recently discovered sketching algorithms for estimating various graph properties relevant to security estimation and management.

Our ongoing and future efforts on this research thrust are described below.

- Our model aims at disrupting an attacker's knowledge in the reconnaissance phase of the cyber kill chain. While this may be the most cost-effective approach to cyber security, no defense is perfect, and we must ensure that multiple layers of defense exist. For instance, MTDs can also be designed to disrupt an attacker's ability to maintain a foothold in the system, and we have started

exploring this direction as part of our work on mitigation of Advanced Persistent Threats (see next subsection). Thus, we plan to extend our framework to model this additional class of MTDs.

- Our framework originally assumed independence between different services. We are now exploring how to capture the relationships between services and the MTDs that can protect them from multi-step attacks.
- Finally, we plan to run additional and extensive experiments on our testbed to validate all theoretical results achieved as part of this research effort.

QUANTIFICATION AND MITIGATION OF ADVANCED PERSISTENT THREATS

Advanced Persistent Threats (APTs) have emerged as increasingly sophisticated cyber-attacks, often waged by state actors or other hostile organizations against high-profile targets. APT actors employ a diversified set of sophisticated tools and advanced capabilities to penetrate target systems, evade detection, and maintain a foothold within compromised systems for extended periods of time. Stealth and persistence enable APT actors to conduct long-term espionage and sabotage operations. Despite significant efforts to develop APT detection and mitigation capabilities, the stealthy nature of APTs poses significant challenges, and defending from such threats is still an open research problem. In particular, quantitative models to capture how APTs may create and maintain a foothold within a target system were still lacking. To address this gap, we developed a quantitative framework to (i) assess the cost incurred by APT actors to compromise and persist within a target system; (ii) estimate the value they gain over time by persisting in the system; (iii) simulate how the footprint of an APT evolves over time when, to maintain stealth, attackers have constraints on the volume of potentially detectable activity they can engage in. We also proposed a preliminary defender model, and results from the evaluation show that our approach is promising, thus encouraging further research in this direction.

Our ongoing and future efforts on this research thrust include extending the model to account for various factors that can cause the value of an asset to decrease over time. Originally, we assumed that the value of a node remains constant over time. However, some nodes may quickly become less valuable to the attacker after the initial exploitation. Similarly, the defender may move data between nodes, rendering some nodes less valuable. Therefore, the attacker may choose to relinquish such nodes and use the available budget to compromise new and more valuable nodes. Our extended model will capture such dynamics.

DEFEATING DATA STRUCTURE MANIPULATION ATTACKS

In the wake of the research community gaining deep understanding about control-hijacking attacks, data-oriented attacks have emerged. Among data-oriented attacks, data structure manipulation attack (DSMA) is a major category. Pioneering research was conducted and shows that DSMA is able to circumvent the most effective defenses against control-hijacking attacks — DEP, ASLR and CFI. Up to this day, only two defense techniques have demonstrated their effectiveness: Data Flow Integrity (DFI) and Data Structure Layout Randomization (DSLRL). However, DFI has high performance overhead, and dynamic DSLRL has two main limitations. **L-1:** Randomizing a large set of data structures will significantly affect the performance. **L-2:** To be practical, only a fixed sub-set of data structures are randomized. In the case that the data structures targeted by an attack are not covered, dynamic DSLRL is essentially ineffective.

To address these two limitations, we propose a novel technique, *feedback-control-based adaptive DSLR* and build a system named SALADSPPlus. SALADSPPlus seeks to optimize the trade-off between security and cost through feedback control. Using a novel feedback-control-based adaptive algorithm extended from the Upper Confidence Bound (UCB) algorithm, the defender (controller) uses the feedbacks (cost-effectiveness) from previous randomization cycles to adaptively choose the set of data structures to randomize (the next action). Different from dynamic DSLR, the set of randomized data structures are adaptively changed based on the feedbacks. To obtain the feedbacks, SALADSPPlus inserts canary in each data structure at the time of compilation. We have implemented SALADSPPlus based on gcc-4.5.0. Experimental results show that the runtime overheads are 1.8%, 3.7%, and 5.3% when the randomization cycles are selected as 10s, 5s, and 1s respectively. In next year, we will study the more powerful attack model where the attacker can observe past defense actions and choose attack actions using arbitrary strategies. For this case, we plan to leverage the technique of adversarial bandit.

BAYESIAN ATTACK GRAPHS

We consider the ACD problem in a computer network where an intelligent attacker exploits combinations of multiple known or zero-day vulnerabilities to compromise machines in the network. The defender faces two challenges: (1) partial observability: the defender can only observe the states of a subset of the machines due to limited detection capabilities; (2) unknown utility functions: the defender can only receive some utility values; i.e., feedbacks, after some defense actions are taken, but does not know the utility functions; i.e., the mapping from actions and environment to the feedbacks because the locations and the impacts of zero-day vulnerabilities are not available. In fact, these two challenges are ubiquitous when we design ACD schemes against zero-day attacks. For example, the system states are the security-sensitive data structures and difficult to observe because it requires huge time and human effort to find out the security-sensitive ones among all the data structures. And the segmentation faults are determined by the defender's actions; i.e., guard page locations, the attacker's actions; i.e., over-read memory, and the environment; i.e., the allocated buffers on the heap. The defender can only access the induced utility values but does not know the utility functions which involve the allocated buffers and the attacker's actions.

We propose online algorithms to solve partially observable Markov decision process (POMDP) problems with unknown utility functions. In particular, given previous actions, induced utility values and observations, the defender forms a belief about the current state of the system and takes the empirical average of utility values to estimate of the utility functions. Then the agent performs value iterations with the estimates of the utility functions to approximate the optimal value function of each belief state at each time step. A limitation of this algorithm is high computational complexity. In our example, each update of estimates of the utility functions take hours when the action space has 2,240 actions and observation space has 280 observations. To mitigate the computational complexity, we introduce a Q-learning based algorithm which updates one Q value per iteration. The update time reduces to seconds with the same sizes. In next year, we will study the more practical scenario where the transition probabilities of POMDP are unknown and instead learn on-the-fly.

APPLYING DEEP LEARNING TO ACCURATELY DETECT ROP PAYLOADS

Return-oriented programming (ROP) is a code reuse attack that chains short snippets of existing code (known as gadgets) to perform arbitrary operations on target machines. Since ROP is Turing complete and since ROP attacks only require the target program to have some buffer overflow vulnerabilities, ROP attacks are one of the most important types of remote exploits in real world. Although ASLR, i.e., Randomization of ASL (Address Space Layout), makes a computer more resilient to ROP attacks, several kinds of ROP attacks can still successfully circumvent ASLR defense. As a result, how to detect ROP attacks has attracted a lot of attention in the research community.

Existing detection methods against ROP often rely on certain heuristic rules and/or require instrumentations to the program or the compiler. As a result, they exhibit low detection efficiency and/or have high runtime overhead. Our review of the existing techniques shows that no existing ROP detection technique simultaneously meets the following four high-desired requirements: (1) can detect most ROP attacks; (2) close to zero false positive rate; (3) low runtime overhead; (4) non-intrusive.

In this work, we developed ROPNN, a unique blend of ASL-guided disassembly and deep neural networks, to detect ROP payloads in HTTP requests, PDF files, and images, etc. This unique blend is based on a key insight that *deep learning cannot be directly applied* to detect ROP attacks in a way that simultaneously meets the four requirements mentioned above.

After a deep neural network is trained against a particular application, ROPNN is doing its intrusion detection job as follows: the disassembler treats the application's input data as code pointers to potential gadgets and finds the potential gadget chains, which are then classified by the deep neural network as benign or malicious. We propose novel methods to generate the benign and malicious training datasets, respectively, and process huge amount (TB-level) of raw input data to obtain sufficient training data.

Our experiments show that ROPNN has very high detection rate (98.3%) while maintaining very low false positive rate (0.01%). We tested ROPNN against ROP exploits that are collected in-the-wild, created manually or created by ROP exploit generation tools Ropper and ROPC. ROPNN successfully detects all of the 80 real exploits. Meanwhile, ROPNN is completely non-intrusive and does not incur any runtime overhead to the target program. In sum, ROPNN is a very practical intrusion detection technique. It could be deployed as a plugin for commercial network intrusion detection systems such as Snort and Bro.

During the next year, we will conduct in-depth analysis on (a) why the built deep neural network can accurately detect ROP payloads, (b) why the built deep neural network is better than a classifier trained by other machine learning techniques, and (c) how resilient the built deep neural network is when defending against adversarial samples.

ADAPTIVE MANAGEMENT OF SYSTEM VULNERABILITIES

Vulnerability remediation is an ongoing task in operational software and network security management. A key challenge in remediation is the dynamic prioritization of remediation efforts based on system importance, vulnerability severity, duration of the exposed vulnerability and resources available to the system administration team. We have developed and implemented an effective vulnerability management strategy, called VULCON (VULnerability CONtrol), to formulate and implement solutions to

this problem, based on periodic vulnerability scans that most organizations perform routinely already. The strategy is based on two fundamental performance metrics: (1) time-to-vulnerability remediation (TVR) and (2) total vulnerability exposure (TVE). VULCON takes as input real vulnerability scan reports, metadata about the discovered vulnerabilities, asset criticality, and personnel resources. VULCON uses a mixed-integer multi-objective optimization algorithm to prioritize vulnerabilities for patching, such that the above performance metrics are optimized subject to the given resource constraints. VULCON has been tested on multiple months of real scan data from a cyber-security operations center (CSOC). Results indicate an overall TVE reduction of 8.97% when VULCON optimizes a realistic security analyst workforce's effort. Additionally, VULCON demonstrates that it can determine monthly resources required to maintain a target TVE score. As such, VULCON provides valuable operational guidance for improving vulnerability response processes in CSOCs.

FORMULATING AND ANALYZING DYNAMIC CYBER-SECURITY GAME MODELS

The aim of our work on game modeling is to extend the complexity envelope so that realistic cyber-security scenarios can be effectively treated to game-theoretic analysis. There are several dimensions to complexity, but perhaps the most important for adaptive cyber-defense are dynamics and information. The interaction between dynamics and information is especially crucial for ACD, since it is the progression of attacker information over time, and the adaptation of defender to make that information obsolete that defines the power of ACD techniques.

Our approach to this goal has been to explore a series of domains, each illustrating a core issue in network security, and presenting basic challenges in representation and reasoning. The two domains we focused on this year both build on prior work (in this project and external) on security modeling and decision support.

The first is in the domain of botnet data exfiltration, building on prior work in this project. In this model, the strategic attacker decision is routing: what path to take in the botnet to exfiltrate data, and the strategic defender decision is sensor placement. Prior defender strategies were based on network analysis and heuristics for optimizing coverage with a fixed budget of sensors. By formulating the situation as a game, we are able to consider placement strategies that expressly respond to expected attacker routing strategies. Our main finding from analysis of the game model we constructed is that these defense strategies can be much more effective than strategies developed purely based on network structure and without taking the attacker routing strategy into account.

The second main area of investigation this year was on games defined by multi-stage attack graphs. Again, this work builds on a standard representation structure (attack graphs) that have been developed in various forms by prior researchers, originally well before this project but further refined in this project. Our work specifically builds on the formal Bayesian attack graph model developed in this project. Unlike the POMDP model, the game model assumes that both attackers and defenders strive for optimal actions on the attack graph. We developed sophisticated heuristics based on relaxed optimization problems, and found they support tractable and effective defenses. We have also begun to explore deep reinforcement learning as a method for generating strategies for this game. Preliminary results are highly promising, and we expect to spend the remaining time on this project pursuing that approach.

FUNDAMENTAL CONTROL- AND GAME-THEORETIC METHODS

The goal of this line of research is to develop fundamental control- and game theoretic techniques for solving suitably formulated cyber security problems.

Results on the control approach. We developed two methods to solve cyber-security control problems. The first is a minmax approach according to which the controller attempts to minimize the worst-case damage inflicted to the cyber network by the attacker. The main advantage of this method is its scalability. The second is a Partially Observed Markov Decision Process (POMDP)-based approach according to which the defender learns, through his observations, the attacker's strategy and attempts to minimize the average damage inflicted to the network by the attacker. The main advantage of this approach is discovery of an optimal adaptive defense strategy.

Results on the game-theoretic approach. We developed the "Sufficient Information Approach to Dynamic Games with Asymmetric Information" where the agents participating in the game possess private and common information. This approach leads to: (i) a systematic methodology for sequentially decomposing (over time) stochastic dynamic games with asymmetric information; (ii) an algorithm for determining a class of equilibria called "Sufficient Information Based-Perfect Bayesian Equilibria" (SIB-PBE). These features are the key characteristics and advantages of the approach.

7. SCIENTIFIC PROGRESS IN YEAR 6

The major goal of this effort in the past year was to develop the scientific basis and associated understanding of autonomous cyber operations and specifically how adaptive cyber defenses play a role in such operations. In particular, we sought to apply reinforcement learning and software security techniques to solve important adaptive cyber defense problems. In addition, we investigated techniques for processing high-volume streams of network and host security data with the goal of extracting graph and network properties quickly and with reasonable memory utilization for appropriately controlling the defenses' adaptation.

For the control-theoretic reasoning component, the focus was on completing the research on new representational and algorithmic methods. This research had been developed in the course of several PhD dissertations, so the emphasis in the past year was on summarizing and disseminating the key theoretical results and other technical accomplishments. For the game-theoretic reasoning component, the key goal was to demonstrate the ability to leverage the power of new deep reinforcement learning (deep RL) techniques in this context. In particular, the goal was to extend our prior modeling and analysis of cyber-security domains defined by attack graphs, by integrating deep RL with game-theoretic search methods.

Finally, an overarching goal was that of finalizing the production of a book summarizing all the results achieved during the entire project.

A detailed analysis of some fundamental results about adaptive autonomy in adversarial engagements was performed, based on seminal disparate results in game theory and machine learning spanning 20 years of research. It was concluded that autonomous cyber operations in which both attackers and defenders are adapting during operation can lead to unstable, nonconvergent and suboptimal dynamics

within the cyber systems. Ideas for dominating in such an environment were identified and documented. In addition, new algorithms and analyses were discovered for computing approximate solutions to network centrality problems based on streaming data.

We completed all control theory threads of the project, and published or drafted papers describing the project's contributions. We also demonstrated the deep RL integration with empirical game-theoretic analysis (EGTA). Our new EGTA+RL algorithm produced strikingly superior results for the attack-graph game scenario. The algorithm also incorporated novel approaches to combinatorial action spaces and strategy exploration, as documented in a paper presented at the 20th ACM Conference on Economics and Computations (EC 2019).

GAME-THEORETIC REINFORCEMENT LEARNING

We have also solved the convergence rate problem with game theoretic reinforcement learning when this machine learning technique is used to build better dynamic platforms (i.e., dynamically mutating the operating system of a server targeted by the attacker). In particular, our accomplishment is summarized as follows:

- Each operating system has a set of specific vulnerabilities. To prevent the attacker from having sufficient knowledge about these specific vulnerabilities, dynamic platforms – a moving target defense (MTD) technique – seek to dynamically mutate the operating system of a server targeted by the attacker. To making this MTD technique more effective, game theoretic reinforcement learning can be leveraged to make better mutating decisions. However, the convergence rate problem with the corresponding class of multi-player discrete games had not yet been solved in the literature.
- Our Automatica 2019 paper investigates this class of multi-player discrete games where each player aims to maximize its own utility function. Each player does not know the other players' action sets, their deployed actions or the structures of its own or the others' utility functions. Instead, each player only knows its own deployed actions and its received utility values in recent history. We proposed a reinforcement learning algorithm which converges to the set of action profiles which have maximal stochastic potential with probability one. Furthermore, an upper bound on the convergence rate is derived and is minimized when the exploration rates are restricted to p-series. To validate the algorithm, we conducted a case study in dynamic platforms.

REPACKAGING ATTACKS AGAINST MOBILE APPS

Additionally, we have developed a novel technique to defend from repackaging attacks against mobile apps. In particular, this accomplishment is summarized as follows:

- Artificial diversity through code obfuscation is a main MTD technique. Our IEEE Transactions on Mobile Computing 2018 paper proposed a novel code obfuscation method for defending repackaging attacks against millions of mobile apps. Our method is the first MTD technique that leverages a unique information asymmetry between developers and attackers, and lets a repackaged app automatically expose itself.
- By simply adding malicious code or advertisements in legitimate smartphone apps, attackers could benefit a lot from repackaging. The existing license protection mechanisms can be easily subverted by repackaged apps. A major defense is to detect. However, detection requires finding at least two “similar” apps simultaneously. We propose a self-defending approach: let a

repackaged app automatically expose itself. However, it is very challenging to achieve this goal. If developers and smartphones/users do not share any secret, attackers' app repackaging studio would be able to do whatever legitimate smartphones/users are able to do. We find that there exists a unique information asymmetry between developers and attackers. Leveraging this asymmetry, our new self-defending code (SDC) approach encrypts parts of the app code at compile time and dynamically decrypts the ciphertext code at run-time. Different from previous work, the key is derived from both the information asymmetry and the app's checksum. Once the app is repackaged, the changed checksum will let the app run abnormally, further exposing the repackaging. The information asymmetry protects the key from being attacked. We build a smartphone anti-repackaging system prototype. To the best of our knowledge, this is the first work that lets repackaged apps automatically malfunction while having none effect on a benign app's function.

DEFENDING AGAINST ZERO-DAY BUFFER OVER-READ ATTACKS

We have also addressed a main design flaw in our on-going work on defending against zero-day buffer over-read attacks. In particular, this accomplishment is summarized as follows:

- Buffer over-read attacks are unauthorized reads going beyond the boundaries of the vulnerable buffers. Many buffer over-read attacks are zero-day attacks where the locations of vulnerable buffers are usually unknown to security analysts. Zero-day buffer over-read attacks can produce severe damages to servers. 24-55% of HTTPS servers in the Alexa Top 1 Million were initially vulnerable to Heartbleed.
- Our 2019 submission to IEEE Transactions on Computers has solved an open problem in protecting a web server against Heartbleed and other zero-day buffer over-read attacks: Is there any defense which can bound the security damage (e.g., the amount of data disclosed to the attacker)?
- Our solution consists of (a) a novel adaptive cyber defense scheme which combines a new multi-arm bandit algorithm and a simple modification of the memory allocator; (b) a theorem which mathematically proves that the security damage is bounded by the proposed defense scheme; and (c) a way to mathematically predict the bound. We have designed and implemented a system prototype of the proposed defense scheme, and the experiment results show that the measured security damage is always below the mathematically predicted bound.

CYBER SECURITY OPERATION CENTERS (CSOCs)

Large organizations including Army Research Laboratory (ARL) operate Cyber Security Operation Centers (CSOCs), which form the first line of cyber-defense. This year we produced several papers related to effective management of CSOCs under adverse conditions. Two major results are as follows:

- The inspection of cyber-alerts is a critical part of CSOC operations (defender or blue team). In previous year, we proposed a reinforcement learning (RL) based approach for the defender's decision making to prevent the cyber-alert queue length from growing large and overwhelming the defender. In this year, we performed a red team (*adversarial*) evaluation of our approach. With the recent attacks on learning-based decision making systems, it is even more important to test the limits of the defender's RL approach. In this work, we learn several adversarial alert generation policies and the best response against them for various defender's inspection policy. Surprisingly, we find the defender's policies to be quite robust to the best response of the attacker. In order to explain this observation, we extend the earlier defender's RL model to a game model with adversarial RL, and show that there exist defender policies that can be robust against

any adversarial policy. We also derive a competitive baseline from the game theory model and compare it to the defender's RL approach. However, when we go further to exploit the assumptions made in the Markov Decision Process (MDP) in the defender's RL model, we discover an attacker policy that overwhelms the defender. We use a double oracle like approach to retrain the defender with episodes from this discovered attacker policy. This made the defender robust to the discovered attacker policy and no further harmful attacker policies were discovered. Overall, the adversarial RL and double oracle approach in RL are general techniques that are applicable to other RL usage in adversarial environments.

- Under normal operating conditions, a CSOC's analyst staffing level (defines the service rate) is designed for a certain level of alert arrival rate. However, alerts may arrive at a much higher rate than the service rate at the CSOC due to several adverse events. Such adverse events include an external attack, lack of sufficient staffing due to absenteeism, or other network events such as a temporary breakdown in communication link between IDSs and the CSOC. The consequence of high alert generation rate above the service rate leads to a buildup in alert backlog, which in turn reduces the Level of Operational Effectiveness (LOE) of the CSOC. Though CSOC managers have a few action options, with some being ad hoc, to maintain the LOE during the shift, more often than not, the LOE reaches a critical low due to high alert backlogs. At this point, the only option is to discard alerts in order to restore the LOE of the CSOC. In this work, we developed an intelligent strategy that uses reinforcement learning to make decisions on when and how many alerts to discard during the shift operation. Our experiments show that the RL-based intelligent strategy discards the lowest number of alerts over several experimental simulation runs, determined the right timing to discard the alerts, and maintained the AvgTTA below the threshold set by the CSOC. It is also concluded that at the time of several adverse events and with limited analyst resource, a CSOC manager can use the RL-based strategy as a decision-support system that guarantees the AvgTTA below the threshold value in every hour of shift operation. Such a guarantee with minimum number of alerts discarded is a paradigm shift in how CSOCs with limited resources (both regular and on-call analysts combined) can efficiently operate in the face of uncertainties that affect the length of the alert backlog for investigation. Several meta-principles provide insights into the process of decision-making that is involved in discarding alerts, which could guide CSOC managers in maintaining the LOE under adverse events.

EDITED BOOK

Finally, we produced and edited a book summarizing all the results achieved during the project, which was recently published by Springer. The book is titled ***Adversarial and Uncertain Reasoning for Adaptive Cyber Defense: Control- and Game-Theoretic Approaches to Cyber Security*** and was edited by Sushil Jajodia, George Cybenko, Peng Liu, Cliff Wang, and Michael Wellman.

8. SCIENTIFIC PROGRESS IN YEAR 7

Building on previous results in machine learning, adversarial modeling, and moving target defense quantification, we have developed adaptive cyber defenses based on reinforcement learning and deep learning approaches that push the boundary of existing autonomous cyber defense capabilities. This effort is paving the way for the development of next-generation autonomous cyber operations.

Specifically, we developed: (i) regret-based Q-Learning reinforcement learning algorithms for spatially distributed multi-agent systems where both parties in an adversarial engagement are adapting; (ii) the first ACD scheme to defend against multi-stage attacks using a reinforcement-learning-based POMDP (partially observable Markov decision process); (iii) the first deep-learning-based technique to detect

previously-unseen ROP payloads without using any heuristics; and (iv) the first reinforcement-learning-based ACD technique to ensure quantitatively bounded damages (from cyber-attacks) in effectively defending against Heartbleed-like attacks.

We also made progress on (i) modeling how stealthy attackers establish and expand their foothold within target systems, and (ii) ensuring a predetermined level of operational effectiveness (LOE) in Cyber Security Operations Centers (CSOCs) under all adverse conditions.

Details about achievements in each area are provided in the subsections below.

MODELING STEALTHY ADVERSARIES

In recent years, threat actors have increasingly prioritized stealth as a critical feature to ensure persistence within a target system. An analysis of the Stuxnet malware, purportedly deployed to sabotage centrifuges required for the Iranian nuclear program, showed that as much as 50% of the development cost was devoted to stealth-related features. Furthermore, while Stuxnet possessed the capability to catastrophically damage all affected centrifuges at any given time, the threat actors opted to subtly degrade the centrifuges in a manner which would not be immediately evident. This allowed the malware to delay the overall progress of the Iranian nuclear program to an extent greater than a conventional military strike.

While Stuxnet provides an early case study, many other advanced stealthy threat actors are the subject of ongoing research. Earlier detection of malicious activity has tangible economic benefits for an organization. The Ponemon Institute conducts yearly studies on the cost of organizational data breaches. In 2017, it concluded that organizations which required over 100 days to identify that a breach had occurred faced 36% higher costs compared to organizations which identified breaches in under 100 days. The delay in detecting malicious activity can result in adversaries having the opportunity to gain control of a larger portion of the network and exfiltrate larger amounts of sensitive information – which may include critical trade secrets or closely-guarded strategic data. Clearly, mitigating the threats posed by these stealthy attackers has incredible value for any organization.

While intrusion detection is a vast area of research, relatively little work has been done in examining how stealthy attackers establish their view of the network over time, especially with respect to the factors that may incentivize or deter them from selecting specific targets in the network. In our previous work, we developed a model to study how advanced, omniscient attackers may establish, expand, and maintain their footprint within a target system. By comparison, the work reported here relaxes the assumption that attackers have complete knowledge of the target network prior to start compromising nodes within that network, and more realistically assumes that such knowledge is acquired incrementally, as the attacker further penetrates the network.

With our work, we aim at developing a deeper understanding of how a stealthy attacker operates, which in turn can guide the development of adequate autonomous defenses, including approaches based on moving target defense and cyber deception. In summary, the main contributions of this work are (i) a framework to model reconnaissance by stealthy attackers, and (ii) a practical approach to define cost and reward metrics used by adversaries to drive the selection of their targets when moving laterally within the network.

MAINTAINING THE LEVEL OF OPERATIONAL EFFECTIVENESS OF A CSOC UNDER ALL ADVERSE CONDITIONS

There are several disruptive factors that can have adverse impact on cyber security operations centers (CSOCs). These factors can be external or internal. Some examples of the external events include a zero-day attack, sensor downtime, broken communication link between the IDS and the CSOC, and increased attacker activity on a network that is monitored by the sensors. The internal events include equipment shutdowns at a CSOC, power disruptions, analyst absenteeism, required analyst credentials not available in a shift such as secret clearances, and analyst scheduling issues that does not provide the required expertise mix for alert analysis.

To respond to these disruptive events, several defensive solutions have been proposed: Bring in on-call analysts, offload some of the alerts to partner site(s), keep raising the threshold for the priority level of alerts so fewer and fewer alerts are added to the queue for analysis, and others. *Unfortunately, for each of these defenses, one can construct scenarios under which CSOC will be overwhelmed.*

The objective of this work is to propose an approach that can maintain the predetermined level of operational effectiveness (LOE) under **ALL** adverse conditions. Our approach is based on discarding alerts after all action options have been exhausted by the CSOC manager. We present the design and evaluation of a novel reinforcement learning (RL)-based intelligent strategy for discarding alerts. We compared a total of four strategies for discarding alerts, and results indicate that, overall, the RL strategy is the best performer among all strategies. We measure the following desiderata for comparing the strategies: 1) minimizing the number of alerts discarded, 2) ensuring highest utilization of analysts, 3) determining the optimal time at which the alerts must be discarded in a shift, and 4) in maintaining the best possible LOE closest to the baseline target LOE. Several meta-principles provide insights into the process of decision-making that is involved in discarding alerts, which could guide CSOC managers in maintaining the LOE under adverse events.

This is a joint work with Hasan Cam of Army Research Laboratory and a paper summarizing the results have been submitted to the ACM Transactions on Intelligent Systems and Technology (TIST).

DEFENDING AGAINST MULTI-STAGE ATTACKS USING REINFORCEMENT LEARNING-BASED POMDP

Growing multi-stage attacks in computer networks impose significant security risks and necessitate the development of effective defense schemes which are able to autonomously respond to intrusions during the vulnerability window (i.e., the time window in which the corresponding patch is being developed). However, the defender faces several real-world challenges; e.g., unknown state transition probabilities and unknown impacts of successful exploits.

In this work, we leverage reinforcement learning to develop an innovative adaptive cyber defense to maximize the cost-effectiveness subject to the aforementioned challenges. In particular, we use Bayesian attack graphs to model the interactions between the attacker and networks. Then we formulate the defense problem of interest as a partially observable Markov decision process (POMDP) problem where the defender maintains belief states to estimate system states, leverages Thompson sampling to estimate transition probabilities and utilizes reinforcement learning to choose optimal defense actions using measured utility values.

The algorithm performance is evaluated via numerical simulations based on real-world attacks. The results show that when compared to the baseline value-iteration-based algorithm, our reinforcement

learning algorithm can reduce the computation complexity significantly while achieving similar aggregate utilities. The results also show that the transition probability estimation error of the new algorithm exponentially decreases and reduces by around 9.52% after 100 episodes; i.e., 10,000 time steps, while that of the baseline “solely history” policy decreases at the beginning then keeps increasing as simulation goes on.

This work is the first one that systematically addresses the two above-mentioned **real-world challenges**, namely unknown state transition probabilities and unknown impacts of successful exploits. In contrast, the existing POMDP-based defense schemes (against multi-stage attacks) either assume that the state transition probabilities are given or assume that the impacts of successful exploits are known.

DEEP LEARNING-BASED TECHNIQUE TO DETECT PREVIOUSLY UNSEEN ROP PAYLOADS WITHOUT USING ANY HEURISTICS

Return-oriented programming (ROP) is a most important code reuse attack that chains short snippets of existing code to perform arbitrary operations on target machines. Existing detection methods against ROP exhibit unsatisfactory detection accuracy and/or have high runtime overhead. It seems that all the most practical existing ROP detection methods rely on heuristics provided by subject-matter experts (SME).

In this work, we developed DEEPRETURN which is to our best knowledge the first deep-learning-based technique to detect previously-unseen ROP payloads. The significance of this new technique is indicated not only by that it achieves very high accuracy, but also by that it does not make use of any heuristics. DEEPRETURN can potentially solve the “soliciting knowledge from SME” challenge, which is one of the biggest challenges in defending against code reuse attacks.

DEEPRETURN innovatively combines address space layout guided disassembly and deep neural networks to detect ROP payloads. The disassembler treats application input data as code pointers and aims to find any potential gadget chains, which are then classified by a deep neural network as benign or malicious. Our experiments show that DEEPRETURN has high detection rate (99.3%) and a very low false positive rate (0.01%). DEEPRETURN successfully detects all of the 100 real-world ROP exploits that are collected in-the-wild, created manually or created by ROP exploit generation tools. DEEPRETURN is non-intrusive and does not incur any runtime overhead to the protected program.

REINFORCEMENT LEARNING-BASED ACD TECHNIQUE TO ENSURE QUANTITATIVELY BOUNDED DAMAGES IN DEFENDING HEARTBLEED-LIKE ATTACKS

Heartbleed-like attacks are one of the most serious attacks against mission-critical servers. In recent years, several practical defense techniques have been proposed to defend such attacks; however, no existing technique can ensure **quantitatively bounded damages** (from the attacks). As a result, the existing techniques sometimes could do a great job with minimal damage, but sometimes they could do a very poor job with unacceptable amount of damages.

In this work, we have developed a reinforcement-learning-based ACD technique to ensure quantitatively bounded damages in defending Heartbleed-like attacks. To the best of our knowledge, this new technique is the first one that can ensure bounded damages in defending Heartbleed.

This new technique creatively combines customized memory allocation and reinforcement-learning. It can defend against a class of zero-day buffer over-read attacks (including Heartbleed) that follow unknown stationary probability distributions. In particular, first, we have developed an improved Upper

Confidence Bound (UCB) algorithm, which is a widely-used algorithm for fulfilling a simplified reinforcement learning task (i.e., solving a stochastic multi-armed bandit problem). The improved UCB algorithm adaptively allocates guard pages on a heap based on induced damage of the guard pages so as to minimize the accumulated damage over time. Second, an efficient server modification is introduced to randomly allocate buffers.

In terms of ensuring bounded damages, the damages associated with the improved UCB algorithm are proven to be always below a temporal bound without knowing which attack is launched when the buffer allocation follows a certain stationary probability distribution. Moreover, the damages of our scheme asymptotically converge to those of the optimal defense policy where the launched attacks and their distributions are known in advance.

The new defense scheme is evaluated with several real-world Heartbleed attacks. The experiment results demonstrate the validity of the upper bound and show that the adaptive defense is effective against all the attacks of interest with runtime overheads as low as 5%.

9. SCIENTIFIC PROGRESS IN YEAR 8

During this reporting period, we primarily focused on (i) studying the problem of on-line learning to adapt a deployed moving target cyber defense during autonomous operation, (ii) applying machine learning techniques, including deep learning and transfer learning techniques, to develop novel ACD techniques, (iii) improving our game-theoretic approaches.

The work on on-line learning and adaptation has led to the definition on a new area of research, which we called *Attritable Machine Learning*. Using machine learning, we have successfully developed three novel ACD techniques. First, we have developed a new anomaly detection methodology based on deep learning. Second, we have developed the first deep-learning-based technique to identify silent buffer overflow vulnerabilities in third-party program binaries without source code. Third, we have developed a novel domain adaptation technique to address the imbalanced data issue in applying deep learning to detect previously unseen ROP payloads.

Details about achievements in each area are provided in the subsections below.

ATTRITABLE MACHINE LEARNING

We have been studying the problem of on-line learning to adapt a deployed moving target cyber defense during autonomous operation. This problem has several challenges including:

- a. On-line learning and adaptation require trying possibly suboptimal defensive configurations to fully explore the available defensive actions.
- b. Using such suboptimal defenses will increase the compromise rate of blue force defensive assets; Blue force adaptation and underlying mission success have to be sustained in spite of those compromised assets.
- c. Mission success in spite of compromised assets requires redundant, diverse subsystems/agents.
- d. The ability to continue mission operations and adapt successfully in the presence of compromised and/or inoperative systems.

This has led us to introduce a new research area which we call *Attritable Machine Learning*. Approaches to Attritable Machine learning must address the above issues and provide a mechanism for continued assured operation even if some compute and information assets are no longer trustworthy.

Accordingly, we have been investigating the use of blockchain-like technology in concert with Byzantine agreement methods to allow attritable systems to continue their missions while adapting through learning.

GAME-THEORETICAL ANALYSIS

We made significant progress on the methodology of combining deep RL integration with empirical game-theoretic analysis (EGTA). Specifically, we extended our EGTA+RL algorithm, which we had already found to produce strikingly superior results for the attack-graph game scenario. The important accomplishments in extending this work have been:

- Improved techniques for dealing with combinatorial action spaces, which are an important feature of attack-graph games and other cyber-security situations (paper to be submitted).
- Q-mixing method for constructing a policy against any probability distribution of opponent policies (arXiv:2009.14180).
- Use of Q-mixing in EGTA+RL for dramatically improved performance (published as spotlight in ICLR-21).
- Principles for analyzing strategy exploration methods (arXiv:2105.10423, under review for conference publication).
- Regularization methods for strategy exploration (under review for conference publication).

We showed that we could integrate deep RL with game-theoretic search methods to achieve new levels of performance on cyber-security domains defined by attack graphs. The objective in this year's add-on research was to consolidate those gains by leveraging and extending these powerful methods.

NOVEL ANOMALY DETECTION APPROACH FOR DISCRETE EVENT LOGS

Anomaly detection is a fundamentally important aspect of adversarial and uncertain reasoning for adaptive cyber defense. Anomaly detection is also one of the most challenging problems in the field of intrusion detection for discrete event logs. While most prior work focused on applying unsupervised learning upon engineered features, most recent work has started to resolve this challenge by applying deep learning methodology to abstraction of discrete event entries. Inspired by natural language processing, LSTM-based anomaly detection models were proposed. They try to predict upcoming events, and raise an anomaly alert when a prediction fails to meet a certain criterion. However, such a predict-next-event methodology has a fundamental limitation: event predictions may not be able to fully exploit the distinctive characteristics of sequences. This limitation leads to high false positives (FPs).

In this work, we have developed a novel deep-learning-based anomaly detection approach for discrete event logs. Based on the **key observation** that it is also critical to examine the structure of sequences and the **bi-directional causality** among individual events, we propose a new methodology: Recomposing event sequences as anomaly detection. We propose **DabLog**, a LSTM-based Deep Autoencoder-Based anomaly detection method for discrete event Logs. The fundamental difference is that, rather than predicting upcoming events, our approach determines whether a sequence is normal or abnormal by analyzing (encoding) and reconstructing (decoding) the given sequence.

To the best of our knowledge, DabLog is the first work that detects anomalies through Deep-Autoencoder-based event sequence recomposing. This work provides new understanding about why bi-directional causality could be very important in detecting abnormal events in enterprise networks.

Our evaluation results show that our new methodology can significantly reduce the numbers of FPs, hence achieving a higher F1 score. For example, when being compared to a representative existing deep-learning-based anomaly detection method (denoted Baseline), DabLog has **2,419 less** false positives, in exchange of 83 less true positives when detecting abnormal events in HDFS (Hadoop file system) systems. In addition, DabLogs reaches a F1 score of 80.25%, whereas Baseline can only reach 56.77%.

DEEP-LEARNING-BASED TECHNIQUE TO IDENTIFY SILENT BUFFER OVERFLOW

VULNERABILITIES IN PROGRAM BINARIES WITHOUT SOURCE CODE

A software vulnerability could be exploited without any visible symptoms. When no source code is available, although such silent program executions could cause very serious damage, analyzing silent yet harmful exeutions is still an open problem. When source code is unavailable, the existing techniques and tools (including TaintCheck, Memcheck, Fuzzing, and Symbolic Execution) cannot identify silent stack BOF (buffer overflow vulnerabilities); cannot identify global BOF, and cannot identify heap BOF (except that Memcheck can partially cope with silent heap BOF).

In this work, we propose a graph neural network (GNN) assisted data flow analysis method for spotting silent buffer overflows in execution traces. The new method combines a novel graph structure (denoted DFG+) beyond data-flow graphs, a tool to extract DFG+ from execution traces, and a modified Relational Graph Convolutional Network as the GNN model to be trained. The evaluation results show that a well-trained model can be used to analyze vulnerabilities in execution traces (of previously-unseen programs) without support of any source code.

Leveraging deep learning, the proposed method is, to our best knowledge, the first general-purpose analysis method for identifying silent buffer overflows. It is also the first method to spot silent buffer overflows in global variables, stack variables, or heap variables without crossing the boundary of allocated chunks. It seems also the first work that conducts deep-neural-network-assisted data flow analysis for cybersecurity purposes.

Our model achieves 94.39% accuracy on the test data, and successfully locates 29 out of 30 real-world silent buffer overflow vulnerabilities. The paper is currently under review by ACM CCS 2021.

TACKLING IMBALANCED DATA IN CYBERSECURITY WITH TRANSFER LEARNING: A CASE WITH ROP PAYLOAD DETECTION

In recent years, deep learning gained proliferating popularity in the cybersecurity application domain, since when being compared to traditional machine learning, it usually involves less human effort, produces better results, and provides better generalizability. However, the **imbalanced data issue** is very common in cybersecurity, which can substantially deteriorate the performance of the deep learning models.

In this work, we have developed a transfer-learning-based method to tackle the imbalanced data issue in cybersecurity using Return-Oriented Programming (ROP) payload detection as a case study. Return-oriented programming (ROP) is a most important code reuse attack that chains short snippets of existing code to perform arbitrary operations on target machines. Existing detection methods against ROP exhibit unsatisfactory detection accuracy and/or have high runtime overhead. It seems that all the most practical existing ROP detection methods rely on heuristics provided by subject-matter experts (SME). During the last reporting period of time, we have developed DEEPRETURN, to our best knowledge the first deep-learning-based technique to detect previously unseen ROP payloads.

In this work, we apply **domain adaptation**, a widely-considered transfer learning technique, to address the imbalanced data issue in training DEEPRETURN neural networks for detecting ROP payloads. A main limitation of the DEEPRETURN technique is that it is extremely costly to generate benign training data samples, although it is straightforward to generate malicious data samples. To address this daunting challenge, we apply the mean-maximum-discrepancy-based domain adaptation method in a way specific to DEEPRETURN neural networks. As a result, the deep neural network designated for the target domain (i.e., a different program vulnerable to ROP attack) will be trained using not only the **imbalanced target** domain data, but also the source domain data which are actually balanced data. Through domain adaptation, useful **knowledge** (about how to detect ROP payloads) is **transferred** from the source domain to the target domain, and the accuracy of the target domain neural network model can be substantially improved when compared to independent model training in the target domain with imbalanced data.

The evaluation results show that our approach achieved 3.3% average false positive rate, 97.18% average F1 score and 94.18% average detection rate on 3 different target domain programs using 2 different source domain programs, with 0 benign training data samples in the target domain. The performance improvement compared to the baseline is a trade-off between false positive rate and detection rate. Using our approach, **the number of false positives is reduced by 23.20%**, and as a trade-off, the number of detected malicious samples is reduced by 0.50%.

10. TECHNOLOGY TRANSFER

- Technical interchange meetings with ARL, AFRL/RISB, MITRE, and two companies (one small and one large business)
- Interactions with the Army Research Lab
 - Presentation at ARL, Adelphi on July 12, 2016
 - Presentation at ARL, Adelphi on July 25, 2017
 - Several visit by ARL researchers to GMU
 - Gave the CSOC software to ARL
- Interaction with ACS, Northrop Grumman, and General Dynamics Information Technology
- Interactions with the Army CRA on Cyber Security:
 - Presented a moving target survey at the CRA Bootcamp at UC Davis in July 2017
 - Participated in the Cyber CRA annual review, August 2017
- Presented MURI concepts at Lawrence Livermore national Lab in August 2017
- We gave seminars at ARL in July 2017 and August 2017.
- We delivered the CSOC scheduling software to ARL.
- We published several joint papers with ARL researchers Hasan Cam and Robert Erbacher.
- Several researchers – Hasan Cam, Jennifer Cowley, Stephen Hutchinson, Ken Yu, and Travis Parker – visited George Mason University.
- We interacted with the members of the ARL ITA Project:
 - IBM TJ Watson (Liang Ma) and ARL (Hasan Cam, Michelle Vani, Sue Kase)
- We interacted with the ARL CRA on Cyber Security Project:

- We presented a moving target survey at the CRA Bootcamp at the University of California Davis in July 2017.
- We participated in the Cyber CRA annual review in August 2017
- The project team was involved in a number of activities and outreach initiative that amplified the impact of the research effort well beyond the four participating institutions.
- During Year 6, we also continued our close collaboration with the researchers at the Army Research Laboratory. Drs. Jajodia and Ganesan worked closely with ARL researchers Drs. Hasan Cam, Jennifer Cowley, Stephen Hutchinson. This resulted in several several papers jointly coauthored with Dr Cam. Another paper coauthored with Drs. Cam, Cowley, and Hutchinson is in the works.
- Drs. Peng Liu and Minghui Zhu visited the Army Research Lab and gave two presentations on “reinforcement learning based adaptive cyber defenses” and “applying deep learning to defend against zero-day cyber-attacks,” respectively.

11. BROADER IMPACT

The project team was involved in a number of activities and outreach initiative that amplified the impact of the research effort well beyond the four participating institutions.

- We founded and closely involved in the organization of the ACM Workshops on Moving Target Defense, which was co-located with ACM CCS. Sushil Jajodia is a member of the Steering Committee.
- We were involved in the organization of The 5th International Workshop on Graphical Models for Security (GraMSec 2018). Sushil Jajodia is a member of the Steering Committee, and George Cybenko served as a program co-chair.
- We organized an ARO Invitational Workshop on Foundations of Autonomous Adaptive Cyber Systems at George Mason University, February 27-28, 2018. An edited book was published as an output of the workshop.
- We engaged is significant interaction and technical exchanges with Vencore Labs (now Perspecta Labs), Northrop Grumman, and General Dynamics Information Technology.
- Presented results of the MURI effort at Lawrence Livermore National Lab in August 2017.
- We gave presentations and invited talks at several international conferences and various other events (see 2.(c).f for details).

12. STUDENT/SUPPORTED PERSONNEL

GRADUATE STUDENTS

1. Ankit Shah, George Mason University
2. Keith Pham, George Mason University
3. Sridhar Venkatesan, George Mason University
4. Giridhar Kaushik, George Mason University
5. Prakruthi Karuna, George Mason University
6. Jooyeon Lee, George Mason University
7. Jing Liu, George Mason University

8. Warren Connell, George Mason University
9. Fabio De Gaspari, George Mason University 9%
10. Gabriel Stocco, Dartmouth College
11. Patrick Sweeney, Dartmouth College
12. Roger Hallman, Dartmouth College
13. Kate Farris, Dartmouth College
14. Benjamin Priest, Dartmouth College
15. Zhisheng Hu, Penn State University
16. Jun Xu, Penn State University
17. Feng Xiao, Penn State University
18. Pinyao Guo , Penn State University
19. Chuangang Ren, Penn State University
20. Wenhui Hu, Penn State University
21. Wenhui Zhang, Penn State University, 100%
22. Jun Wang, Penn State University
23. Xusheng Li, Penn State University, 50%
24. Mohammad Rasouli, University of Michigan
25. Achintya Prakash, University of Michigan
26. Erik Miehl, University of Michigan
27. Mason Wright, University of Michigan, 50% GSRA

TOTAL # 27

POST DOCTORATES

1. Kareem Amin, University of Michigan
2. Thanh Nguyen, University of Michigan
3. Ping Chen, Penn State University
4. Chen Cao, Penn State University
5. Luigi Mancini, George Mason University
6. Arunesh Sinha, University of Michigan
7. Le Guan, Penn State University
8. Chen Cao, Penn State University

TOTAL # 8

FACULTY

1. Sushil Jajodia, George Mason University
2. Massimiliano Albanese, George Mason University
3. Kun Sun, George Mason University
4. Noseong Park, George Mason University, 1 mo.
5. Hemant Purohit, George Mason University, .5 mo.
6. George Cybenko, Dartmouth College

7. Peng Liu, Penn State University
8. Minghui Zhu, Penn State University
9. Satinder Singh, University of Michigan
10. Michael Wellman, University of Michigan
11. Demosthenis Teneketzis, University of Michigan

TOTAL #: 11

UNDERGRADUATE STUDENTS

1. Feynman Liang, Dartmouth College
2. John Sullivan, Dartmouth College
3. Ruihao Zhao, University of Michigan (unfunded)
4. Zhenyang Gong, University of Michigan
5. Alejandro Cuevas Penn State University
6. Samuel Philip, unfunded summer intern
7. Alejandro Cuevas, Penn State University
8. John Sullivan, Dartmouth, research intern

TOTAL #: 9

UNDERGRADUATE DEGREES AWARDED

Number who graduated during this period: 2

Number who graduated during this period with a degree in science, mathematics, engineering, or technology fields: 2

TOTAL #: 2

MASTERS DEGREES AWARDED

1. Mohammad Rasouli, University of Michigan
2. Pinyao Guo, Penn State University
3. Maram Jibreel, George Mason University

TOTAL #: 3

PH.D.S AWARDED

1. Sridhar Venkatesan, George Mason University
2. Prakruthi Karuna, George Mason University
3. Ankit Shah, George Mason University
4. Keith Pham, George Mason University
5. Warren Connell, George Mason University
6. Gabriel Stocco, Dartmouth College
7. Patrick Sweeney, Dartmouth College
8. Kate Farris, Dartmouth College

9. Mingyi Zhao, Penn State University
10. Ben Priest, Dartmouth College
11. Erik Miehl, University of Michigan
12. Mohammad Rasouli, University of Michigan
13. Mason Wright, University of Michigan
14. Hamidreza Tavafooghi, University of Michigan
15. Jun Xu, Penn State University
16. Lannan Luo, Penn State University
17. Pinyao Guo, Penn State University
18. Zhisheng Hu, Penn State University

TOTAL #: 18

13. HONORS AND AWARDS

- Michael Wellman received the 2014 SIGAI Autonomous Agents Research Award.
- Massimiliano Albanese received the 2014 Mason Emerging Researcher/Scholar/Creator Award.
- Peng Liu received the 2015 Penn State College of Information Science and Technology Senior Faculty Excellence in Research Award.
- George Cybenko serves on the US Army Cyber Institute Advisory Board.
- George Cybenko serves on the US Air Force Scientific Advisory Board.
- **Peng Liu** received a *Best Paper Runner-up Award* at IEEE CNS 2016 for the paper:
Xiaoyan Sun, Jun Dai, Peng Liu, Anoop Singhal, and John Yen, ***"Towards Probabilistic Identification of Zero-day Attack Paths,"*** Proceedings of IEEE CNS, October 2016.
- **Mohammad Rasouli** achieved 1st Place for "Scalable Control of Cyber-Physical Systems" at the University of Michigan 2016 Engineering Graduate Symposium, Systems and Communication Engineering session, November 2016.
- **George Cybenko** received the 2016 SPIE Eric A. Leifeld Award. The award recognizes outstanding contributions to global homeland security commemorating the victims of the 9/11 attacks on America, as a reminder and inspiration to all to apply technology to better secure our homelands.
- **George Cybenko** received the US Air Force Commander's Public Service Award.
- **Sushil Jajodia** received the 2016 Federal Information Systems Security Educators' Association (FISSEA) Educator of the Year Award
- **Sushil Jajodia** received the 2016 IEEE Computer Society Technical Achievement Award
- **Best paper Award.** L.H. Pham, M. Albanese, and B.W. Priest, "A Quantitative Framework to Model Advanced Persistent Threats," in Proceedings of the 15th International Conference on Security and Cryptography (SECRYPT 2018), pages 282-293, SciTePress, 2018.
- **Best Paper Award,** Le Guan, Shijie Jia, Bo Chen, Fengwei Zhang, Bo Luo, Jingqiang Lin, Peng Liu, Xinyu Xing and Luning Xia, "Supporting Transparent Snapshot for Bare-metal Malware Analysis on Mobile Devices," in Proceedings of the 33rd Annual Conference on Computer Security Applications (ACSAC 2017), 2017.

- **Penn State Alumni Association Dissertation Award and Distinguished Doctoral Scholar Medal**, Jun Xu, Penn State University, 2017-2018 .
- **Raymond G. Tronzo, M.D. Professorship in Cybersecurity**, Peng Liu, Penn State University, 2018.
- Minghui Zhu received an NSF CAREER award.
- Minghui Zhu was promoted to Associate Professor with tenure.
- Undergraduate student Alejandro Cuevas was offered full scholarship by the CMU CyLab PhD Program.
- Best Paper Award, DBSEC 2019
- George Cybenko was named Dartmouth's first SIAM Fellow
- Sushil Jajodia was recipient of the 2020 IEEE Computer Society W. Wallace McDowell Award
- DSN 10-Year Test-of-Time Award. The following paper, co-authored by Peng Liu, has received the 10-Year Test-of-Time Award from the IEEE/IFIP International Conference on Dependable Systems and Networks (DSN) in summer 2020:
- Peng Xie, Jason H Li, Xinming Ou, Peng Liu and Renato Levy, "Using Bayesian Networks for Cyber Security Analysis," In Proceedings of IEEE DSN-DCCS 2010, June 2010, Chicago, USA

14. PUBLICATIONS

PAPERS PUBLISHED IN PEER-REVIEWED JOURNALS

1. **DOI:** <http://dx.doi.org/10.22215/timreview/712>
Title: Quantitative Metrics and Risk Assessment: The Three Tenets Model of Cybersecurity
Authors: J. Hughes and G. Cybenko
Publication Title: *Technology Innovation Management Review*
Month: August
Year: 2013

2. **DOI:**
Title: Pareto-optimal Adversarial Defense of Enterprise Systems
Authors: E. Serra, S. Jajodia, A. Pugliese, A. Rullo, and V. S. Subrahmanian
Publication Title: ACM Transactions on Information and System Security, vol. 17, no. 3
Month: March
Year: 2015

3. **DOI:**
Title: Optimal Scheduling of Cybersecurity Analysts for Minimizing risk
Authors: R. Ganesan, S. Jajodia, and H. Cam
Publication Title: ACM Transactions on Intelligent Systems and Technology, vol. 8, no. 4
Month:
Year: 2017

4. **DOI:** <https://doi.org/10.1145/2882969>
Title: Dynamic Scheduling of Cybersecurity Analysts for Minimizing Risk using Reinforcement Learning
Authors: Rajesh Ganesan, Sushil Jajodia, Ankit Shah, and Hasan Cam
Publication Title: ACM Transactions on Intelligent Systems and Technology, vol. 8, no. 1
Month: January
Year: 2017

5. **DOI:**
Title: Network Diversity: A Security Metric for Evaluating the Resilience of Networks against Zero-Day Attacks
Authors: L. Wang, M. Zhang, S. Jajodia, A. Singhal, and M. Albanese
Publication Title: IEEE Transactions on Information Forensics & Security
Month:
Year: 2015

6. **DOI:**
Title: Program Characterization Using Runtime Values and Its Application to Software Plagiarism Detection
Authors: Y.-C. Jhi, X. Wang, X. Jia, S. Zhu, Peng Liu, and D. Wu
Publication Title: IEEE Transactions on Software Engineering
Month:

Year: 2015

7. **DOI:**

Title: A methodology to measure and monitor level of operational effectiveness in a CSOC

Authors: Ankit Shah, Rajesh Ganesan, Sushil Jajodia, and Hasan Cam

Publication Title: International Journal of Information Security, Springer

Month:

Year: 2017

8. **DOI:**

Title: Using temporal probabilistic logic for optimal monitoring of security events with limited resources

Authors: Sushil Jajodia, Noseong Park, Edoardo Serra, and V.S. Subrahmanian

Publication Title: Journal of Computer Security, vol. 24, no. 6

Month:

Year: 2016

9. **DOI:**

Title: Impeding Behavior-based Malware Analysis via Replacement Attacks to Malware Specifications

Authors: Jiang Ming, Zhi Xin, Pengwei Lan, Dinghao Wu, Peng Liu, and Bing Mao

Publication Title: Journal of Computer Virology and Hacking Techniques, Springer

Month:

Year: 2016

10. **DOI:**

Title: Towards Service Continuity for Transactional Applications via Diverse Device Drivers

Authors: Shengzhi Zhang, Xiaoqi Jia, and Peng Liu

Publication Title: International Journal of Information and Computer Security

Month: November

Year: 2016

11. **DOI:**

Title: Dynamically Discovering Likely Memory Layout to Perform Accurate Fuzzing

Authors: Kai Chen, Yingjun Zhang, and Peng Liu

Publication Title: IEEE Transactions on Reliability

Month:

Year: 2016

12. **DOI:**

Title: Dynamic Games with Asymmetric Information: Common Information Based Perfect Bayesian Equilibria and Sequential Decomposition

Authors: E. Miehling, M. Rasouli, and D. Teneketzis

Publication Title: IEEE Transactions on Automatic Control

Month: January

Year: 2017

13. DOI: <https://doi.org/10.1109/MSP.2018.1331034>
Title: Defending from Stealthy Botnets Using Moving Target Defenses
Authors: Massimiliano Albanese, Sushil Jajodia, and Sridhar Venkatesan
Publication Title: IEEE Security and Privacy Magazine, vol. 16, no. 1
Month: January/February
Year: 2018
-
14. DOI: <https://doi.org/10.1177/1548512917706043>
Title: A Graphical Model to Assess the Impact of Multi-Step Attacks
Authors: Massimiliano Albanese and Sushil Jajodia
Publication Title: Journal of Defense Modeling and Simulation, vol. 15, no. 1
Month: January
Year: 2018
Note: Selected by the Guest Editor, Alexander Kott, as an article of particular value
-
15. DOI: <https://doi.org/10.1109/TIFS.2018.2819967>
Title: A POMDP Approach to the Dynamic Defense of Large-Scale Cyber Networks
Authors: E. Miehling, M. Rasouli, and D. Teneketzis
Publication Title: IEEE Transactions on Information Forensics and Security
Month: October
Year: 2018
-
16. DOI: <https://doi.org/10.1145/3173457>
Title: Dynamic Optimization of the Level of Operational Effectiveness of a CSOC under Adverse Conditions
Authors: Ankit Shah, Rajesh Ganesan, Sushil Jajodia, and Hasan Cam
Publication Title: ACM Transactions on Intelligent Systems and Technology
Month: July
Year: 2018
-
17. DOI: <https://doi.org/10.1007/s11761-018-0235-3>
Title: Adaptive Reallocation of Cybersecurity Analysts to Sensors for Balancing Risks between Sensors
Authors: Ankit Shah, Rajesh Ganesan, Sushil Jajodia, and Hasan Cam
Publication Title: Service Oriented Computing and Applications
Month: June
Year: 2018
-
18. DOI: <https://doi.org/10.1109/JSYST.2018.2809506>
Title: Optimal Assignments of Sensors to Analysts in a Cybersecurity Operations Center
Authors: Ankit Shah, Rajesh Ganesan, Sushil Jajodia, and Hasan Cam
Publication Title: IEEE Systems Journal
Month: April
Year: 2018
-
19. DOI: <https://doi.org/10.1186/s42400-018-0003-x>

Title: Feedback control can make data structure layout randomization more cost-effective under zero-day attacks.

Authors: Ping Chen, Zhisheng Hu, Jun Xu, Minghui Zhu, and Peng Liu

Publication Title: Cybersecurity

Month: June

Year: 2018

20. DOI: <https://doi.org/10.1109/TMC.2017.2782249>

Title: Leveraging Information Asymmetry to Transform Android Apps into Self-Defending Code against Repackaging Attacks

Authors: Kai Chen, Yingjun Zhang, and Peng Liu

Publication Title: IEEE Transactions on Mobile Computing

Month: August

Year: 2018

21. DOI: <https://doi.org/10.1109/TIFS.2018.2821095>

Title: Using Bayesian Networks for Probabilistic Identification of Zero-day Attack Paths

Authors: X. Sun, J. Dai, P. Liu, A. Singhal, and J. Yen

Publication Title: IEEE Transactions on Information Forensics and Security

Month: March

Year: 2018

22. DOI: <https://doi.org/10.1109/JSYST.2018.2828832>

Title: Learning from Experts' Experience: Towards Automated Cyber Security Data Triage

Authors: Chen Zhong, John Yen, Peng Liu, and Robert F. Erbacher

Publication Title: IEEE Systems Journal

Month: May

Year: 2018

23. DOI: <https://doi.org/10.1109/JIOT.2018.2847733>

Title: The Effect of IoT New Features on Security and Privacy: New Threats, Existing Solutions, and Challenges Yet to Be Solved

Authors: Wei Zhou, Yan Jia, Anni Peng, Yuqing Zhang, and Peng Liu

Publication Title: IEEE Internet of Things Journal

Month: June

Year: 2018

24. DOI: <https://doi.org/10.1016/j.cose.2018.02.011>

Title: A cyber security data triage operation retrieval system

Authors: C. Zhong, T. Lin, P. Liu, J. Yen, and K. Chen

Publication Title: Computers & Security (Elsevier)

Month: July

Year: 2018

25. DOI: <https://doi.org/10.1109/TBDATA.2018.2790439>

Title: A Large-scale Study of Android Malware Development Phenomenon on Public Malware Submission and Scanning Platform

- Authors:** H. Huang, C. Zheng, J. Zeng, W. Zhou, S. Zhu, P. Liu, I. Molloy, S. Chari, C. Zhang, and Q. Guan
Publication Title: IEEE Transactions on Big Data
Month: January
Year: 2018
-
26. **DOI:** <https://doi.org/10.1002/spe.2576>
Title: A Policy-Centric Approach to Protecting OS Kernel from Vulnerable LKMs
Authors: D. Tian, X. Xiong, C. Hu, and P. Liu
Publication Title: Software: Practice and Experience
Month: March
Year: 2018
-
27. **DOI:** <https://doi.org/10.1145/3196884>
Title: VULCON: A System for Vulnerability Prioritization, Mitigation, and Management.
Authors: Katheryn A. Farris, Ankit Shah, George Cybenko, Rajesh Ganesan, and Sushil Jajodia
Publication Title: ACM Transactions on Privacy and Security
Month: July
Year: 2018
-
28. **DOI:** <https://doi.org/10.1016/j.future.2019.05.032>
Title: Unsupervised Multi-Stage Attack Detection Framework without Details on Single-State Attacks
Authors: J. Shin, S. H. Choi, P. Liu, Y. H. Choi
Journal Title: Elsevier Future Generation of Computer Systems
Month: November
Year: 2019
-
29. **DOI:** <https://doi.org/10.1109/TMC.2019.2936561>
Title: Tainting-Assisted and Context-Migrated Symbolic Execution of Android Framework for Vulnerability Discovery and Exploit Generation
Authors: Lannan Luo, Qiang Zeng, Chen Cao, Kai Chen, Jian Liu, Limin Liu, Neng Gao, Min Yang, Xinyu Xing, and Peng Liu
Journal Title: IEEE Transactions on Mobile Computing
Month: August
Year: 2019
-
30. **DOI:** <https://doi.org/10.1109/TSE.2019.2939528>
Title: POMP++: Facilitating Postmortem Program Diagnosis with Value-set Analysis
Authors: Mu, D., Du, Y., Xu, J., Xu, J., Xing, X., Mao, B., Liu, P.
Journal Title: IEEE Transactions on Software Engineering
Month: August
Year: 2019
-
31. **DOI:** <https://doi.org/10.1109/TCC.2020.2970183>
Title: Semi-synchronized Non-blocking Concurrent Kernel Cruising
Authors: Authors: D. Tian, Q. Zeng, D. Wu, P. Liu, C. Hu

Journal Title: IEEE Transactions on Cloud Computing
Year: 2020

32. DOI: <https://doi.org/10.1145/3377554>

Title: Two can play that game: An adversarial evaluation of a cyber-alert inspection system

Authors: Ankit Shah, Arunesh Sinha, Rajesh Ganesan, Sushil Jajodia, Hasan Cam

Journal Title: ACM Transactions on Intelligent Systems and Technology

Month: May

Year: 2020

33. DOI: <https://doi.org/10.1109/TDSC.2020.2975789>

Title: A study of account recovery in the wild and its security implications

Authors: Yue Li, Zeyu Chen, Haining Wang, Kun Sun, Sushil Jajodia

Journal Title: IEEE Trans. on Dependable and Secure Computing

Month: March

Year: 2020

34. DOI: <https://doi.org/10.1109/JSYST.2020.2964985>

Title: Disclose or exploit? A game theoretic approach to strategic decision making in cyber warfare

Authors: Haiping Chen, Qian Han, Sushil Jajodia, Roy Lindelauf, V. S. Subrahmanian, Yanhai Xiong

Journal Title: IEEE Systems Journal

Month: September

Year: 2020

35. DOI: <https://doi.org/10.1145/3372498>

Title: An outsourcing model for alert analysis in a cybersecurity operations center

Authors: Ankit Shah, Rajesh Ganesan, Sushil Jajodia, Hasan Cam

Journal Title: ACM Trans. on the Web (TWEB)

Month: January

Year: 2020

36. DOI: <https://doi.org/10.1109/TIFS.2019.2954737>

Title: Understanding the manipulation on recommender systems through web injection

Authors: Yubao Zhang, Jidong Xiao, Shuai Hao, Haining Wang, Sencun Zhu, Sushil Jajodia

Journal Title: IEEE Trans. on Information Forensics and Security (TIFS)

Month: November

Year: 2020

37. DOI: <https://doi.org/10.1109/TPDS.2019.2927977>

Title: Adaptive alert management for balancing optimal performance among distributed CSOCs using reinforcement learning

Authors: Ankit Shah, Rajesh Ganesan, Sushil Jajodia, Pierangela Samarati, Hasan Cam

Journal Title: IEEE Trans. on Parallel and Distributed Systems (TPDS)

Month: January

Year: 2020

38. DOI: <https://doi.org/10.1177/1548512919874129>
Title: Vulnerability selection for remediation: An empirical analysis
Authors: Ankit Shah, Katheryn A. Farris, Rajesh Ganesan, Sushil Jajodia
Journal Title: Journal of Defense Modeling and Simulation
Month: September
Year: 2019
-
39. DOI: <https://doi.org/10.3233/JCS-191306>
Title: Mitigating the insider threat of remote administrators in clouds through maintenance task assignments
Authors: Nawaf Alhebaishi, Lingyu Wang, Sushil Jajodia, Anoop Singhal
Journal Title: Journal of Computer Security
Month: July
Year: 2019
-
40. DOI: <https://doi.org/10.1109/MIS.2018.2877280>
Title: AI and Fake News
Authors: Anne K. Cybenko and George Cybenko
Journal Title: IEEE Intelligent Systems
Month: December
Year: 2018
-
41. DOI: <https://doi.org/10.1155/2018/2864873>
Title: Multistage attack graph security games: Heuristic strategies, with empirical game-theoretic analysis
Authors: TH Nguyen, M Wright, MP Wellman, S Singh
Journal Title: Security and Communication Networks
Month: December
Year: 2018
-
42. DOI: <https://doi.org/10.1109/TIFS.2018.2819967>
Title: A POMDP Approach to the Dynamic Defense of Large-Scale Cyber Networks
Authors: E Miehl, M Rasouli, D Teneketzis
Journal Title: IEEE Transactions on Information Forensics and Security
Month: October
Year: 2018
-
43. DOI: <https://doi.org/10.1109/TDSC.2018.2882825>
Title: Performance Modeling of Moving Target Defenses with Reconfiguration Limits
Authors: Warren Connell, Daniel A. Menasce, and Massimiliano Albanese
Journal Title: IEEE Transactions on Dependable and Secure Computing
Month: November
Year: 2018
-
44. DOI: <https://doi.org/10.1002/spy2.51>
Title: DeBot: A novel network-based mechanism to detect exfiltration by architectural stealthy botnets

Authors: Sridhar Venkatesan, Massimiliano Albanese, Cho-Yu Jason Chiang, Angelo Sapello, and Ritu Chadha
Journal Title: Wiley Security and Privacy
Month: December
Year: 2018

45. **DOI:** <https://doi.org/10.1016/j.automatica.2019.02.032>
Title: On convergence rates of game theoretic **reinforcement** learning algorithms
Authors: Zhisheng Hu, Minghui Zhu, Ping Chen, and Peng Liu
Journal Title: Automatica
Month: June
Year: 2019
-
46. **DOI:** <https://doi.org/10.1109/TDSC.2018.2861756>
Title: Building a Trustworthy Execution Environment to Defeat Exploits from both Cyber Space and Physical Space for ARM
Authors: Le Guan, Chen Cao, Peng Liu, Xinyu Xing, Xinyang Ge, Shengzhi Zhang, Meng Yu and Trent Jaeger
Journal Title: IEEE Transactions on Dependable and Secure Computing
Month: May
Year: 2019
-
47. **DOI:** <https://doi.org/10.1109/TMC.2017.2782249>
Title: Leveraging Information Asymmetry to Transform Android Apps into Self-Defending Code against Repackaging Attacks
Authors: Kai Chen, Yingjun Zhang, Peng Liu
Journal Title: IEEE Transactions on Mobile Computing
Month: August
Year: 2018
-
48. **DOI:** <https://doi.org/10.1016/j.future.2019.05.032>
Title: Unsupervised multi-stage attack detection framework without details on single-stage attacks
Authors: Jinmyeong Shin, Seok Hwan Choi, Peng Liu, Yoon Ho Choi
Journal Title: Elsevier Future Generation Computer Systems
Month: November
Year: 2019
-
49. **DOI:** <https://doi.org/10.1145/3173457>
Title: Dynamic Optimization of the Level of Operational Effectiveness of a CSOC under Adverse Conditions
Authors: Ankit Shah, Rajesh Ganesan, Sushil Jajodia, Hasan Cam,
Journal Title: ACM Trans. on Intelligent Systems and Technology (TIST)
Volume: 9
Number: 5
Year: 2018

50. DOI: <https://doi.org/10.1007/s11761-018-0235-3>

Title: Adaptive Reallocation of Cybersecurity Analysts to Sensors for Balancing Risks between Sensors

Authors: Ankit Shah, Rajesh Ganesan, Sushil Jajodia, Hasan Cam

Journal Title: Service Oriented Computing and Applications (SOCA), Springer

Volume: 12

Number: 2

Year: 2018

Pages: 122-135

51. DOI: <https://doi.org/10.1109/JSYST.2018.2809506>

Title: Optimal Assignments of Sensors to Analysts in a Cybersecurity Operations Center

Authors: Ankit Shah, Rajesh Ganesan, Sushil Jajodia, Hasan Cam

Journal: IEEE Systems Journal

Volume: 13

Number: 1

Year: 2019

Pages: 1060-1071

52. DOI: <https://doi.org/10.3233/JCS-171094>

Title: Hybrid Adversarial Defense: Merging Honeypots and Traditional Security Methods

Authors: Tanmoy Chakraborty, Sushil Jajodia, Noseong Park, Andrea Pugliese, Edoardo Serra, V. S. Subrahmanian

Journal Title: *Journal of Computer Security (JCS)*

Volume: 26

Number: 5

Year: 2018

Pages: 615-645

53. DOI: <https://doi.org/10.1109/TDSC.2019.2898661>

Title: FORGE: A Fake Online Repository Generation Engine for Cyber Deception

Authors: Tanmoy Chakraborty, Sushil Jajodia, Jonathan Katz, Antonio Picariello, Giancarlo Sperli, V. S. Subrahmanian

Journal Title: *IEEE Trans. on Dependable and Secure Computing (TDSC)*

Year: 2019

54. DOI: <https://doi.org/10.3233/JCS-191306>

Title: Mitigating the insider threat of remote administrators in clouds through maintenance task assignments

Authors: Nawaf Alhebaishi, Lingyu Wang, Sushil Jajodia, Anoop Singhal

Journal Title: *Journal of Computer Security*

Volume: 27

Number: 4

Year: 2019

Pages: 427-458

55. DOI: <https://doi.org/10.1016/j.comcom.2019.06.004>

Title: Optimizing the network diversity to improve the resilience of networks against the unknown attacks

Authors: Daniel Borbor, Lingyu Wang, Sushil Jajodia, Anoop Singhal

Journal: *Computer Communications*

Volume: 145

Year: 2019

Pages: 96-112

56. DOI: <https://doi.org/10.1109/TIFS.2018.2886465>

Title: A two-step approach to optimal selection of alerts for investigation in a CSOC

Authors: Ankit Shah, Rajesh Ganesan, Sushil Jajodia, Hasan Cam

Journal Title: *IEEE Trans. on Information Forensics and Security (TIFS)*

Vol. 14

No. 7,

Year: 2019

Pages: 1857-1870

57. DOI: <https://doi.org/10.1007/s41109-019-0209-1>

Title: Analytics for directed contact networks

Keywords: Temporal networks, directed contact networks, Markov models

Abstract: *Directed contact networks* (DCNs) are temporal networks that are useful for analyzing and modeling phenomena in transportation, communications, epidemiology and social networking. Specific sequences of contacts can underlie higher-level behaviors such as *flows* that aggregate contacts based on some notion of semantic and temporal proximity. We describe a simple inhomogeneous Markov model to infer flows and taint bounds associated with such higher-level behaviors, and also discuss how to aggregate contacts within DCNs and/or dynamically cluster their vertices. We provide examples of these constructions in the contexts of information transfers within computer and air transportation networks, thereby indicating how they can be used for data reduction and anomaly detection.

Authors: George Cybenko and Steve Huntsman

Journal Title: *Applied Network Science*

Month: November

Year: 2019

58. DOI: <https://doi.org/10.1109/MCS.2020.2976390>

Title: Linear Algebra and Learning From Data

Keywords: Linear algebra, machine learning

Abstract: This book contains the key linear algebra and optimization techniques at the forefront of active data-science and machine learning practice today. This is an appropriate choice of content because while state-of-the-art machine learning applications can change each month (as in reinforcement learning, language translation, game playing, or image classification), the

underlying mathematical concepts and algorithms do not.

Authors: George Cybenko

Journal Title: IEEE Control Systems Magazine

Month: June

Year: 2020

59. **DOI:** <https://doi.org/10.1145/3418897>

Title: Adaptive cyber defense against multi-stage attacks using learning-based POMDP

Keywords: Adaptive cyber defense, reinforcement learning, POMDP

Abstract: Growing multi-stage attacks in computer networks impose significant security risks and necessitate the development of effective defense schemes which are able to autonomously respond to intrusions during vulnerability windows. However, the defender faces several real-world challenges; e.g., unknown likelihoods and unknown impacts of successful exploits. In this paper, we leverage reinforcement learning to develop an innovative adaptive cyber defense to maximize the cost-effectiveness subject to the aforementioned challenges. In particular, we use Bayesian attack graphs to model the interactions between the attacker and networks. Then we formulate the defense problem of interest as a partially observable Markov decision process problem where the defender maintains belief states to estimate system states, leverages Thompson sampling to estimate transition probabilities and utilizes reinforcement learning to choose optimal defense actions using measured utility values. The algorithm performance is verified via numerical simulations based on real-world attacks.

Authors: Z. Hu, M. Zhu and P. Liu

Journal Title: ACM Transactions on Privacy and Security

Year: 2020

60. **DOI:** <https://doi.org/10.3233/JCS-191368>

Title: DEEPRETURN: A deep neural network can learn how to detect previously-unseen ROP payloads without using any heuristics

Keywords: Return-oriented programming, intrusion detection system, deep learning

Abstract: Return-oriented programming (ROP) is a code reuse attack that chains short snippets of existing code to perform arbitrary operations on target machines. Existing detection methods against ROP exhibit unsatisfactory detection accuracy and/or have high runtime overhead. In this paper, we present DEEPRETURN, which innovatively combines address space layout guided disassembly and deep neural networks to detect ROP payloads. The disassembler treats application input data as code pointers and aims to find any potential gadget chains, which are then classified by a deep neural network as benign or malicious. Our experiments show that DEEPRETURN has high detection rate (99.3%) and a very low false positive rate (0.01%). DEEPRETURN successfully detects all of the 100 real-world ROP exploits that are collected in-the-wild, created manually or created by ROP exploit generation tools. DEEPRETURN is non-intrusive and does not incur any runtime overhead to the protected program.

Authors: X. Li, Z. Hu, H. Wang, Y. Fu, P. Chen, M. Zhu and P. Liu.

Journal Title: Journal of Computer Security

Year: 2020

61. DOI: <https://doi.org/10.1016/j.future.2019.05.032>

Title: Unsupervised Multi-Stage Attack Detection Framework without Details on Single-State Attacks

Keywords: intrusion detection, cyber kill-chain, detection rule generation

Abstract: Majority of network attacks currently consist of sophisticated multi-stage attacks, which break down network attacks into several single-stage attacks. The early multi-stage attack detection methods focused on describing the detection rule based on the occurrence sequence of each single-stage attacks. That is, such works assumed that after details on single-stage attack behavior are obtained from attack knowledge, attack semantics or attack statistical analysis, the detection rules can be generated from their possible occurrence sequence. However, their practical usage is limited due to the high false negative ratio while detecting multi-stage attack that consists of diverse combinations of single-stage attacks during the long time period. In this paper, we propose a new multi-stage attack detection framework, which consists of multi-stage attack detection rule generation phase and multi-stage attack detection phase. After comparing the incoming traffics with the generated multi-stage attack detection rules, various multi-stage attack patterns are detected without pre-observed details on the single stage attack behavior. From DARPA LLS DDoS dataset, we show that all the possible multi-stage attack patterns are correctly detected. Also, from datasets in CTU-13 including the large volume of multi-stage attack patterns, we observe F1-measure of 0.938 at maximum.

Authors: J. Shin, S. H. Choi, P. Liu, Y. H. Choi

Journal Title: Elsevier Future Generation of Computer Systems

Month: November

Year: 2019

62. DOI: <https://doi.org/10.1109/TMC.2019.2936561>

Title: Tainting-Assisted and Context-Migrated Symbolic Execution of Android Framework for Vulnerability Discovery and Exploit Generation

Keywords: symbolic execution, vulnerability discovery, exploit generation, Android framework

Abstract: Android Application Framework is an integral and foundational part of the Android system. Each of the two billion (as of 2017) Android devices relies on the system services of Android Framework to manage applications and system resources. Given its critical role, a vulnerability in the framework can be exploited to launch large-scale cyber attacks and cause severe harms to user security and privacy. Recently, many vulnerabilities in Android Framework were exposed, showing that it is indeed vulnerable and exploitable. While there is a large body of studies on Android application analysis, research on Android Framework analysis is very limited. In particular, to our knowledge, there is no prior work that investigates how to enable symbolic execution of the framework, an approach that has proven to be very powerful for vulnerability discovery and exploit generation. We design and build the first system, CENTAUR, that enables symbolic execution of Android Framework. Due to the middleware nature and technical peculiarities of the framework that impinge on the analysis, many unique challenges arise and are addressed in CENTAUR. The system has been applied to discovering new vulnerability instances, which can be exploited by recently uncovered attacks against the

framework, and to generating PoC exploits.

Authors: Lannan Luo, Qiang Zeng, Chen Cao, Kai Chen, Jian Liu, Limin Liu, Neng Gao, Min Yang, Xinyu Xing, and Peng Liu

Journal Title: IEEE Transactions on Mobile Computing

Month: August

Year: 2019

63. DOI: <https://doi.org/10.1109/TSE.2019.2939528>

Title: POMP++: Facilitating Postmortem Program Diagnosis with Value-set Analysis

Keywords: Postmortem Program Diagnosis, Failure Diagnosis, Reverse Execution, Value-set Analysis

Abstract: With the emergence of hardware-assisted processor tracing, execution traces can be logged with lower runtime overhead and integrated into the core dump. In comparison with an ordinary core dump, such a new post-crash artifact provides software developers and security analysts with more clues to a program crash. However, existing works only rely on the resolved runtime information, which leads to the limitation in data flow recovery within long execution traces. In this work, we propose POMP++, an automated tool to facilitate the analysis of post-crash artifacts. More specifically, POMP++ introduces a reverse execution mechanism to construct the data flow that a program followed prior to its crash. Furthermore, POMP++ utilizes Value-set Analysis, which helps to verify memory alias relation, to improve the ability of data flow recovery. With the restored data flow, POMP++ then performs backward taint analysis and highlights program statements that actually contribute to the crash. We have implemented POMP++ for Linux system on x86-32 platform, and tested it against various crashes resulting from 31 distinct real-world security vulnerabilities. The evaluation shows that, our work can pinpoint the root causes in 29 cases, increase the number of recovered memory addresses by 12% and reduce the execution time by 60% compared with existing reverse execution. In short, POMP++ can accurately and efficiently pinpoint program statements that truly contribute to the crashes, making failure diagnosis significantly convenient.

Authors: Mu, D., Du, Y., Xu, J., Xu, J., Xing, X., Mao, B., Liu, P.

Journal Title: IEEE Transactions on Software Engineering

Month: August

Year: 2019

64. DOI: <https://doi.org/10.1109/TCC.2020.2970183>

Title: Semi-synchronized Non-blocking Concurrent Kernel Cruising

Keywords: kernel heap buffer overflow, concurrent monitoring

Abstract: Abstract: Kernel heap buffer overflow vulnerabilities have been exposed for decades, but there are few practical countermeasures that can be applied to OS kernels. Previous solutions either suffer from high performance overhead or compatibility problems with mainstream kernels and hardware. In this paper, we present KRUISER, a concurrent kernel heap buffer overflow monitor. Unlike conventional methods, the security enforcement of which is usually inlined into the kernel execution, Kruiser migrates security enforcement from the kernel's normal execution to a concurrent monitor process, leveraging the increasingly popular

multi-core architectures. To reduce the synchronization overhead between the monitor process and the running kernel, we design a novel semi-synchronized non-blocking monitoring algorithm, which enables efficient runtime detection on live memory without incurring false positives. To prevent the monitor process from being tampered and provide guaranteed performance isolation, we utilize the virtualization technology to run the monitor process out of the monitored VM, while heap memory allocation information is collected inside the monitored VM in a secure and efficient way. The hybrid VM monitoring technique combined with the secure canary that cannot be counterfeited by attackers provides guaranteed overflow detection with high efficiency. We have implemented a prototype of KRUISER based on Linux and the Xen/KVM hypervisor. The evaluation shows that Kruiser can detect realistic kernel heap buffer overflow attacks in cloud environment effectively with minimal cost.

Authors: Authors: D. Tian, Q. Zeng, D. Wu, P. Liu, C. Hu

Journal Title: IEEE Transactions on Cloud Computing

Year: 2020

65. **DOI:** <https://doi.org/10.1145/3377554>

Title: Two can play that game: An adversarial evaluation of a cyber-alert inspection system

Keywords: Cyber-security operations center, adversarial reinforcement learning, game theory

Abstract: Cyber-security is an important societal concern. Cyber-attacks have increased in numbers as well as in the extent of damage caused in every attack. Large organizations operate a Cyber Security Operation Center (CSOC), which forms the first line of cyber-defense. The inspection of cyber-alerts is a critical part of CSOC operations (defender or blue team). Recent work proposed a reinforcement learning (RL) based approach for the defender's decision-making to prevent the cyber-alert queue length from growing large and overwhelming the defender. In this article, we perform a red team (adversarial) evaluation of this approach. With the recent attacks on learning-based decision-making systems, it is even more important to test the limits of the defender's RL approach. Toward that end, we learn several adversarial alert generation policies and the best response against them for various defender's inspection policy. Surprisingly, we find the defender's policies to be quite robust to the best response of the attacker. In order to explain this observation, we extend the earlier defender's RL model to a game model with adversarial RL, and show that there exist defender policies that can be robust against any adversarial policy. We also derive a competitive baseline from the game theory model and compare it to the defender's RL approach. However, when we go further to exploit the assumptions made in the Markov Decision Process (MDP) in the defender's RL model, we discover an attacker policy that overwhelms the defender. We use a double oracle like approach to retrain the defender with episodes from this discovered attacker policy. This made the defender robust to the discovered attacker policy and no further harmful attacker policies were discovered. Overall, the adversarial RL and double oracle approach in RL are general techniques that are applicable to other RL usage in adversarial environments.

Authors: Ankit Shah, Arunesh Sinha, Rajesh Ganesan, Sushil Jajodia, Hasan Cam

Journal Title: ACM Transactions on Intelligent Systems and Technology

Month: May

Year: 2020

66. DOI: <https://doi.org/10.1109/TDSC.2020.2975789>

Title: A study of account recovery in the wild and its security implications

Keywords: Password Management, Authentication, Secure Communications

Abstract: Account recovery (usually through a password reset) on many websites has mainly relied on accessibility to a registered email, due to its favorable deployability and usability. However, it makes a user's online accounts vulnerable to a single point of failure when the registered email account is compromised. While previous research focuses on strengthening user passwords, the security risk imposed by email-based password recovery has not yet been well studied. In this paper, we first conduct a measurement study to characterize the password recovery activities in the wild. Specifically, we examine the authentication and password recovery protocols from 239 traffic-heavy websites, confirming that most of them use emails for password recovery. We further scrutinize the security policy of leading email service providers and show that a significant portion of them takes no or marginal effort to protect user email accounts, leaving compromised email accounts readily available for mounting password recovery attacks. Then, we conduct case studies to assess potential losses caused by such attacks. Finally, we propose and implement a lightweight email security enhancement called Secure Email Account Recovery (SEAR) to defend against password recovery attacks by adding an extra layer of protection to password recovery emails.

Authors: Yue Li, Zeyu Chen, Haining Wang, Kun Sun, Sushil Jajodia

Journal Title: IEEE Trans. on Dependable and Secure Computing

Month: March

Year: 2020

67. DOI: <https://doi.org/10.1109/JSYST.2020.2964985>

Title: Disclose or exploit? A game theoretic approach to strategic decision making in cyber warfare

Keywords: Cyber-security, decision support system, game theory, Nash equilibrium

Abstract: Today, countries are engaged in the cyber-arms race. With over 16 K new hardware/software vulnerabilities reported in 2018 alone, an important question confronts senior government decision makers when their cyber-warfare units discover a new vulnerability. Should they disclose the vulnerability to the vendor who produced the vulnerable product? Or should they “stockpile” the vulnerability, holding it for developing exploits (i.e., cyber-weapons) that can be targeted at an adversary? Choosing the first option may be important when the affected company is a corporation in the nation state that discovers the vulnerability and/or if that nation state would have a big exposure to that vulnerability. Choosing the second option has obvious advantages to the discovering nation's defense. We formulate the cyber-competition between countries as a repeated cyber-warfare game (RCWG), where two countries (players) compete over a series of vulnerabilities by deciding, at the time of vulnerability discovery, 1) whether to exploit or disclose it and 2) how long to exploit it if they decide to exploit. We define the equilibrium state of the RCWG as a pure strategy Nash

equilibrium, and propose a learning-while-competing framework to compute the pure strategy Nash equilibrium of the formulated RCWG. Although testing our results with real data in the murky world of cyber-warfare is challenging, we were able to obtain real statistics from other sources and demonstrate the effectiveness of our proposed algorithm through a set of simulation results under different scenarios using these third-party statistics. We also report on our DiscX system that can help support government decision makers in their decision whether to disclose or exploit a vulnerability that they find.

Authors: Haiping Chen, Qian Han, Sushil Jajodia, Roy Lindelauf, V. S. Subrahmanian, Yanhai Xiong

Journal Title: IEEE Systems Journal

Month: September

Year: 2020

68. **DOI:** <https://doi.org/10.1145/3372498>

Title: An outsourcing model for alert analysis in a cybersecurity operations center

Keywords: Cyber Security Operations Center, Outsourcing, on-demand analysts, optimization

Abstract: A typical Cybersecurity Operations Center (CSOC) is a service organization. It hires and trains analysts, whose task is to perform analysis of alerts that were generated while monitoring the client's networks. Due to ever-increasing financial and infrastructure burden on a CSOC driven by the rapidly growing demand for security services, it would become prohibitively expensive to continually expand the size of a CSOC to meet the demands in the future. An alternative solution is to outsource the alert analysis process to on-demand analysts, to provide scalable CSOC service to its clients with features, such as (1) higher throughput, (2) higher quality, and (3) more economical service than the current in-house service. The current outsourcing model is not cost effective and an exact optimization model is computationally inefficient. This article presents a novel two-step sequential mixed integer programming optimization method that is used in the development of a new decision-support business model for outsourcing the alert analysis process. It is demonstrated that through this model, a CSOC can effectively deliver its alert management services with the above-mentioned features. Results indicate that the model is scalable, computationally viable, real-time implementable, and can deliver CSOC services that meet the service-level agreement (SLA) between the CSOC and its client. In addition, the article provides valuable insights into the cost of operating the new business process outsourcing model for cybersecurity services.

Authors: Ankit Shah, Rajesh Ganesan, Sushil Jajodia, Hasan Cam

Journal Title: ACM Trans. on the Web (TWEB)

Month: January

Year: 2020

69. **DOI:** <https://doi.org/10.1109/TIFS.2019.2954737>

Title: Understanding the manipulation on recommender systems through web injection

Keywords: Recommender systems, YouTube, web injection

Abstract: Recommender systems have been increasingly used in a variety of web services, providing a list of recommended items in which a user may have an interest. While important,

recommender systems are vulnerable to various malicious attacks. In this paper, we study a new security vulnerability in recommender systems caused by web injection, through which malicious actors stealthily tamper any unprotected in-transit HTTP webpage content and force victims to visit specific items in some web services (even running HTTPS), e.g., YouTube. By doing so, malicious actors can promote their targeted items in those web services. To obtain a deeper understanding on the recommender systems of our interest (including YouTube, Yelp, Taobao, and 360 App market), we first conduct a measurement-based analysis on several real-world recommender systems by leveraging machine learning algorithms. Then, web injection is implemented in three different types of devices (i.e., computer, router, and proxy server) to investigate the scenarios where web injection could occur. Based on the implementation of web injection, we demonstrate that it is feasible and sometimes effective to manipulate the real-world recommender systems through web injection. We also present several countermeasures against such manipulations.

Authors: Yubao Zhang, Jidong Xiao, Shuai Hao, Haining Wang, Sencun Zhu, Sushil Jajodia

Journal Title: IEEE Trans. on Information Forensics and Security (TIFS)

Month: November

Year: 2020

70. DOI: <https://doi.org/10.1109/TPDS.2019.2927977>

Title: Adaptive alert management for balancing optimal performance among distributed CSOCs using reinforcement learning

Keywords: Distributed cybersecurity operations center (CSOC), centralized alert management, level of operational effectiveness, reinforcement learning, and adaptive resource allocation

Abstract: Large organizations typically have Cybersecurity Operations Centers (CSOCs) distributed at multiple locations that are independently managed, and they have their own cybersecurity analyst workforce. Under normal operating conditions, the CSOC locations are ideally staffed such that the alerts generated from the sensors in a work-shift are thoroughly investigated by the scheduled analysts in a timely manner. Unfortunately, when adverse events such as increase in alert arrival rates or alert investigation rates occur, alerts have to wait for a longer duration for analyst investigation, which poses a direct risk to organizations. Hence, our research objective is to mitigate the impact of the adverse events by dynamically and autonomously re-allocating alerts to other location(s) such that the performances of all the CSOC locations remain balanced. This is achieved through the development of a novel centralized adaptive decision support system whose task is to re-allocate alerts from the affected locations to other locations. This re-allocation decision is non-trivial because the following must be determined: (1) timing of a re-allocation decision, (2) number of alerts to be reallocated, and (3) selection of the locations to which the alerts must be distributed. The centralized decision-maker (henceforth referred to as agent) continuously monitors and controls the level of operational effectiveness-LOE (a quantified performance metric) of all the locations. The agent's decision-making framework is based on the principles of stochastic dynamic programming and is solved using reinforcement learning (RL). In the experiments, the RL approach is compared with both rule-based and load balancing strategies. By simulating real-

world scenarios, learning the best decisions for the agent, and applying the decisions on sample realizations of the CSOC's daily operation, the results show that the RL agent outperforms both approaches by generating (near-) optimal decisions that maintain a balanced LOE among the CSOC locations. Furthermore, the scalability experiments highlight the practicality of adapting the method to a large number of CSOC locations.

Authors: Ankit Shah, Rajesh Ganesan, Sushil Jajodia, Pierangela Samarati, Hasan Cam

Journal Title: IEEE Trans. on Parallel and Distributed Systems (TPDS)

Month: January

Year: 2020

71. DOI: <https://doi.org/10.1177/1548512919874129>

Title: Vulnerability selection for remediation: An empirical analysis

Keywords: Vulnerability management, vulnerability remediation, cyber-security operations center, multiple attribute value optimization, mixed integer programming

Abstract: Vulnerabilities are security flaws in software and network systems that criminal hackers can exploit to gain an asymmetric advantage. Cyber-Security Operations Centers must routinely triage and patch vulnerabilities in their system(s) to minimize external exposure to attackers. The personnel resources required to address vulnerability remediation tasks are limited and constrained, thus motivating the need for optimization approaches to improve the efficiency of the vulnerability selection process. This paper investigates two different approaches to vulnerability selection for mitigation through (a) Individual Attribute Value Optimization and (b) Multiple Attribute Value Optimization. The former approach presents a methodology that optimizes the selection of vulnerabilities for mitigation with respect to an individual attribute, while the latter approach considers multiple attributes in the vulnerability selection decision-making. Real scan data from a Cyber-Security Operations Center are used to compare the results between the two mathematical approaches. Furthermore, comparisons are made with the results obtained from (a) the actual (baseline) Cyber-Security Operations Center performance, and (b) a vulnerability prioritization algorithm called VULCON that appeared in recent literature.

Authors: Ankit Shah, Katheryn A. Farris, Rajesh Ganesan, Sushil Jajodia

Journal Title: Journal of Defense Modeling and Simulation

Month: September

Year: 2019

72. DOI: <https://doi.org/10.3233/JCS-191306>

Title: Mitigating the insider threat of remote administrators in clouds through maintenance task assignments
Keywords: Security metric, cloud security, insider attack, attack graph, service dependency graph

Abstract: Abstract: Today's cloud providers strive to attract customers with better services and less downtime in a highly competitive market. The need for minimizing the operational cost unavoidably leads cloud providers to rely on third party remote administrators for fulfilling regular maintenance tasks. In such a scenario, the lack of trust in those third party remote administrators paired with the extra privileges granted to them to complete the maintenance

tasks usually implies undesirable security threats. A dishonest remote administrator, or an attacker armed with the stolen credential of a remote administrator, can pose severe insider threats to both the cloud provider and its tenants. In this paper, we take the first step towards understanding and mitigating such insider threats of remote administrators in clouds. Specifically, we first model the maintenance task assignments and their corresponding security impact due to privilege escalation. We then mitigate such impact through optimizing the task assignments with respect to given constraints. Finally, the simulation results demonstrate the effectiveness of our solution in various scenarios.

Authors: Nawaf Alhebaishi, Lingyu Wang, Sushil Jajodia, Anoop Singhal

Journal Title: Journal of Computer Security

Month: July

Year: 2019

PAPERS PUBLISHED IN NON-PEER-REVIEWED JOURNALS

1. Nuclear Command and Control (classified study report of the Air Force Scientific Advisory Board, George Cybenko co-authored the section on cybersecurity), 2015.
2. Cyber Vulnerabilities of Embedded Systems (classified study report of the Air Force Scientific Advisory Board, George Cybenko co-authored the section on commercial approaches), 2016.

PRESENTATIONS AT MEETINGS, BUT NOT PUBLISHED IN CONFERENCE PROCEEDINGS

1. Date: October 28, 2013
Speaker(s): Patrick Sweeney and George Cybenko
Title: Quantifying covertness (invited talk)
Event/venue: MIT Cyber Security Metrics meeting
2. Date: November 21, 2013
Speaker(s): George Cybenko
Title: Cyber Modeling and Evaluation (invited talk)
Event/venue: DARPA, Arlington VA
3. Date: December 5, 2013
Speaker(s): George Cybenko
Title: Quantifying Covertness (invited talk)
Event/venue: Systems Technology Research, Woburn MA
4. Date: December 19, 2013
Speaker(s): George Cybenko
Title: Quantifying Covertness (seminar)
Event/venue: Applied Communications Systems, New Jersey
5. Date: January 3, 2014
Speaker(s): George Cybenko
Title: Quantitative Cyber Security Risk Assessment (seminar)
Event/venue: CSE, Canada, Ottawa
6. Date: May 13, 2014
Speaker(s): George Cybenko
Title: Data Integrity (seminar)
Event/venue: Kinetic, Reston VA. (ODNI C3E Meeting)
7. Date: December 18, 2014
Speaker(s): Sushil Jajodia
Title: Adversarial and Uncertain Reasoning for Adaptive Cyber Defense: Building the Scientific Foundation
Event/venue: 10th International Conference on Information Systems Security, Hyderabad, India
8. Date: February 11, 2015
Speaker(s): Sushil Jajodia and Rajesh Ganesan
Title: Optimal Scheduling of Cybersecurity Analysts for Minimizing Risk
Event/venue: Army Research Laboratory, Adelphi, MD

9. Date: September 30, 2014
 Speaker(s): George Cybenko
 Title: Deep Learning of Behaviors^[1]for Security
 Event/venue: MIT Lincoln Laboratory Processing, Exploitation and Dissemination Seminar
10. Date: October 8, 2014
 Speaker(s): George Cybenko
 Title: Cyber Security Metrics
 Event/venue: Sprott School of Business, Carleton University Ottawa Canada Colloquium Series
11. Date: October 21, 2014
 Speaker(s): George Cybenko
 Title: Semantic and Syntactic Data Integrity
 Event/venue: Georgia Tech Research Institute, I3P Workshop
12. Date: November 3, 2014
 Speaker(s): George Cybenko
 Title: No free lunch in cyber security
 Event/venue: Keynote at First ACM Workshop on Moving Target Defense, Scottsdale AZ
13. Date: November 19, 2014
 Speaker(s): George Cybenko
 Title: Deep Learning of Behaviors^[1]for Security
 Event/venue: Computer Science Colloquium, Purdue University, W. Lafayette IN
14. Date: December 3, 2014
 Speaker(s): George Cybenko
 Title: Cyber Attack Kill Chains
 Event/venue: DARPA Workshop, Arlington VA.
15. Date: March 4, 2015
 Speaker(s): George Cybenko
 Title: Deep Learning of Behaviors for Security
 Event/venue: Keynote at 2015 ACM International Workshop on International Workshop on Security and Privacy Analytics, San Antonio TX
16. Date: March 14, 2015
 Speaker(s): George Cybenko
 Title: Deep Learning of Behaviors for Security
 Event/venue: Keynote at CRIPS 2015, Deauville France
17. Date: April 21, 2015
 Speaker(s): George Cybenko
 Title: Deep Learning of Behaviors for Security
 Event/venue: Keynote at SPIE Defense and Security Symposium 2015, Baltimore MD
18. Date: April 21, 2015
 Speaker(s): George Cybenko (with Jeff Hughes)
 Title: Enterprise scale cyber risk assessment and management
 Event/venue: SPIE Defense and Security Symposium 2015, Baltimore MD
19. Date: June 29, 2015
 Speaker(s): George Cybenko

- Title: Deep Learning of Behaviors for Security
Event/venue: Department of Computer Science Colloquium, ULM University, Ulm Germany
20. Date: July 1, 2015
Speaker(s): George Cybenko
Title: Deep Learning of Behaviors for Security
Event/venue: Keynote at 12th International Workshop on Multiple Classifier Systems, Ulm Germany
21. Date: July 14, 2015
Speaker(s): George Cybenko
Title: Deep Learning of Behaviors for Security
Event/venue: Colloquium at Center for Communications Research, Princeton NJ
22. Date: July 21-22, 2015
Speaker(s): George Cybenko
Title: Short Course on Adversarial and Adaptive Cyber Defenses (3 lectures)
Event/venue: Graduate Summer School: Games and Contracts for Cyber-Physical Security, IPAM UCLA
23. Date: August 19, 2017
Speaker(s): Sushil Jajodia
Title: Adversarial and Uncertain Reasoning for Adaptive Cyber Defense
Event/venue: IEEE 5G Summit
24. Date: March 23, 2017
Speaker(s): Sushil Jajodia
Title: Data Analytics/Internet of Things
Event/venue: President's Forum on Data, University at Albany
25. Date: March 2, 2017
Speaker(s): Sushil Jajodia
Title: A Mission-centric Framework for Cyber Situation Awareness
Event/venue: DoD Cyber Modeling & Simulation Group
26. Date: October 31, 2016
Speaker(s): Michael Wellman
Title: Experience with Autonomous Agents: What Can We Learn
Event/venue: Hastings Center Experts Workshop on Control and Responsible Innovation in the Development of Autonomous Machines
27. Date: January 2017
Speaker(s): Michael Wellman
Title: Strategic considerations for learning agents
Event/venue: Panel on Machine Learning and Security (and Privacy), NSF SaTC Program Meeting, Washington, DC
28. Date: February 24, 2017
Speaker(s): Michael Wellman
Title: Adverse Outcomes for Financial Markets: AI, Trading Systems, and Market Manipulation
Event/venue: Origins Project Workshop on Envisioning and Addressing Adverse Outcomes of Artificial Intelligence, Arizona State University, Tempe, AZ
29. Date: June 11, 2017
Speaker(s): Erik Miehl

- Title:* A Dependency Graph Formalism for the Dynamic Defense of Large-Scale Cyber Networks (poster)
Event/venue: 3rd Multidisciplinary Conference on Reinforcement Learning and Decision Making (RLDM 2017), University of Michigan, Ann Arbor
30. *Date:* July 6, 2017
Speaker(s): Michael Wellman
Title: Why is AI important now, and what should we, as a nation, be doing about it?
Event/venue: Comptroller General Forum on Artificial Intelligence, Washington, DC
31. Jajodia, *"Hybrid adversarial defense: Merging traditional security defenses and honeypots,"* Invited talk, ARO Workshop on HoneyThings: Autonomous and Resilient Cyber Deception, University of North Carolina Charlotte, January 29, 2018
32. Jajodia, *"Maintaining level of operational effectiveness of a CSOC under adverse conditions,"* Invited talk, NSF Organizational Science and Cyber Security Workshop, Fairfax, July 2018.
33. Jajodia, *"Maintaining level of operational effectiveness of a CSOC under adverse conditions,"* Keynote, 13th International Conference on Information System Security, Mumbai, India, December 18, 2017
34. Jajodia, *"Adversarial and uncertain reasoning for adaptive cyber defense: Building the scientific foundations,"* Invited Talk, Fall Seminar Series, Florida International University, Miami, October 27, 2017
35. Jajodia, *"Adversarial and uncertain reasoning for adaptive cyber defense: Building the scientific foundations,"* Keynote, IEEE International 5G Summit, Reston, August 19, 2017
36. Jajodia, *"Optimal management of cyber security operations center,"* Invited talk, Cyber Security Technology Summit, General Dynamics Information Technology, Springfield, VA, October 18, 2017
37. Teneketzis, *"Control and Game Theoretic Approaches to Cyber Security,"* Plenary lecture, IEEE Symposium on Privacy and Security, Montréal, November 2017.
38. Sinha, *"Putting Learning in Context,"* ARO Invitational Workshop on Foundations of Autonomous Adaptive Cyber Systems, George Mason University, February 2018.
39. Wellman, *"Simulation-Based Game Reasoning for Autonomous Cyber-Defense,"* ARO Invitational Workshop on Foundations of Autonomous Adaptive Cyber Systems, George Mason University, February 2018.
40. Liu, *"Can Cyber Operations be Made Autonomous? An Answer from the Situational Awareness Viewpoint,"* ARO Invitational Workshop on Foundations of Autonomous Adaptive Cyber Systems, George Mason University, February 2018.
41. Liu, *"Honeypots in Firmware-enabled Hardware Devices: Basic Concepts and A Proposal,"* ARO Invitational Workshop on Honey Things, UNCC, January 2018.
42. Cybenko, *"Quantifying Cyber Resilience,"* keynote at US Army Cyber Bootcamp, UC Davis, August 2017
43. Cybenko, *"Machine Learning in Cyber Security,"* keynote talk at the Data Science and Security Conference, Imperial College, London UK, September 2017.
44. Cybenko, *"Machine Learning and Cyber Deception,"* invited talk at the ACM Conference on Computer Security, Dallas TX, October 2017.
45. Cybenko, *"Foundations of Cyber Autonomy,"* US Army Workshop on Cyber Autonomy, George Mason University, February, 2018.
46. Cybenko, *"Machine Learning in Cyber Security,"* SEI CTO Distinguished Lecture, Carnegie Mellon University, Software Engineering Institute, May, 2018.
47. Cybenko, *"Machine Learning and Security" CIA Summer Symposium, Langley, Virginia, June 2018.*

PEER-REVIEWED CONFERENCE PROCEEDING PUBLICATIONS

1. **DOI:** <https://doi.org/10.1109/CNS.2014.6997517>
Title: Manipulating the Attacker's View of a System's Attack Surface
Authors: M. Albanese, E. Battista, S. Jajodia, and V. Casola
Conference Title: 2nd IEEE Conference on Communications and Network Security (IEEE CNS 2014)
Location: San Francisco, California, USA
Conference Dates: October 29-31, 2014

2. **DOI:** https://doi.org/10.1007/978-3-319-11212-1_28
Title: Modeling Network Diversity for Evaluating the Robustness of Networks against Zero-Day Attacks
Authors: L. Wang, M. Zhang, S. Jajodia, A. Singhal, and M. Albanese
Conference Title: 19th European Symposium on Research in Computer Security (ESORICS 2014)
Location: Wroclaw, Poland
Conference Dates: September 7-11, 2014

3. **DOI:** <https://doi.org/10.5220/0005013800190030>
Title: Keeping Intruders at Large: A Graph-Theoretic Approach to Reducing the Probability of Successful Network Intrusions
Authors: P. Shakarian, D. Paulo, M. Albanese, and S. Jajodia
Conference Title: 11th International Conference on Security and Cryptography (SECRYPT 2014)
Location: Vienna, Austria
Conference Dates: August 28-30, 2014

4. **DOI:** N/A
Title: Three tenets for secure cyber-physical system design and assessment
Authors: J. Hughes and G. Cybenko
Conference Title: SPIE Cyber Sensing
Conference Dates: May 2014

5. **DOI:** N/A
Title: Comparing Different Moving Target Defense Techniques
Authors: J. Xu, P. Guo, M. Zhao, R. F. Erbacher, M. Zhu, P. Liu
Conference Title: ACM MTD Workshop 2014
Conference Dates: Nov. 2014

6. **DOI:** N/A
Title: Reinforcement learning algorithms for adaptive cyber defense against Heartbleed
Authors: Minghui Zhu, Zhisheng Hu and Peng Liu
Conference Title: ACM MTD Workshop 2014
Conference Dates: Nov. 2014

7. **DOI:** N/A
Title: Achieving Accuracy and Scalability Simultaneously in Detecting Application Clones on Android Markets

Authors: Kai Chen, Peng Liu, Yingjun Zhang

Conference Title: IEEE International Conference on Software Engineering (ICSE 2014)

Conference Dates: June 2014

8. **DOI:** N/A

Title: Detangling Resource Management Functions from the TCB in Privacy-Preserving Virtualization

Authors: Min Li, Zili Zha, Wanyu Zang, Meng Yu, Peng Liu, Kun Bai

Conference Title: ESORICS 2014

Conference Dates: Sept. 7-11, 2014

9. **DOI:** N/A

Title: Inferring the Stealthy Bridges between Enterprise Network Islands in Cloud Using Cross-Layer Bayesian Networks

Authors: Xiaoyan Sun, Jun Dai, Anoop Singhal, Peng Liu

Conference Title: SecureComm 2014

Location: Beijing, China

Conference Dates: Sept. 23-26, 2014

10. **DOI:** N/A

Title: System Call Redirection: A Practical Approach to Meeting Real-world VMI Needs

Authors: R. Wu, P. Chen, P. Liu, B. Mao

Conference Title: DSN 2014

Conference Dates: June 2014

11. **DOI:** N/A

Title: RootkitDet: Practical End-to-End Defense against Kernel Rootkits in a Cloud Environment

Authors: Lingchen Zhang, Sachin Shetty, Peng Liu, Jiwu Jing

Conference Title: ESORICS 2014

Conference Dates: Sept. 7-11, 2014

12. **DOI:**

Title: Duet: Library Integrity Verification for Android Applications

Authors: Wenhui Hu, Damien Ocateau, Patrick McDaniel, and Peng Liu

Conference Title: ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec)

Location: Oxford, United Kingdom

Conference Dates: July 2014

13. **DOI:**

Title: ViewDroid: Towards Obfuscation-Resilient Mobile Application Repackaging Detection,

Authors: Fangfang Zhang, Heqing Huang, Sencun Zhu, Dinghao Wu and Peng Liu,

Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

14. **DOI:**

Title: Program Logic Based Software Plagiarism Detection,

Authors: Fangfang, Zhang, Dinghao Wu, Peng Liu, and Sencun Zhu,

Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

15. DOI:

Title: Simultaneous input and state smoothing for linear discrete-time stochastic systems with unknown inputs,

Authors: Sze Zheng Yong, Minghui Zhu and Emilio Frazzoli,

Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

16. DOI:

Title: Generalized innovation and inference algorithms for hidden mode switched linear stochastic systems with unknown inputs,

Authors: Sze Zheng Yong, Minghui Zhu and Emilio Frazzoli,

Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

17. DOI:

Title: Droidmarking: Resilient Software Watermarking for Impeding Android Application Repackaging,

Authors: Chuangang Ren, Kai Chen, Peng Liu,

Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

18. DOI:

Title: Computing solutions in infinite-horizon discounted adversarial patrolling games,

Authors: Y. Vorobeychik, B. An, M. Tambe, and S. Singh,

Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

19. DOI:

Title: A supervisory control approach to dynamic cyber-security,

Authors: M. Rasouli, E. Miehling, and D. Teneketzis,

Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

20. DOI:

Title: Empirical game-theoretic analysis of an adaptive cyber-defense scenario (Preliminary report),

Authors: M.P. Wellman and A. Prakash,

Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

21. DOI:

Title: Protecting enterprise networks through attack surface expansion,

Authors: Kun Sun, Sushil Jajodia,

Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and

Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

22. DOI:

Title: Adversarial and uncertain reasoning for adaptive cyber defense: Building the cyber foundation (invited paper),

Authors: George Cybenko, Sushil Jajodia, Michael P. Wellman, and Peng Liu,

Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

23. DOI:

Title: Quantification of moving target cyber defenses,

Authors: K. Farris and G. Cybenko,

Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

24. DOI:

Title: No free lunch in cyber security,

Authors: Cybenko, George and Jeff Hughes,

Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

25. DOI:

Title: Protecting enterprise networks through attack surface expansion (short paper),

Authors: Kun Sun and Sushil Jajodia,

Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

26. DOI:

Title: Empirical game-theoretic analysis of an adaptive cyber-defense scenario (Preliminary report),

Authors: A. Prakash and M. P. Wellman,

Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

27. DOI:

Title: A Supervisory Control Approach to Dynamic Cyber-Security,

Authors: M. Rasouli, E. Miehling, and D. Teneketzis,

Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

28. DOI:

Title: Optimal Defense Policies for Partially Observable Spreading Processes on Bayesian Attack Graphs,

Authors: E. Miehling, M. Rasouli, and D. Teneketzis,

Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

29. DOI:
Title: Empirical game-theoretic analysis for moving target defense,
Authors: A. Prakash and M. P. Wellman,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.
30. DOI:
Title: A Deception Based Approach for Defeating OS and Service Fingerprinting,
Authors: M. Albanese, E. Battista, and S. Jajodia,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.
31. DOI:
Title: Disrupting Stealthy Botnets through Strategic Placement of Detectors,
Authors: S. Venkatesan, M. Albanese, and S. Jajodia.
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.
32. DOI:
Title: A Practical Approach for Adaptive Data Structure Layout Randomization,
Authors: P. Chen, J. Xu, Z. Lin, D. Xu, B. Mao and P. Liu,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.
33. DOI:
Title: Replacement Attacks: Automatically Impeding Behavior-based Malware Specifications,
Authors: J. Ming, Z. Xin, P. Lan, D. Wu, P. Liu, and B. Mao,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.
34. DOI:
Title: HeapTherapy: An Efficient End-to-end Solution against Heap Buffer Overflows,
Authors: Q. Zeng, M. Zhao, and P. Liu,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.
35. DOI:
Title: Risk Assessment of Buffer 'Heartbleed' Over-read Vulnerabilities
Authors: J. Wang, M. Zhao, Q. Zeng, D. Wu, and P. Liu,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.
36. DOI:
Title: Between Mutual Trust and Mutual Distrust: Practical Fine-grained Privilege Separation in Multithreaded Applications,
Authors: Jun Wang, Xi Xiong, and Peng Liu,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and

Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

37. DOI:

Title: Towards Discovering and Understanding Task Hijacking in Android,
Authors: C. Ren, Y. Zhang, H. Xue, T. Wei, and P. Liu,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

38. DOI:

Title: Private Browsing Mode Not Really That Private: Dealing with Privacy Breach Caused by Browser Extensions,
Authors: B. Zhao and P. Liu,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

39. DOI:

Title: An Empirical Study of Web Vulnerability Discovery Ecosystems,
Authors: M. Zhao, J. Grossklags, and P. Liu,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

40. DOI:

Title: From System Services Freezing to the System Server Shutdown in Android: All You Need Is a Loop in an Application,
Authors: H. Huang, S. Zhu, P. Liu,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

41. DOI:

Title: How Your Phone Camera Can Be Used to Stealthily Spy on You: Transplantation Attacks against Android Camera Service,
Authors: Z. Zhang, P. Liu, J. Xiang, J. Jing, and L. Lei,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

42. DOI:

Title: Simultaneous input and state estimation for linear time-invariant continuous-time stochastic systems with unknown inputs,
Authors: S. Z. Yong, M. Zhu, and E. Frazzoli,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

43. DOI:

Title: Simultaneous input and state estimation with a delay,
Authors: S. Z. Yong, M. Zhu, and E. Frazzoli,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

44. DOI:
Title: Simultaneous input and state estimation of linear discrete-time stochastic systems with input aggregate information,
Authors: S. Z. Yong, M. Zhu, and E. Frazzoli,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.
45. DOI:
Title: Resilient state estimation against switching attacks on stochastic cyber-physical systems,
Authors: S. Z. Yong, M. Zhu, and E. Frazzoli,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.
46. DOI:
Title: A Framework for Moving Target Defense Quantification,
Authors: W. Connell, M. Albanese, and S. Venkatesan,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.
47. DOI:
Title: Moving Target Defense against DDoS Attacks: An Empirical Game-Theoretic Analysis,
Authors: M. Wright, S. Venkatesan, M. Albanese, and M.P. Wellman,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.
48. DOI:
Title: A Moving Target Defense Approach to Disrupting Stealthy Botnets,
Authors: S. Venkatesan, M. Albanese, G. Cybenko, and S. Jajodia,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.
49. DOI:
Title: A Quantitative Risk Assessment Framework for Adaptive Intrusion Detection in the Cloud,
Authors: L.H. Pham, M. Albanese, and S. Venkatesan,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.
50. DOI:
Title: A Moving Target Defense Approach to Mitigate DDoS Attacks against Proxy-Based Architectures,
Authors: S. Venkatesan, M. Albanese, K. Amin, S. Jajodia, and M. Wright,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.
51. DOI:
Title: Securing networks against unpatchable and unknown vulnerabilities using heterogeneous hardening options,
Authors: Daniel Borbor, Lingyu Wang, Sushil Jajodia, Anoop Singhal,

Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

52. DOI:

Title: Vulnerability survival analysis: a novel approach to vulnerability management,
Authors: Katheryn A. Farris, John Sullivan, and George Cybenko,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

53. DOI:

Title: Approximating centrality in evolving graphs: toward sublinearity,
Authors: Benjamin W. Priest and George Cybenko,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

54. DOI:

Title: Parallel Computing for Machine Learning in Social Network Analysis,
Authors: George Cybenko,
Conference Title: in Proc. of the ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), July 2014. Oxford, United Kingdom.

55. DOI:

Title:
Authors: Z. Hu, P. Chen, Y. Lu, M. Zhu and P. Liu. Towards a science for adaptive defense: Revisit server protection. IEEE International Conference on Collaboration and Internet Computing, pages 112-121, Pittsburgh, November 2016.
Conference Title:

56. DOI:

Title:
Authors: H. Kim, P. Guo, M. Zhu and P. Liu. On attack-resilient estimation of switched nonlinear cyber-physical systems. American Control Conference, Seattle, pages 4328-4333, May 2017.
Conference Title:

57. DOI:

Title:
Authors: P. Chen, J. Xu, Z. Hu, X. Xing, M. Zhu, B. Mao and P. Liu. What you see is not what you get! Thwarting Just-in-time ROP with Chameleon. 2017 IEEE/IFIP International Conference on Dependable Systems and Networks, June 2017.
Conference Title:

58. DOI:

Title:
Authors: P. Guo, H. Kim, L. Guan, M. Zhu and P. Liu. VCIDS: Collaborative intrusion detection of actuator and sensor attacks on connected vehicles. 13th International Conference on Security and Privacy in Communication Networks, Oct 2017. Accepted.
Conference Title:

59. DOI:
Title: VaultIME: Regaining User Control For Password Manager through Auto-correction,
Authors: L. Guan, S. Farhang, Y. Pu, P. Guo, J. Grossklags, P. Liu,
Conference Title: to appear in Proceedings of SecureComm 2017.
60. DOI:
Title: POMP: Postmortem Program Analysis with Hardware-Enhanced Post-Crash Artifacts,
Authors: Xu, J., Mu, D., Xing, X., Liu, P., Chen, P., Mao, B.,
Conference Title: in Proceedings of the 26th USENIX Security Symposium (USENIX Security),
Vancouver, Canada, August 2017.
61. DOI:
Title: Towards Actionable Mission Impact Assessment in the Context of Cloud computing,
Authors: Xiaoyan Sun, Anoop Singhal and Peng Liu,
Conference Title: DBSEC 2017.
62. DOI:
Title: TrustShadow: Secure execution of unmodified applications with ARM TrustZone,
Authors: L. Guan, P. Liu, X. Xing, X. Ge, S. Zhang, M. Yu, and T. Jaeger,
Conference Title: ACM MobiSys 2017.
63. DOI:
Title: System Service Call-oriented Symbolic Execution of Android Framework with Applications
to Vulnerability Discovery and Exploit Generation,
Authors: Lannan Luo, Qiang Zeng, Chen Cao, Kai Chen, Jian Liu, Limin Liu, Neng Gao, Min Yang,
Xinyu Xing, and Peng Liu,
Conference Title: ACM MobiSys 2017.
64. DOI:
Title: WindowGuard: Systematic Protection of GUI Security in Android,
Authors: Chuangang Ren, Peng Liu, Sencun Zhu,
Conference Title: NDSS 2017.
65. DOI:
Title: Dancing with Wolves: Towards Practical Event-driven VMM Monitoring,
Authors: D. Liang, P. Liu, J. Xu, P. Chen, Q. Zeng,
Conference Title: VEE 2017.
66. DOI:
Title: Android Malware Development on Public Malware Scanning Platforms: A Large-scale Data-
driven Study,
Authors: Heqing Huang, Cong Zheng, Junyuan Zeng, Wu Zhou, Sencun Zhu, Peng Liu, Suresh
Chari, Ce Zhang,
Conference Title: IEEE BigData 2016.
67. DOI:
Title: Sanitizing Data Is Not Enough! Towards Sanitizing Structural Artifacts in Flash Media,
Authors: Bo Chen, Shijie Jia, Luning Xia, and Peng Liu,

Conference Title: ACSAC, 2016.

68. DOI:

Title: From Physical to Cyber: Escalating Protection for Personalized Auto Insurance,
Authors: Guan, L., Xu, J., Wang, S., Xing, X., Lin, L., Huang, H., Liu, P., Lee, W.,
Conference Title: in Proceedings of the 14th ACM Conference on Embedded Networked Sensor Systems (SenSys 2016), Palo Alto, USA, December 2016.

69. DOI:

Title: StraightTaint: Decoupled Offline Symbolic Taint Analysis,
Authors: Jiang Ming, Dinghao Wu, Gaoyao Xiao, Jun Wang, and Peng Liu,
Conference Title: in Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering (ASE 2016), Singapore, September 2016

70. DOI:

Title: CREDAL: Towards Locating a Memory Corruption Vulnerability with Your Core Dump,
Authors: Xu, J., Mu, D., Chen, P., Xing, X., Liu, P.,
Conference Title: ACM CCS 2016.

71. DOI:

Title: Call Me Back! Attacks on System Server and System Apps in Android through Synchronous Callback,
Authors: Kai Wang, Yuqing Zhang, Peng Liu,
Conference Title: ACM CCS 2016.

72. DOI:

Title: A Machine Learning Approach for Detecting Third-Party Trackers on the Web,
Authors: Qianru Wu, Qixu Liu, Yuqing Zhang, Peng Liu and Guanxing Wen,
Conference Title: ESORICS 2016.

73. DOI:

Title: Stay in Your Cage! A Sound Sandbox for Third-Party Libraries on Android,
Authors: Fabo Wang, Yuqing Zhang, Kai Wang, Peng Liu and Wenjie Wang,
Conference Title: ESORICS 2016.

74. DOI:

Title: Towards Probabilistic Identification of Zero-day Attack Paths,
Authors: Xiaoyan Sun, Jun Dai, Peng Liu, Singhal, Anoop, John Yen,
Conference Title: IEEE CNS 2016.

75. DOI:

Title: On stochastic dynamic games with delayed sharing information structure,
Authors: H. Tavafoghi, Y. Ouyang, and D. Teneketzis,
Conference Title: in Proceedings of the 5th IEEE Conference on Decision and Control (CDC 2016), December 2016.

76. DOI:

Title: A Stackelberg Game Model for Botnet Traffic Exfiltration,

Authors: T.H. Nguyen, M.P. Wellman, and S. Singh,
Conference Title: in Proceedings of AAAI-17 Workshop on Artificial Intelligence for Cyber Security, February 2017.

77. **DOI:** 10.5220/0006872602820293

Title: A Quantitative Framework to Model Advanced Persistent Threats [Best Paper Award]

Authors: Luan Huy Pham, Massimiliano Albanese, and Benjamin W. Priest

Publication Title: Proceedings of the 15th International Conference on Security and Cryptography (SECRYPT 2018)

Month: July

Year: 2018

78. **DOI:** 10.1145/3140549.3140550

Title: Performance Modeling of Moving Target Defenses

Authors: Warren Connell, Daniel A. Menasce, and Massimiliano Albanese

Publication Title: Proceedings of the 4th ACM Workshop on Moving Target Defense (MTD 2017)

Month: October

Year: 2017

79. **DOI:** 10.1145/3140549.3140552

Title: Detecting Stealthy Botnets in a Resource-Constrained Environment using Reinforcement Learning

Authors: Sridhar Venkatesan, Massimiliano Albanese, Ankit Shah, Rajesh Ganesan, and Sushil Jajodia

Publication Title: Proceedings of the 4th ACM Workshop on Moving Target Defense (MTD 2017)

Month: October

Year: 2017

80. **DOI:** 10.1145/3140549.3140562

Title: Multi-Stage Attack Graph Security Games: Heuristic Strategies, with Empirical Game-Theoretic Analysis

Authors: T.H. Nguyen, M. Wright, M.P. Wellman, and S. Singh

Publication Title: Proceedings of the 4th ACM Workshop on Moving Target Defense (MTD 2017)

Month: October

Year: 2017

81. **DOI:** 10.1007/978-3-319-68711-7_9

Title: A Stackelberg game model for botnet data exfiltration

Authors: T.H. Nguyen, M.P. Wellman, and S. Singh

Publication Title: Proceedings of the 8th Conference on Decision and Game Theory for Security (GameSec 2017)

Month: October

Year: 2017

82. **DOI:** 10.1109/EuroSP.2018.00035

Title: SoK: Security and Privacy in Machine Learning

Authors: N. Papernot, P. McDaniel, A. Sinha, and M.P. Wellman

Publication Title: Proceedings of the 3rd IEEE European Symposium on Security and Privacy

Month: April

Year: 2018

83. **DOI:** 10.24963/ijcai.2018/775

Title: Stackelberg Security Games: Looking Beyond a Decade of Success

Authors: A. Sinha, F. Fang, B. An, C. Kiekintveld, and M. Tambe

Publication Title: Proceedings of the 27th International Joint Conference on Artificial Intelligence

Month: July

Year: 2018

84. **DOI:** 10.1145/3140549.3140556

Title: Online algorithms for adaptive cyber defense on Bayesian attack graphs

Authors: Zhisheng Hu, Minghui Zhu, and Peng Liu

Publication Title: Proceedings of the 4th ACM Workshop on Moving Target Defense

Month: October

Year: 2017

85. **DOI:** 10.1145/3133956.3134035

Title: FlashGuard: Leveraging Intrinsic Flash Properties to Defend Against Encryption Ransomware

Authors: Jian Huang, Jun Xu, Xinyu Xing, Peng Liu, and Moinuddin K. Qureshi

Publication Title: ACM CCS 2017

Month: October

Year: 2017

86. **DOI:** 10.1145/3133956.3134011

Title: DEFTL: Implementing Plausibly Deniable Encryption in Flash Translation Layer

Authors: S. Jia, L. Xia, B. Chen, and P. Liu

Publication Title: ACM CCS 2017

Month: October

Year: 2017

87. **DOI:** 10.1145/3134600.3134647

Title: Supporting Transparent Snapshot for Bare-metal Malware Analysis on Mobile Devices

Authors: Le Guan, Shijie Jia, Bo Chen, Fengwei Zhang, Bo Luo, Jingqiang Lin, Peng Liu, Xinyu Xing, and Luning Xia

Publication Title: ACSAC 2017

Month: December

Year: 2017

88. **DOI:** 10.1109/DSN.2018.00065

Title: RoboADS: Anomaly detection against sensor and actuator misbehaviors in mobile robots

Authors: P. Guo, H. Kim, N. Virani, J. Xu, M. Zhu, and P. Liu

Publication Title: DSN 2018

Month: June

Year: 2018

89. **DOI:** 10.1145/3212480.3212503
Title: Android STAR: Interaction-Preserving Messenger-Usage Inspection
Authors: L. Yuan, P. Liu, and S. Zhu
Publication Title: Proceedings of the 11th ACM Conference on Security & Privacy in Wireless and Mobile Networks (ACM WiSec 2018)
Month: June
Year: 2018
-
90. **DOI:**
Title: Estimating Edge-Local Triangle Count Heavy Hitters in Edge-Linear Time and Almost-Vertex-Linear Space
Authors: Benjamin W. Priest, Roger Pearce, and Geoffrey Sanders
Publication Title: IEEE High Performance Extreme Computing (HPEC) Graph Challenge
Month: September
Year: 2018
-
91. **DOI:** <https://doi.org/10.1145/3328526.3329634>
Title: Iterated Deep Reinforcement Learning in Games: History-Aware Training for Improved Stability
Authors: M Wright, Y Wang, MP Wellman
Conference Title: 20th ACM Conference on Economics and Computation
Location: Phoenix, AZ, USA
Conference Dates: June 2019
-
92. **DOI:** <https://doi.org/10.1109/CDC.2018.8619040>
Title: A Sufficient Information Approach to Decentralized Decision Making
Authors: H Tavafoghi, Y Ouyang, D Teneketzis
Conference Title: 57th IEEE Conference on Decision and Control
Location: Miami, FL, USA
Conference Dates: December 2018
-
93. **DOI:** <https://doi.org/10.1609/aaai.v33i01.33012133>
Title: Deception in finitely repeated security games
Authors: TH Nguyen, Y Wang, A Sinha, MP Wellman
Conference Title: 33rd AAAI Conference on Artificial Intelligence
Location: Honolulu, HI, USA
Conference Dates: January 2019
-
94. **DOI:**
Title: Probably almost stable strategy profiles in simulation-based games
Authors: M Wright, MP Wellman
Conference Title: AAMAS-19 Workshop on Games, Agents, and Incentives
Location: Montreal, Canada
Conference Dates: May 2019
-
- DOI:** https://doi.org/10.1007/978-3-030-01554-1_26
Title: A learning and masking approach to secure learning

Authors: L Nguyen, S Wang, A Sinha

Conference Title: 9th Conference on Decision and Game Theory for Security

Location: Seattle, WA, USA

Conference Dates: October 2018

95. **DOI:** <https://doi.org/10.1109/ISI.2018.8587326>

Title: Measuring the Effectiveness of Network Deception

Keywords: cyber deception, metrics

Abstract: Cyber reconnaissance is the process of gathering information about a target network for the purpose of compromising systems within that network. Network-based deception has emerged as a promising approach to disrupt attackers' reconnaissance efforts. However, limited work has been done so far on measuring the effectiveness of network-based deception. Furthermore, given that Software-Defined Networking (SDN) facilitates cyber deception by allowing network traffic to be modified and injected on-the-fly, understanding the effectiveness of employing different cyber deception strategies is critical. In this paper, we present a model to study the reconnaissance surface of a network and model the process of gathering information by attackers as interactions with a cyber defensive system that may use deception. To capture the evolution of the attackers' knowledge during reconnaissance, we design a belief system that is updated by using a Bayesian inference method. For the proposed model, we present two metrics based on KL-divergence to quantify the effectiveness of network deception. We tested the model and the two metrics by conducting experiments with a simulated attacker in an SDN-based deception system. The results of the experiments match our expectations, providing support for the model and proposed metrics.

Authors: Shridatt Sugrim and Sridhar Venkatesan and Jason Youzwak and Cho-Yu Chiang and Ritu Chadha and Massimiliano Albanese and Hasan Cam

Conference Title: IEEE International Conference on Intelligence and Security Informatics (IEEE ISI 2018)

Location: Miami, FL, USA

Conference Dates: November 2018

96. **DOI:**

Title: Towards Large-Scale Hunting for Android Negative-Day Malware

Keywords: Negative-Day Malware, Android, detection

Abstract: Android malware writers often utilize online malware scanners to check how well their malware can evade detection, and indeed we can find malware scan reports that were generated before the major outbreaks of such malware. If we could identify in-development malware before malware deployment, we would have developed effective defense mechanisms to prevent malware from causing devastating consequences. To this end, we propose Lshand to discover undiscovered malware before day zero, which we refer to as negative-day malware. The challenge includes scalability and the fact that malware writers would apply detection evasion techniques and submission anonymization techniques. Our approach is based on the observation that malware development is a continuous process and thus malware variants inevitably will share certain characteristics throughout its development process. Accordingly, Lshand clusters scan

reports based on selective features and then performs further analysis on those seemingly benign apps that share similarity with malware variants. We implemented and evaluated Lshand with submissions to VirusTotal. Our results show that Lshand is capable of hunting down undiscovered malware in a large scale, and our manual analysis and a third-party scanner have confirmed our negative-day malware findings to be malware or grayware.

Authors: Lun-Pin Yuan, Wenjun Hu, Ting Yu, Peng Liu, Sencun Zhu

Conference Title: RAID 2019

Location: Beijing, China

Conference Dates: Sept. 23-25, 2019

97. **DOI:**

Title: Discovering and Understanding the Security Hazards in the Interactions between IoT Devices, Mobile Apps, and Cloud on Smart Home Platforms

Keywords: vulnerability analysis, IoT platforms

Abstract: A smart home connects tens of home devices to the Internet, where an IoT cloud runs various home automation applications. While bringing unprecedented convenience and accessibility, it also introduces various security hazards to users. Prior research studied smart home security from several aspects. However, we found that the complexity of the interactions among the participating entities (i.e., devices, IoT clouds, and mobile apps) has not yet been systematically investigated. In this work, we conducted an in-depth analysis of five widely-used smart home platforms. Combining firmware analysis, network traffic interception, and blackbox testing, we reverse-engineered the details of the interactions among the participating entities. Based on the details, we inferred three legitimate state transition diagrams for the three entities, respectively. Using these state machines as a reference model, we identified a set of unexpected state transitions. To confirm and trigger the unexpected state transitions, we implemented a set of phantom devices to mimic a real device. By instructing the phantom devices to intervene in the normal entity-entity interactions, we have discovered several new vulnerabilities and a spectrum of attacks against real-world smart home platforms.

Authors: W. Zhou, Y. Jia, Y. Yao, L. Zhu, L. Guan, Y. Mao, P. Liu, Y. Zhang

Conference Title: USENIX Security Symposium 2019

Location: Santa Clara, CA, USA

Conference Dates: August 14-16, 2019

98. **DOI:** <https://doi.org/10.1145/3317549.3319721>

Title: Protecting Mobile Devices from Physical Memory Attacks with Targeted Encryption

Keywords: Physical memory attacks, memory encryption

Abstract: Sensitive data in a process could be scattered over the memory of a computer system for a prolonged period of time. Unfortunately, DRAM chips were proven insecure in previous studies. The problem becomes worse in the mobile environment, in which users' smartphones are easily lost or stolen. The powered-on phones may contain sensitive data in the vulnerable DRAM chips. In this paper, we propose MemVault, a mechanism to protect sensitive data in Android devices against physical memory attacks. MemVault keeps track of the propagation of well-marked sensitive data sources, and selectively encrypts tainted sensitive memory contents

in the DRAM chip. When a tainted object is accessed, MemVault redirects the access to the internal RAM (iRAM), where the cipher-text object is decrypted transparently. iRAM is a system-on-chip (SoC) component which is by nature immune to physical memory exploits. We have implemented a MemVault prototype system, and have evaluated it with extensive experiments. Our results validate that MemVault effectively eliminates the occurrences of clear-text sensitive objects in DRAM chips, and imposes acceptable overheads.

Authors: Le Guan, Chen Cao, Sencun Zhu, Jingqiang Lin, Peng Liu, Yubin Xia, Bo Luo

Conference Title: Proc. 12th ACM Conference on Security and Privacy in Wireless and Mobile Networks, WiSec'19

Location: Miami, FL, USA

Conference Dates: May 15-17, 2019

99. **DOI:** <https://doi.org/10.1145/3292006.3300041>

Title: MimosaFTL: Adding Secure and Practical Ransomware Defense Strategy to Flash Translation Layer

Keywords: ransomware, flash side defense

Abstract: Ransomware attacks have become prevalent nowadays due to sudden flourish of cryptocurrencies. Most existing defense strategies for ransomware, however, are vulnerable to privileged ransomware who can compromise the operating system and hence any backup data stored locally. The out-of-place-update and the isolation nature of flash memory storage, for the first time, makes it possible to design a defense strategy which is secure against the privileged ransomware. In this work, we propose MimosaFTL, a secure and practical ransomware defense strategy for mobile computing devices equipped with flash memory as external storage. MimosaFTL is secure against the privileged malware by taking advantage of unique characteristics of flash storage. In addition, it is more practical (compared to prior work) for real-world deployments by: 1) incorporating a fine-grained detection scheme which can detect presence of ransomware accurately; and 2) allowing the victim to efficiently restore the infected external storage to the exact point when the malware starts to perform corruption. Experimental evaluation shows that, MimosaFTL can mitigate ransomware attacks effectively with a small negative impact on both I/O performance and lifetime of flash storage.

Authors: Peiying Wang, Shijie Jia, Bo Chen, Luning Xia and Peng Liu

Conference Title: 9th ACM Conference on Data and Application Security and Privacy, CODASPY 2019

Location: Richardson, TX, USA

Conference Dates: Mar 25 - 27 2019

100. **DOI:** <https://doi.org/10.1145/3274694.3274733>

Title: Take It or Leave It: A Survey Study on Operating System Upgrade Practices

Keywords: operating system upgrade, survey study

Abstract: Software upgrades play a pivotal role in enhancing software performance, and are a critical component of resolving software bugs and patching security issues. However, consumers' eagerness to upgrade to the newest operating system is often tempered after release. In this paper, we focus on the upgrade perceptions and practices of users utilizing Microsoft Windows,

with particular consideration given to the current upgrade cycle to Windows 10, which was, for a time, offered at no monetary cost to many users. To better understand the relevant factors for upgrade decisions, we deployed a structured survey, including several open-ended questions to add additional depth. We collected data from 239 Microsoft Windows users and utilized qualitative and quantitative methods to analyze user upgrade practices. Important themes include how to best notify users of upcoming upgrade opportunities, how users perceive privacy issues associated with OS upgrade decisions, and whether security constitutes a significant decision-making factor. We also explore how end-of-life dates, indicating the end of support by the vendor, are perceived by users.

Authors: S. Farhang, J. Weidman, M. M. Kamani, J. Grossklags, P. Liu

Conference Title: Proc. of the 34th Annual Computer Security Applications Conference (ACSAC 2018)

Location: San Juan, Puerto Rico, USA

Conference Dates: December 03 - 07, 2018

101. **DOI:** https://doi.org/10.1007/978-3-030-00470-5_18

Title: CryptMe: Data Leakage Prevention for Unmodified Programs on ARM Devices

Keywords: data leakage prevention, memory encryption, ARM devices

Abstract: Sensitive data (e.g., passwords, health data and private videos) can be leaked due to many reasons, including (1) the misuse of legitimate operating system (OS) functions such as core dump, swap and hibernation, and (2) physical attacks to the DRAM chip such as cold-boot attacks and DMA attacks. While existing software-based memory encryption is effective in defeating physical attacks, none of them can prevent a legitimate OS function from accidentally leaking sensitive data in the memory. This paper introduces CryptMe that integrates memory encryption and ARM TrustZone-based memory access controls to protect sensitive data against both attacks. CryptMe essentially extends the Linux kernel with the ability to accommodate the execution of unmodified programs in an isolated execution domain (to defeat OS function misuse), and at the same time transparently encrypt sensitive data appeared in the DRAM chip (to defeat physical attacks). We have conducted extensive experiments on our prototype implementation. The evaluation results show the efficiency and added security of our design.

Authors: Chen Cao, Le Guan, Ning Zhang, Neng Gao, Jingqiang Lin, Bo Luo, Peng Liu, Ji Xiang, Wenjing Lou

Conference Title: RAID 2018

Location: Heraklion, Crete, Greece

Conference Dates: Sept. 10-12, 2018

102. **DOI:** <https://doi.org/10.1109/CNS48642.2020.9162298>

Title: A Quantitative Framework to Model Reconnaissance by Stealthy Attackers and Support Deception-Based Defenses

Keywords: advanced persistent threats, lateral movement, cyber deception, reconnaissance, adversarial modeling

Abstract: In recent years, persistent cyber adversaries have developed increasingly sophisticated techniques to evade detection. Once adversaries have established a foothold within the target

network, using seemingly-limited passive reconnaissance techniques, they can develop significant network reconnaissance capabilities. Cyber deception has been recognized as a critical capability to defend against such adversaries, but, without an accurate model of the adversary's reconnaissance behavior, current approaches are ineffective against advanced adversaries. To address this gap, we propose a novel model to capture how advanced, stealthy adversaries acquire knowledge about the target network and establish and expand their foothold within the system. This model quantifies the cost and reward, from the adversary's perspective, of compromising and maintaining control over target nodes. We evaluate our model through simulations in the CyberVAN testbed, and indicate how it can guide the development and deployment of future defensive capabilities, including high-interaction honeypots, so as to influence the behavior of adversaries and steer them away from critical resources.

Authors: Luan Huy Pham, Massimiliano Albanese, Ritu Chadha, Cho-Yu J. Chiang, Sridhar Venkatesan, Charles Kamhoua, and Nandi Leslie

Conference Title: 2020 IEEE Conference on Communications and Network Security (IEEE CNS 2020)

Location: Virtual due to COVID-19

Conference Dates: June 29 – July 1, 2020

103. DOI: <https://doi.org/10.1007/978-3-030-36537-0>

Title: Graph Models in Tracking Behaviors for Cyber-Security

Keywords: Control flow graphs, Electromagnetic emissions, Cyber security

Abstract: This book contains the key linear algebra and optimization techniques at the forefront of active data-science and machine learning practice today. This is an appropriate choice of content because while state-of-the-art machine learning applications can change each month (as in reinforcement learning, language translation, game playing, or image classification), the underlying mathematical concepts and algorithms do not.

Authors: George Cybenko

Conference Title: 6th International Workshop on Graphical Models for Security (GramSec 2019)

Location: Hoboken, New Jersey, USA

Conference Dates: November 2019

104. DOI: <https://doi.org/10.1109/SP40000.2020.00017>

Title: Unexpected Data Dependency Creation and Chaining: A New Attack to SDN

Keywords: Software-Defined Networking, D2C2 attack

Abstract: Software-Defined Networking (SDN) is an emerging network architecture that provides programmable networking through a logically centralized controller. As SDN becomes more prominent, its security vulnerabilities become more evident than ever. Serving as the "brain" of a software-defined network, how the control plane (of the network) is exposed to external inputs (i.e., data plane messages) is directly correlated with how secure the network is. Fortunately, due to some unique SDN design choices (e.g., control plane and data plane separation), attackers often struggle to find a reachable path to those vulnerable logic hidden deeply within the control plane. In this paper, we demonstrate that it is possible for a weak

adversary who only controls a commodity network device (host or switch) to attack previously unreachable control plane components by maliciously increasing reachability in the control plane. We introduce D2C2 (data dependency creation and chaining) attack, which leverages some widely-used SDN protocol features (e.g., custom fields) to create and chain unexpected data dependencies in order to achieve greater reachability. We have developed a novel tool, SVHunter, which can effectively identify D2C2 vulnerabilities. Till now we have evaluated SVHunter on three mainstream open-source SDN controllers (i.e., ONOS, Floodlight, and Opendaylight) as well as one security-enhanced controller (i.e., SE-Floodlight). SVHunter detects 18 previously unknown vulnerabilities, all of which can be exploited remotely to launch serious attacks such as executing arbitrary commands, exfiltrating confidential files, and crashing SDN services.

Authors: Feng Xiao, Jinqun Zhang, Jianwei Huang, Guofei Gu, Dinghao Wu, Peng Liu

Conference Title: 41st IEEE Symposium on Security and Privacy (IEEE S&P 2020)

Location: San Francisco, CA

Conference Dates: May 18-20, 2020

105. **DOI:** 10.1007/978-3-030-49669-2_1

Title: MODELING AND MITIGATING SECURITY THREATS IN NETWORK FUNCTIONS VIRTUALIZATION (NFV)

Keywords: Network functions virtualization, security threats

Abstract: By virtualizing proprietary hardware networking devices, Network Functions Virtualization (NFV) allows agile and cost-effective deployment of diverse network services for multiple tenants on top of the same physical infrastructure. As NFV relies on virtualization, and as an NFV stack typically involves several levels of abstraction and multiple co-resident tenants, this new technology also unavoidably leads to new security threats. In this paper, we take the first step toward modeling and mitigating security threats unique to NFV. Specifically, we model both cross-layer and co-residency attacks on the NFV stack. Additionally, we mitigate such threats through optimizing the virtual machine (VM) placement with respect to given constraints. The simulation results demonstrate the effectiveness of our solution.

Authors: Nawaf Alhebaishi, Lingyu Wang, Sushil Jajodia

Conference Title: IFIP WG 11.3 Annual Conf on Data and Applications Security and Privacy (DBSEC 2020)

Location: Virtual**Conference Dates:** June 25-27, 2020

BOOK CHAPTERS

1. M. Albanese, H. Cam, and S. Jajodia, “**Automated Cyber Situation Awareness Tools for Improving Analyst Performance,**” to appear in *Cybersecurity Systems for Human Cognition Augmentation*, Springer, 2014.
2. M. Zhu and S. Martinez, “**Distributed resilient control of operator-vehicle networks under cyber attacks,**” To appear in *Cyber Security for Industrial Control Systems: From the Viewpoint of Close-loop*, CRC Press, P. Cheng, H. Zhang, and J. Chen (Eds.), 2015.

3. George Cybenko and Patrick Sweeney, “**Identifying and Exploiting the Cyber High Ground with Botnets,**” in *Cyber Warfare: Building the Scientific Foundation*, Sushil Jajodia, Paulo Shakarian, V.S. Subrahmanian, Vipin Swarup, and Cliff Wang (Eds.), Springer 2015.
4. M. Rasouli, E. Miehl, and D. Teneketzis, “**A Scalable Decomposition Method for the Dynamic Defense of Cyber Networks**”, to appear in *Game Theory for Security Risk Management - From Theory to Practice*, Springer 2017.
5. G. Cybenko, “**Metrics for Cyber Resilience**”, to appear in *Cyber Resilience*, A. Kott (Ed.), 2017.
6. **Title:** From Cyber Situational Awareness to Adaptive Cyber Defense: Leveling the Cyber Playing Field

Keywords: In the cyber security landscape, the asymmetric relationship between defender and attacker tends to favor the attacker: while the defender needs to protect a system against all possible ways of breaching it, the attacker needs to identify and exploit only one vulnerable entry point in order to succeed. In this chapter, we show how we can effectively reverse such intrinsic asymmetry in favor of the defender by concurrently pursuing two complementary objectives: increasing the defender's understanding of multiple facets of the cyber landscape -- referred to as Cyber Situational Awareness (CSA) -- and creating uncertainty for the attacker through Moving Target Defense (MTD) or Adaptive Cyber Defense (ACD) techniques. This chapter provides a brief overview of contributions in these areas, and discusses future research directions.

Authors: Massimiliano Albanese

Publication Title: Data and Application Security and Privacy: Status and Prospects - Essays Dedicated to Sushil Jajodia on the Occasion of His 70th Birthday

Month: n/a

Year: 2018

7. **DOI:**

Title: Attack Graph Games and their Asymptotic Equilibria

Keywords: Attack graphs, game theory, asymptotic solutions

Abstract: This paper presents and analyzes an attack graph optimization problem that arises in modeling certain adversarial cyber-attack and -defend scenarios. The problem formulation is based on representing attacks against a system as a finite, weighted, directed graph in which the directed edges represent transitions between states in an attack and edge weights represent the estimated cost to an attacker for traversing the edge. An attacker strives to traverse the graph from a specified start node to a specified end node using the least weight cost directed path between those nodes. On the other hand, the defender seeks to allocate defensive measures in such a way as to maximize the attacker's minimal cost attack path. We study the role that minimal cut sets play in hardening the attack graph and prove that under this simple model minimal cut sets are optimal defensive investments in the limit even though minimal cut sets may not play a role in hardening a system initially. Viewing attackers and defenders as players in a two person, non-zero sum game, the results in this paper describe the asymptotic behavior of optimal solutions to the game under certain conditions.

Authors: George Cybenko and Gabriel Stocco

Publication Title: From Database to Cyber Security

Month: n/a

Year: 2018

8. **DOI:** https://doi.org/10.1007/978-3-319-75268-6_4

Title: A Scalable Decomposition Method for the Dynamic Defense of Cyber Networks

Keywords: Cyber security, minmax control information state, dynamic programming, influence graphs, scalability

Abstract: We investigate the problem of defending a cyber network against progressive attacks from an adversary. The defender is unable to perfectly observe attacks and the network's security status and instead must use its imperfect observations to determine a defense strategy. The nature of the defender's imperfect information is assumed to be non-probabilistic. Thus, the defender takes a conservative (minmax) approach to defending the network, attempting to construct a defense policy that minimizes the worst-case damage. Determining an optimal minmax defense strategy proves to be computationally intractable even for small-scale networks. To address this dimensionality issue, we propose a scalable decomposition method which involves the construction of multiple local defense problems, each equipped with a corresponding local defense policy. The local defense policies communicate information with one another with the goal of achieving network-wide security. The local defense problem's construction is based on a decomposition of the network into clusters. For the decomposition, we use the notion of an influence graph to describe the dependencies among the security states of the network's nodes. These dependencies, along with the available computational capability, are used to determine clusters of nodes, with each cluster corresponding to a local defense problem. After clusters are specified, we design the information structure of the network, that is, the information each local defense problem has over time to defend its own cluster; this information includes the data the local defense problem gathers from the environment along with the data communicated by other local defense policies. We illustrate the decomposition methodology with an example.

Authors: M. Rasouli, E. Miehling, and D. Teneketzis

Publication Title: Game Theory for Security and Risk Management

Month: May

Year: 2018

9. **DOI:** 10.1007/978-3-319-77492-3

Title: Quantifying Cyber Resilience: Metrics based on the system performance perspective

Abstract: Cyber resiliency has become an increasingly relevant and timely research and operational concept in cyber security. While several metrics have been proposed for quantifying cyber resiliency, a considerable gap remains between those metrics and operationally measurable and meaningful concepts that can be empirically determined in an objective manner. This work describes a concrete quantitative notion of cyber resiliency that can be tailored to meet specific needs of organizations that seek to introduce resiliency into their assessment of their cyber security posture.

Authors: George Cybenko

Publication Title: Cyber Resilience of Systems and Networks

Month: July

Year: 2018

-
10. DOI: https://doi.org/10.1142/9781786345646_011

Title: Large-scale Analogue Measurements and Analysis for Cyber-Security

Keywords: Security, RF, signal processing

Abstract: All digital logic and data in a computing system are subject to compromise by an attacker. As a result, any security monitoring or analysis done within the same communications and computational fabric is also subject to compromise. Moreover, many embedded processors do not have the additional memory or computational resources to coexist with security solutions. As a result, there is interest in using out-of-band analogue side-channel measurements and their analyses to monitor expected programme execution. This chapter describes an approach to this problem that involves measuring external electromagnetic emanations from a commodity processor and analysing those measurements to estimate various aspects of application software execution. It is of particular importance to alert the user to deviations from the logically constrained execution paths determined by the programme structure. Our approach uses analogue RF measurements made in multiple bands at GHz rates which generate enormous amounts of data at high sampling rates. These measurements are used initially for developing the appropriate classification and analytics, and subsequently for implementing those classifiers and trackers operationally in real time. We describe our implementation of machine learning and information theory-based algorithms for identifying the most useful features for tracking an application's control flow graph. Moreover, we describe a hierarchy of control flow graph models in which different models have different execution tracking performance properties.

Authors: George Cybenko and Gil M. Raz

Conference Title: Data Science for Cyber-Security, Security Science and Engineering 3

Month: November

Year: 2018

106. DOI: https://doi.org/10.1007/978-3-319-77492-3_2

Title: Metrics Based on the System Performance Perspective

Keywords: Security, metrics, resiliency

Abstract: This part of the book presents two alternative – but not incompatible – views on how to quantify cyber resilience via suitable metrics. This chapter – the first of the two – takes the perspective in which system performance is central to the metrics. As discussed in the introduction chapter of this book, cyber resiliency has become an increasingly important, relevant, and timely research and operational concept in cyber security. Although multiple metrics have been proposed for quantifying cyber resiliency, a connection remains to be made between those metrics and operationally measurable and meaningful resilience concepts that can be empirically determined in an objective manner. This chapter describes a concrete quantitative and measureable notion of cyber resiliency that can be tailored to meet specific needs of organizations seeking to introduce resiliency into their assessment of their cyber security posture.

Authors: George Cybenko

Conference Title: Cyber Resilience of Systems and Networks

Month: April

Year: 2019

DOI: https://doi.org/10.1007/978-3-030-04834-1_5

Title: Asymptotic Behavior of Attack Graph Games

Keywords: Security, graphs, asymptotic analysis

Abstract: This paper presents and analyzes an attack graph optimization problem that arises in modeling certain adversarial cyber attack and defend scenarios. The problem formulation is based on representing attacks against a system as a finite, weighted, directed graph in which the directed edges represent transitions between states in an attack and edge weights represent the estimated cost to an attacker for traversing the edge. An attacker strives to traverse the graph from a specified start node to a specified end node using the least weight cost directed path between those nodes. On the other hand, the defender seeks to allocate defensive measures in such a way as to maximize the attacker's minimal cost attack path. We study the role that minimal cut sets play in hardening the attack graph and prove that under this simple model minimal cut sets are optimal defensive investments in the limit even though minimal cut sets may not play a role in hardening a system initially. Viewing attackers and defenders as players in a two person, non-zero sum game, the results in this paper describe the asymptotic behavior of optimal solutions to the game under certain conditions.

Authors: George Cybenko and Gabriel Stocco

Conference Title: From Database to Cyber Security

Month: November

Year: 2018

DOI: https://doi.org/10.1007/978-3-030-22479-0_18

Title: CASFinder: Detecting Common Attack Surface

Keywords: Attack surface, vulnerability

Abstract: Code reusing is a common practice in software development due to its various benefits. Such a practice, however, may also cause large scale security issues since one vulnerability may appear in many different software due to cloned code fragments. The well known concept of relying on software diversity for security may also be compromised since seemingly different software may in fact share vulnerable code fragments. Although there exist efforts on detecting cloned code fragments, there lack solutions for formally characterizing their specific impact on security. In this paper, we revisit the concept of software diversity from a security viewpoint. Specifically, we define the novel concept of *common attack surface* to model the relative degree to which a pair of software may be sharing potentially vulnerable code fragments. To implement the concept, we develop an automated tool, *CASFinder*, in order to efficiently

identify common attack surface between any given pair of software with minimum human intervention. Finally, we conduct experiments by applying our tool to real world open source software applications. Our results demonstrate many seemingly unrelated software applications indeed share significant common attack surface.

Authors: Mengyuan Zhang, Yue Xin, Lingyu Wang, Sushil Jajodia, Anoop Singhal

Conference Title: *33rd Annual IFIP WG 11.3 Conference on Data and Applications Security (DBSEC 2019)*

Year: 2019

BOOKS

DOI: 10.1007/978-3-319-66505-4

Title: Network Security Metrics

Authors: Lingyu Wang, Sushil Jajodia, and Anoop Singhal

Year: 2017

Sushil Jajodia, Paulo Shakarian, V. S. Subrahmanian, Vipin Swarup, and Cliff Wang (Eds.), "[Cyber Warfare: Building the Scientific Foundation](#)," ISBN 978-3-319-1408-4, 340 pages, Springer, Berlin, 2015.

M. Zhu and S. Martinez, "**Distributed optimization-based control of multi-agent networks in complex environments**," Springer Briefs in Electrical and Computer Engineering (Control, Automation, and Robotics), Springer, 2015.

Sushil Jajodia and Kun Sun (Eds.), "**Proceedings of the First ACM Workshop on Moving Target Defense**," ACM, 2014.

EDITED BOOK

We produced an edited a book titled **Adaptive Autonomous Secure Cyber Systems** and edited by Sushil Jajodia, George Cybenko, V. S. Subrahmanian, Cliff Wang, and Michael Wellman which was published by Springer.

Autonomy in physical and cyber systems promises to revolutionize cyber operations. The ability of autonomous systems to execute at scales, scopes, and tempos exceeding those of humans and human-controlled systems will introduce entirely new types of cyber defense strategies and tactics, especially in highly contested physical and cyber environments. The development and automation of cyber strategies that are responsive to autonomous adversaries pose basic new technical challenges for cyber security.

The recent advances in adaptive cyber defense (ACD) have developed a range of new ACD techniques and methodologies for reasoning in an adaptive environment. Building on that work, this book explores fundamental scientific problems essential for autonomous cyber defense. Specific areas include:

- The extent to which autonomous cyber systems can be designed and operated in a framework that is significantly different from the human-based systems we now operate;

- Game and control theory-based moving target defenses (MTDs) and ACDs for fully autonomous cyber operations;
- On-line learning algorithms, including deep recurrent networks and reinforcement learning, for the kinds of situation awareness and decisions that autonomous cyber systems will require;
- Use machine learning (ML) and cyber deception methods and to reason about the situation and appropriate responses;
- Defending against attacks on ML-based autonomous cyber defensive agents.

DOI: <http://doi.org/10.1007/978-3-030-33432-1>

Book Title: Adaptive Autonomous Secure Cyber Systems

Editors: Sushil Jajodia, George Cybenko, V. S. Subrahmanian, Vipin Swarup, Cliff Wang, Michael Wellman

Publisher: Springer

Year: 2020

BOOK CHAPTERS

DOI: <https://doi.org/10.1002/9781119593386.ch5>

Book Title: Modeling and Design of Secure Internet of Things

Chapter Title: Adaptation and Deception in Adversarial Cyber Operations

Keywords: adversarial cyber operations, computational game theory, nontrivial games, real-world adversarial situations

Abstract: This chapter outlines a variety of challenges related to adaptation in adversarial operations. Such issues arise in autonomous systems that operate in hostile environments, where the adversaries can adapt and shape the environment as well. This includes national security systems like autonomous cyber operations, battlefield Internet of Things, unmanned vehicles, electronic warfare, and social media information operations. We note that recent progress in reinforcement learning has resulted in remarkable advances in game playing technology, to the extent that machines are now demonstrating superior play against human experts in games such as Go and some versions of Texas Hold 'Em poker. Those advances use a variety of combinations of computational game theory and deep reinforcement learning. However, there are several fundamental reasons that such successes do not translate immediately to more complex adversarial interactions such as arise in national security operations. This chapter reviews some of those reasons and how they relate to deceptions that are possible by opponents in adaptive, adversarial systems.

Authors: George Cybenko

Year: 2020

DOI: https://doi.org/10.1007/978-3-030-30719-6_1

Book Title: Adversarial and Uncertain Reasoning for Adaptive Cyber Defense: Control- and Game-Theoretic Approaches to Cyber Security

Chapter Title: Overview of control and game theory in adaptive cyber defenses

Keywords: Adaptive cyber defense, control theory, game theory

Abstract: The purpose of this chapter is to introduce cyber security researchers to key concepts in modern control and game theory that are relevant to Moving Target Defenses and Adaptive Cyber Defense. We begin by observing that there are fundamental differences between control models and game models that are important for security practitioners to understand. Those differences will be illustrated through simple but realistic cyber operations scenarios, especially with respect to the types and amounts of data required for modeling. In addition to modeling differences, there are a variety of ways to think about what constitutes a “solution.” Moreover, there are significant differences in the computational and information requirements to compute solutions for various types of Adaptive Cyber Defense problems. This material is presented in the context of the advances documented in this book, the various chapters of which describe advances made in the 2012 ARO ACD MURI.

Authors: G. Cybenko, M. Wellman, P. Liu and M. Zhu

Publisher: Springer

Year: 2019

DOI: https://doi.org/10.1007/978-3-030-30719-6_4

Book Title:

Chapter Title: Reinforcement learning for adaptive cyber defense against zero-day attacks

Keywords: Adaptive cyber defense, reinforcement learning

Abstract: In this chapter, we leverage reinforcement learning as a unified framework to design effective adaptive cyber defenses against zero-day attacks. Reinforcement learning is an integration of control theory and machine learning. A salient feature of reinforcement learning is that it does not require the defender to know critical information of zero-day attacks (e.g., their attack targets, and the locations of the vulnerabilities). This information is difficult, if not impossible, for the defender to gather in advance. The reinforcement learning based schemes are applied to defeat three classes of attacks: strategic attacks where the interactions between an attacker and a defender are modelled as a non-cooperative game; non-strategic random attacks where the attacker chooses its actions by following a predetermined probability distribution; and attacks depicted by Bayesian attack graphs where the attacker exploits combinations of multiple known or zero-day vulnerabilities to compromise machines in a network.

Authors: Z. Hu, P. Chen, M. Zhu and P. Liu

Publisher: Springer

Year: 2019

DOI: https://doi.org/10.1007/978-3-030-30719-6_7

Book Title: Adversarial and Uncertain Reasoning for Adaptive Cyber Defense: Control- and Game-Theoretic Approaches to Cyber Security

Chapter Title: MTD techniques for memory protection against zero-day attacks

Keywords: Moving target defense, memory protection

Abstract: During the past 25 years, the arms race between attacks exploiting memory corruption and memory protection techniques has drawn tremendous attention. This book chapter seeks to give an in-depth review of the newest research progress made on applying the MTD methodology to protect memory corruption exploits. The new research progress also represents the current phase of the arms race in the MTD perspective. In particular, on one hand, at the frontier of defending against control-hijacking attacks, we will give an in-depth review on the shift of defense strategy from static ASLR to dynamic ASLR. On the other hand, at the frontier of defending against data-oriented attacks, we will give an in-depth review on the shift of defense strategy from static DSLR to dynamic DSLR.

Authors: P. Chen, Z. Hu, J. Xu, M. Zhu, R. Erbacher, S. Jajodia and P. Liu

Publisher: Springer

Year: 2019

DOI: https://doi.org/10.1007/978-3-030-33432-1_4

Book Title: Adaptive Autonomous Secure Cyber Systems

Chapter Title: Can Cyber Operations Be Made Autonomous? An Answer from the Situational Awareness Viewpoint

Keywords: autonomous cyber systems, cyber operations

Abstract: Although many building blocks of today's cyber-defense solutions are already fully automatic, there is still a debate on whether next-generation cyber-defense solutions should be wholly autonomous. In this paper, we contribute to the debate in the context of Cybersecurity Operations Centers (CSOCs), which have been widely established in prominent companies and organizations to achieve cyber situational awareness. Based on the lessons we learned from a recent case study on making CSOC data

triage operations more autonomous, we conclude that instead of asking whether cyber operations can be made autonomous or not, it seems more appropriate to ask the following questions: (a) How to make cyber operations more autonomous? (b) What is the right research roadmap for making cyber operations more autonomous? We also comment on what should be the current frontier in building a significantly better CSOC.

Authors: C. Zhong, J. Yen, P. Liu

Editors: S. Jajodia, G. Cybenko, V. S. Subrahmanian, V. Swarup, C. Wang, and M. Wellman

Publisher: Springer

Year: 2020
