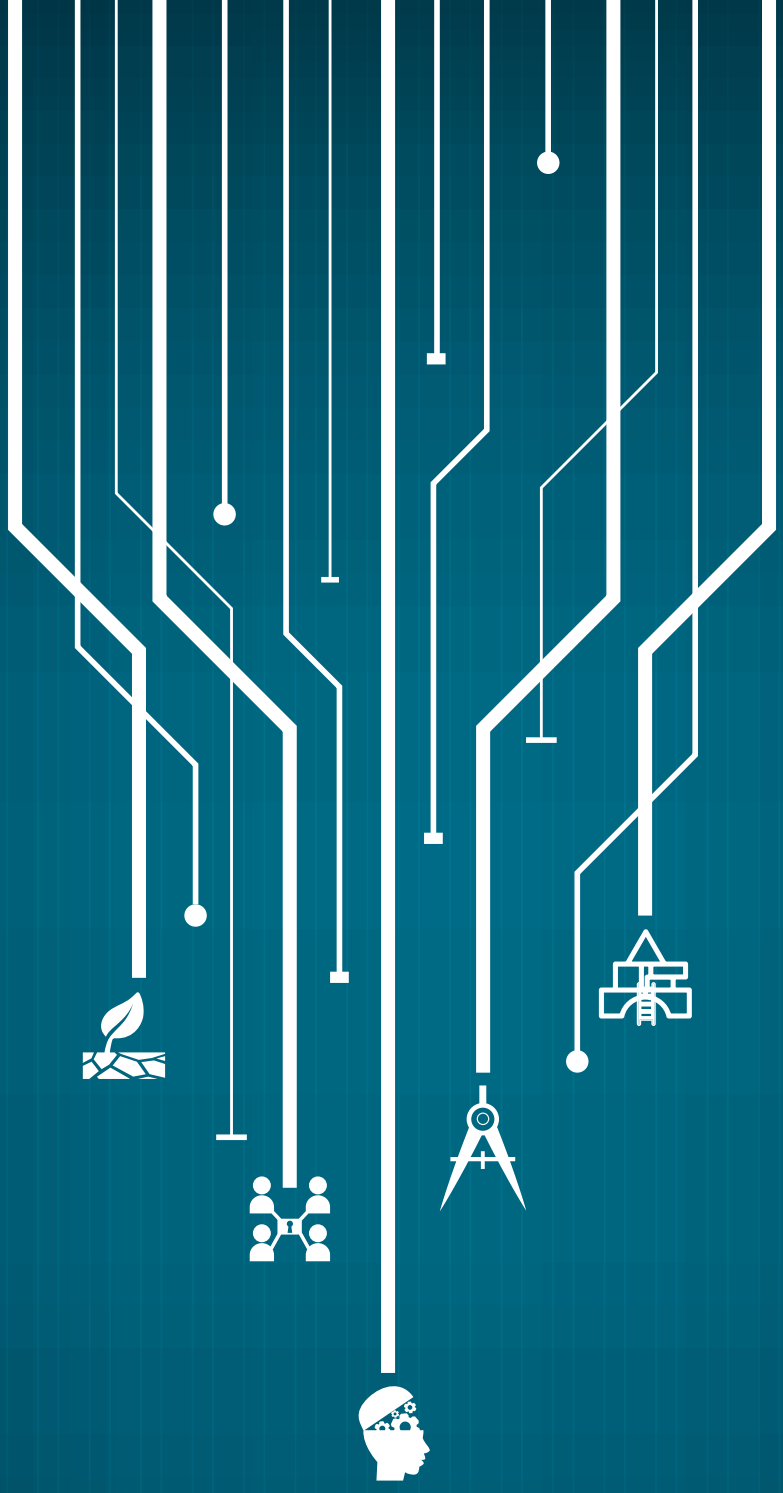


Science of Security and Privacy
2019 Annual Report

VISION

The National Security Agency Research Directorate sponsors the Science of Security and Privacy Initiative for the promotion of a foundational cybersecurity science that is needed to mature the cybersecurity discipline and to underpin advances in cyberdefense.



Science of Security and Privacy

Initiative 2019



Table of Contents

Executive Summary 5

Section 1: Engaging the Academic Community for Foundational Research 7

Hard Problems 8

Science of Security Lablets

Carnegie Mellon University 11
International Computer Science Institute 19
North Carolina State University 31
University of Illinois Urbana-Champaign 42
University of Kansas 53
Vanderbilt University 64

Science of Security Quarterly Meetings 73

Section 2: Promoting Rigorous Scientific Principles 82

Annual Best Cybersecurity Paper Competition 83
Intel International Science and Engineering Fair (ISEF) 86

Section 3: Growing the Science of Security 91

HotSoS 2019 92
SoS in the News 100
Outreach 102
CPS Summer Camp 104

Digital copies are available at https://cps-vo.org/group/sos/annualreport2019

Executive Summary

In 2019 the Science of Security and Privacy initiative (SoS) continued to contribute to the advancement of cybersecurity science through its support of research, commitment to scientific principles, and outreach aimed at growing the community. Under the sponsorship of the National Security Agency (NSA) Research Directorate (RES) whose mission is to secure the future by conducting ground-breaking research in a wide variety of science, technology, engineering, and mathematics areas, the SoS initiative advanced the goal of safeguarding interactions in cyberspace.

The SoS initiative is focused on producing scientifically supported cybersecurity advancement in the establishment of cybersecurity as a science. By replacing ad hoc and common practice approaches to security with scientifically supported best practice methods established through rigorous research, SoS is developing a strategic rather than tactical method of approaching cybersecurity. These strategic results are needed to transform cybersecurity from a cost-disadvantaged, reactionary field to one that is efficient and proactive. Established in 2011, the Science of Security fosters the establishment of security science through the pursuit of its three stated goals:

- Engage the academic community for foundational research
- Promote rigorous scientific principles
- Grow the SoS community

The SoS initiative engaged the academic community for foundational research through continued sponsorship of the third generation of SoS Lablets. The six Lablets focus on projects that address some of the most significant cybersecurity research challenges aligned against the five Hard Problems, the major focus areas identified in 2012 by NSA and the Lablets. The five Hard Problems are:

- Scalability and Composability
- Policy-Governed Secure Collaboration
- Security-Metrics-Driven Evaluation, Design, Development and Deployment
- Resilient Architectures
- Understanding and Accounting for Human Behavior

The six SoS Lablets are Carnegie Mellon University (CMU), the International Computer Science Institute (ICSI), North Carolina State University (NCSU), the University of Illinois Urbana-Champaign (UIUC), the University of Kansas (KU), and Vanderbilt University (VU). In addition to addressing the Hard Problems, ICSI focuses on privacy and VU focuses on Cyber-Physical Systems (CPS) research.

In 2019 NSA SoS leadership continued to promote better engagement between NSA and the Lablets, and the facilitation of NSA/Lablet collaboration and tech transfer arising from Lablet research by designating champions for each of its 28 projects. The projects are as follows:

Carnegie Mellon University:

- Characterizing User Behavior and Anticipating its Effects on Computer Security
- Model-Based Explanation for Human-in-the-Loop Security
- Obsidian: A Language for Secure-by-Construction Blockchain Programs
- Securing Safety-Critical Machine Learning Algorithms

International Computer Science Institute:

- Contextual Integrity for Computer Systems
- Designing for Privacy
- Governance for Big Data
- Operationalizing Contextual Integrity
- Scalable Privacy Analysis

North Carolina State University:

- Coordinated Machine Learning-Based Vulnerability and Security Patching for Resilient Virtual Computing Infrastructure
- Development of Methodology Guidelines for Security Research
- Predicting the Difficulty of Compromise through How Attackers Discover Vulnerabilities
- Principles of Secure Bootstrapping for IoT
- Reasoning about Accidental and Malicious Misuse via Formal Methods

University of Illinois at Urbana-Champaign:

- A Human Agent-Focused Approach to Security Modeling
- An Automated Synthesis Framework for Network Security and Resilience
- Monitoring, Fusion, and Response for Cyber Resilience
- Resilient Control of Cyber-Physical Systems with Distributed Learning
- Uncertainty in Security Analysis

University of Kansas:

- Cloud-Assisted IOT Systems Privacy
- Formal Approaches to the Ontology and Epistemology of Resilience
- Scalable Trust Semantics and Infrastructure
- Secure Native Binary Execution
- Side-Channel Attack Resilience

Vanderbilt University:

- Analytics for Cyber-Physical System Cybersecurity
- Foundations of Cyber-Physical System Resilience
- Mixed Initiative and Collaborative Learning in Adversarial Environments
- Multi-Model Test Bed for the Simulation-Based Evaluation of Resilience

Details about the Hard Problems and 2019 research on the specific projects can be found in Section 1.

While SoS sponsorship of fundamental research also contributes to the achievement of the second goal of promoting rigorous scientific principles, there are several other activities undertaken by SoS that reinforce that effort. The SoS initiative sponsored the 7th Annual Best Scientific Cybersecurity Paper Competition, again sponsored awards at the Intel International Science and Engineering Fair (ISEF), and sponsored an additional Best Paper Award at the Hot Topics in the Science of Security Symposium (HotSoS 2019). There were 37 papers submitted in the 7th Annual Paper Competition bringing the total submissions to over 260 during the competition's seven years. This year's winning paper was *Evaluating Fuzz Testing* by researchers from the University of Maryland and University of Texas at Dallas. This was the fifth consecutive year that SoS has sponsored prizes at ISEF, and this year there were 7 winners including Suha Hussain, who followed up her 2nd place finish in 2018 with a 1st place in 2019. Ms. Hussain's project investigated privacy violations of children's apps.

Details on SoS activities to promote rigorous scientific principles in 2019 can be found in Section 2.

The SoS initiative's support of foundational research through the Lablets and the promotion of rigorous scientific principles both serve to grow the Science of Security, but there are other activities that expand the Science of Security into other communities. HotSoS 2019 brought together close to 120 researchers and practitioners from academia, government, and industry for thought-provoking presentations on the cybersecurity science. The SoS-VO (www.sos-vo.org) grew to over 1700 members and continued as the centralized, online location for researchers and all interested parties to engage in discussion and access the most current research in cybersecurity. Both NSA SoS-sponsored activities and the topic of Science of Security, regularly appeared in social media. Associate Lablets were established to increase participation in the SoS community. SoS personnel attended numerous conferences to raise awareness of Science of Security, and the SoS initiative co-sponsored CPS Security Summer Camps at Vanderbilt University.

Details on what the SoS initiative did in 2019 to grow the Science of Security can be found in Section 3.

In 2020 the SoS initiative will continue to sponsor foundational research at the Lablets and seek to increase the impact of Lablet research on cybersecurity operations at NSA. The initiative will sponsor the 8th Annual Best Scientific Cybersecurity Paper Competition as well as the Intel ISEF awards. HotSoS 2020 will be held at the University of Kansas in April and continue its focus on the advancement of scientific methods in approaching the Hard Problems in cybersecurity. SoS personnel will continue to grow the Science of Security through outreach efforts at all academic levels to raise awareness of the need for foundational cybersecurity science to ensure a mature and reliable cyber defense.



Engaging the Academic Community for Foundational Research



In 2019 the Science of Security and Privacy (SoS) initiative continued its engagement of the academic community for foundational research primarily through its support of the six Science of Security Lablets: Carnegie Mellon University (CMU), the International Computer Science Institute (ICSI), North Carolina State University (NCSU), the University of Illinois Urbana-Champaign (UIUC), the University of Kansas (UK), and Vanderbilt University (VU). The specific research projects undertaken by the Lablets were selected by NSA to create a portfolio of projects that have technical excellence, NSA mission relevance, broad applicability beyond NSA, and in total would span the SoS five Hard Problems.

The six Lablets and their twelve Sub-Lablets performed research on 28 projects in 2019 and published 76 peer-reviewed papers bringing to well over 650 the number of papers published by Lablet researchers since the SoS initiative was established in 2012. The papers have addressed multiple aspects of the five Hard Problems, and have been presented at conferences, symposia, and workshops around the world. The foundational research embodied by the papers has contributed significantly to enhancing the scientific rigor of research into cybersecurity.

The Principal Investigators (PIs) of the Science of Security Lablets, along with the NSA Research organization, developed five Hard Problems. These Hard Problems serve as a means of establishing challenging and critical research goals, establishing a common language, and assessing progress in foundational SoS research. The papers published over the past year provide tangible evidence of the impact that the Lablets' research has had on improving the Science of Security in the five Hard Problem focus areas.

The Lablets also engage regularly in community outreach, participate in international conferences and workshops, and integrate Science of Security principles into their curricula. This year, for example, ICSI hosted the 2nd Symposium of Applications of Contextual Integrity which brought together researchers from computer and information science, communications, political science, and public health to discuss research using the theory of privacy as Contextual Integrity.

In addition to providing quarterly and annual reports on their activities, the Lablets, along with Sub-lablets, collaborators, and NSA researchers, meet quarterly to present updates on their research projects and exchange information about issues related to Science of Security.

The Hard Problems, Lablet projects and activities, and Lablet quarterly meetings are detailed in this section.

Science of Security

Hard Problems



At the very beginning of the Science of Security and Privacy initiative, the Principal Investigators (PIs) of the SoS Lablets, in collaboration with NSA Research, developed the Hard Problems as a means of establishing challenging and critical research goals. The Hard Problems also serve as the beginnings of a common language and a way to assess progress. These problems were selected for their level of technical challenge, their potential operational significance, and the likelihood that these problems would benefit from emphasis on scientific research methods and improved measurement capabilities. The five Hard Problems are not intended to cover all cybersecurity research challenges, but rather five specific areas that need scientific progress.

Hard problems are, by definition, elusive. The problem properties are not readily processed by traditional or well recognized modes of inquiry; the issue area or domain has not been subject to extensive analysis to date; the underlying dynamics reflect a daunting complexity; data-creation is a necessary but not sufficient condition for progress; system boundary may not be readily defined; and temporality may take on many forms--these are only some of the most daunting elements of Hard Problems.

The Hard Problems were designed to be crisply stated and well scoped in order to be able to assess progress towards solutions. Solutions may have the feature of incrementality in that discernible steps will lead towards an overall solution, each step with the potential to result in a corresponding increment of mission impact even when a fully comprehensive solution may remain challenging. Fundamental research undertaken by the Lablets is tied to at least one Hard Problem. See the individual Lablet project write-ups to learn the impact their research has had on the Hard Problems described below.



Resilient Architectures

Resilient Architectures includes the ability of a system to statically withstand attack, the ability of a system to continue to deliver essential services in the midst of an attack, and how quickly a system can be restored to full functionality following an attack. The Hard Problem focuses on designing, analyzing, and building systems that can: 1) withstand attack; 2) continue to deliver essential services (potentially at a diminished level) while under attack; and 3) quickly recover full functionality following an attack. The research goal is to develop the means to design and analyze system architectures that deliver required service in the face of compromised components.



Scalability and Composability

Scalability and Composability deals with the development and analysis of large-scale secure systems and the study of how to improve system security through security improvement of the components. The Hard Problem focuses on developing approaches for reasoning about software systems in a scalable way. The way to achieve scalability is via composability: reasoning approaches that allow us to analyze the

security properties of one component at a time, and then use the results of those analyses to reason about properties of the system as a whole. The research goal is to develop ways to construct systems and reason about system-level security properties using components with known security properties, without having to fully re-analyze the constituent components.



Policy-Governed Secure Collaboration

Policy-Governed Secure Collaboration aims to develop the science underlying methods to express and enforce normative requirements and policies for handling data with differing usage needs and among users in different authority domains. The Hard Problem is about developing the science that underlies methods for expressing and enforcing normative requirements and policies for information handling and privacy. Key challenges in policy are: 1) tackling differing uses, and differing expectations regarding uses, for the information; and 2) bridging across authority domains. The goal of the research is to develop a sociotechnical systems architecture that brings forth the interplay between social and technical elements of cybersecurity, including expressing and reasoning about norms and policies, computing interventions to achieve organizational needs, and predicting their complexity.



Security Metrics and Models

Security Metrics and Models addresses the measurement of properties relevant to cybersecurity, and quantifying the degree to which a system satisfies those properties. The Hard Problem involves techniques for effectively measuring and quantifying the extent to which a given system satisfies a particular set of security properties. Challenges include identifying the appropriate metrics for a given context, performing the measurement, analyzing the measurements and interpreting them with respect to a descriptive model, and understanding the degree of uncertainty which ought to accompany the measurements and their analysis**.** The goal of the research is to develop security metrics and models capable of predicting whether, or confirming that, a given cyber system preserves a given set of security properties (deterministically or probabilistically), in a given context.



Human Behavior

Human Behavior addresses how to handle the unpredictability and complexity of human actors in cybersecurity. These actors include malicious attackers, system users, and software/system developers. The goal of the research is to develop models of human behavior that enable the design, modeling, and analysis of systems with specified security properties.

For additional information on the Hard Problems, see the

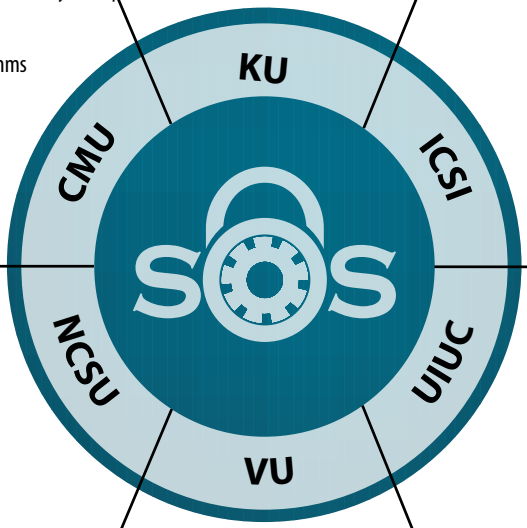
following two documents:

Science of Security Hard Problems: A Lablet Perspective

(November 2012) This document introduces, defines, explains

the rationale of the five Hard Problems and the research needed
View and download from http://cps-vo.org/node/6394.
Science of Security Lablet Progress on the Hard Problems (August 2015) This document highlights major contributions that that SoS Lablets have made towards each of the five Hard Problems. View and download from http://cps-vo.org/node/21590.

The SoS Lablet Initiative



Leads:
Travis Breaux
Jonathan Aldrich

Sub-Lablets:
George Mason University
University of North Carolina

Hard Problems:
- Human Behavior
- Metrics
- Resilient Architectures

Projects:

p. 12 Characterizing user Behavior and Anticipating its Effects on Computer Security with a Security Behavior Observatory
p. 14 Model-Based Explanation For Human-in-the-Loop Security
p. 16 Obsidian Language for Blockchain
p. 18 Securing Safety-Critical Machine Learning Algorithms

Lead:
Perry Alexander

Sub-Lablet:
University of Tennessee

Hard Problems:
- Metrics
- Scalability and Composability
- Resilient Architectures

Projects:

p. 54 Cloud-Assisted IoT Systems Privacy
p. 56 Formal Approaches to the Ontology & Epistemology of Resilience
p.58 Scalable Trust Semantics & Infrastructure
p. 60 Secure Native Binary Execution
p. 62 Side-Channel Attack Resistance

Lead:
Serge Egleman

Sub-Lablets:
Cornell Tech
UC Berkeley

Hard Problems:
- Human Behavior
- Metrics
- Scalability and Composability
- Secure Collaboration

Projects:

p. 21 Contextual Integrity for Computer Systems
p. 23 Designing for Privacy
p. 25 Governance for Big Data
p. 27 Operationalizing Contextual Integrity
p. 29 Scalable Privacy Analysis

Leads:
Laurie Williams
Munindar Singh

Sub-Lablets:
Purdue University
Rochester Institute of Technology
University of Alabama

Hard Problems:
- Metrics
- Resilient Architectures
- Secure Collaboration

Projects:

p. 32 Coordinated Machine Learning-Based Vulnerability & Security Patching for Resilient Virtual Computing Infrastructure
p. 34 Development of Methodology Guidelines for Security Research
p. 36 Predicting the Difficulty of Compromise through How Attackers Discover Vulnerabilities
p. 38 Principles of Secure BootStrapping for IoT
p. 42 Reasoning about Accidental and Malicious Misuse via Formal Methods

Lead:
Xenofon Koutsoukos

Sub-Lablets:
Massachusetts Institute of Technology
UC Berkeley

Hard Problems:
- Human Behavior
- Metrics
- Resilient Architectures
- Scalability and Composability

Projects:

p. 65 Analytics for Cyber-Physical System Cybersecurity
p. 67 Foundations of a CPS Resilience
p. 69 Mixed Initiative and Collaborative Learning in Adversarial Environments
p. 71 Multi-model Test Bed for the Simulation-based Evaluation of Resilience

Leads:
Sayan Mitra
David M. Nicol
William H. Sanders

Sub-Lablets:
Illinois Institute of Technology
UT Austin

Hard Problems:
- Human Behavior
- Metrics
- Resilient Architectures
- Scalability and Composability
- Secure Collaboration

Projects:

p. 43 A Human-Agent-Focused Approach to Security Modeling
p. 45 Automated Synthesis Framework For Network Security and Resilience
p. 47 Monitoring, Fusion, and Response for Cyber Resilience
p. 49 Resilient Control of Cyber-Physical Systems with Distributed Learning
p. 51 Uncertainty in Security Analysis

Carnegie Mellon University

The Carnegie Mellon University (CMU) Science of Security (SoS) Lablet team is led by Principal Investigator (PI) Travis Breaux and Co-PI Jonathan Aldrich. The team of faculty, postdoctoral and PhD student researchers from CMU and Sub-Lablets George Mason University and the University of North Carolina at Chapel Hill are engaged in four research projects. Many of the Lablet faculty and graduate students are affiliated with the CMU CyLab, the coordinating entity for cybersecurity research.

Results of the project research in 2019 have led to papers accepted at the Symposium on Usable Privacy and Security (SOUPS 2019), the 10th ACM/IEEE International Conference on Cyberphysical Systems, the 17th ACM-IEEE International Conference on Formal Methods and Models for Systems Design (MEMCODE\'19), the SPLASH Student Research Competition, and ACM Transactions on Privacy and Security.

Moreover, CMU researchers have given talks on their projects at the IBM Programming Languages Day, the Israel Institute of Technology (Technion), the Etherium Foundation and the Facebook blockchain group, and the Summer school on Real-world Crypto and Privacy in Sibenik, Croatia. The have also briefed Congressional and ARO staff.

Details on the four research projects identified below can be found in the following sections.

- Characterizing User Behavior and Anticipating its Effects on Computer Security with a Security Behavior Observatory
- Model-Based Explanation for Human-in-the Loop Security
- Obsidian Language for Blockchain
- Securing Safety-Critical Machine Learning Algorithms



Travis Breaux



Jonathan Aldrich

Carnegie Mellon University

Characterizing User Behavior and Anticipating its Effects on Computer Security with a Security Behavior Observatory



Lorrie Cranor



Nicolas Christin

PI:	Lorrie Cranor CO-PI: Nicolas Christin
URL:	https://cps-vo.org/CMU_SecurityBehaviorObservatory
HARD PROBLEM:	Human Behavior 

Goal

This research aims to characterize users' choices of password tools and what influence the tools have on password practices. The results of this study will contribute to finding usable solutions for managing authentication.

Abstract

Systems that are technically secure may still be exploited if users behave in unsafe ways. Most studies of user behavior are in controlled laboratory settings or in large-scale between-subjects measurements in the field. Both methods have shortcomings: lab experiments are not in natural environments and therefore may not accurately capture real world behaviors (i.e., low ecological validity), whereas large-scale measurement studies do not allow the researchers to probe user intent or gather explanatory data for observed behaviors, and they offer limited control for confounding factors. The team uses a multi-purpose observational resource, the Security Behavior Observatory (SBO), which was developed to collect data from Windows home computers. The SBO collects a wide array of system, network, and browser data from over 500 home Windows computer users (who participate as human subjects), and this data can be used to investigate a number of aspects of computer security that are especially affected by the hard problem of understanding and accounting for human behavior. Password authentication is one aspect of computer security that is heavily affected by human behavior, since human tendencies and capabilities tend to be directly at odds with what are considered the most secure password practices. This team has previously used data from the Security Behavior Observatory to investigate password practices, including the prevalence of password reuse and the potential effects of password management tools on password habits, and this year they published a follow-up study that used interviews to understand why users are choosing various existing password tools and why those tools are or are not leading to more secure password practices. This team is also conducting ongoing work on a number of related research questions, including investigating how people learn about online data breaches and actions people take in the aftermath of those breaches.

Accomplishments

To follow up on previous findings that suggested that people in the Security Behavior Observatory sample using password managers did not necessarily have stronger passwords or decreased password reuse, this team conducted interviews with a separate sample of 30 participants to better understand password habits and choices of password management tools. These results were published at the Symposium on Usable Privacy and Security. Our results suggested that users of built-in password managers may have different underlying motivations for using password tools (i.e., mostly focused on convenience) and may thus use

those tools to aid their insecure password habits, whereas people using separately installed password managers seem to be more motivated to prioritize security.

We are also conducting ongoing research using the SBO dataset to study how people learn about breaches online and the actions people take in the aftermath of breaches. We study to what extent participants come across breach information in their browsing, and what aspects of people or their browsing increase their likelihood of reading about breach information. We also study the methods by which people come across breach-related pages characteristics of these methods and pages and what influences people to learn more about breaches or take action (either in the form of more browsing or changing passwords), how often people actually change their passwords in the aftermath of a breach, and how constructive these changes are.

Impact on Hard Problems

The Security Behavior Observatory addresses the hard problem of "Understanding and Accounting for Human Behavior" by collecting data directly from people's own home computers, thereby capturing people's computing behavior "in the wild". This data is the closest to the ground truth of the users' everyday security and privacy challenges that the research community has ever collected. We expect the insights discovered by analyzing this data will profoundly impact multiple research domains, including but not limited to behavioral sciences, computer security and privacy, economics, and human-computer interaction.

Education and Outreach

- The SBO dataset has been used in multiple related projects since its inception and will continue to provide researchers across the community with users' actual practices in a variety of security and privacy applications.

Publications

- Sarah Pearman, Shikun Zhang, Lujo Bauer, Nicolas Christin, and Lorrie Cranor, "Why people (don't) use password managers effectively," in Proceedings of the *Symposium on Usable Privacy and Security (SOUPS 2019)*, August 2019.

Model Based Explanation for Human-in-the Loop Security



David Garlan

PI:	David Garlan
URL:	https://cps-vo.org/CMU_Model-BasedHuman-in-the-LoopSecurity
HARD PROBLEM:	Security Metrics and Models, Resilient Architectures, Human Behavior

Goal

Effective response to security attacks often requires a combination of both automated and human-mediated actions. Currently we lack adequate methods to reason about such human-system coordination, including ways to determine when to allocate tasks to each party and how to gain assurance that automated mechanisms are appropriately aligned with organizational needs and policies. This project focuses on combining human and automated actions in response to security attacks, and we will show how probabilistic models and model checkers can be used both to synthesize complex plans that involve a combination of human and automated actions, as well as to provide human-understandable explanations of mitigation plans proposed or carried out by the system.

Abstract

Models that support attack-resiliency in systems need to address the allocation of tasks to humans and systems, and how the mechanisms align with organizational policies. These models include, for example, identification of when and how systems and humans should cooperate, how to provide self-explanation to support human hand-offs, and ways to assess overall effectiveness of coordinated human-system approaches for mitigating sophisticated threats. In this project, we develop a model-based approach to: 1) reason about when and how systems and humans should cooperate with each other; 2) improve human understanding and trust in automated behavior through self-explanation; and 3) provide mechanisms for humans to correct a system's automated behavior when it is inappropriate. We will explore the effectiveness of the techniques in the context of coordinated system-human approaches for mitigating Advanced Persistent Threats (APTs).

Accomplishments

As self-security becomes more automated, it becomes harder for humans who interact with the autonomous system to understand the behavior of the systems. Particularly while optimizing for multiple quality objectives and acting under uncertainty, it can be difficult for humans to understand the system behavior generated by an automated planner. We developed an approach with tool support that aims at clarifying system behavior through interactive explanation by allowing end-users to ask Why and Why-Not questions about specific behaviors of the system, and providing answers in the form of contrastive explanation. Recently we designed and piloted a human study to understand the effectiveness of explanations to human operators.

We conducted a human-subject experiment to evaluate the effectiveness of our explainable planning approach. The experimental design is based on one of our envisioned use-cases of the generated explanation. The use-case scenario aims to address the problem of potential misalignment between the end-user's objectives and the planning agent's reward function, particularly in multi-objective planning settings. The use-case of our explanation approach is to enable the end-user of a planning-based system to identify preference alignment or misalignment between the user's objectives and the planning agent's reward function. The main hypothesis of this experimental study is that our explainable planning approach can improve the end-user's ability to identify such preference (mis)alignment in multi-objective planning settings. This was divided into two sub-hypothesis - H1: Participants who receive the explanations are significantly more likely to correctly determine whether the robot's proposed plan is in line with their (given) preference; and

H2: Participants who receive the explanations have significantly more reliable confidence in their determination (i.e. have higher confidence-weighted scores). In the study, each participant is prescribed a fixed preference over the 3 objectives. The participant is told that the robot may or may not be optimizing for a reward function that aligns with their prescribed preference. When presented with a navigation plan from the robot, the participant is asked to indicate whether they think the robot's plan is the best available option with respect to their prescribed preference. The study participants are divided into 2 groups. The control group does not receive an explanation from the robot, except the prediction of how much time the navigation would take, and how safe and intrusive the navigation is. The experimental group receives a full explanation from the robot. We measured the participant's accuracy of identifying preference alignment and misalignment, the confidence in their answer, and the amount of time they take to come to an answer. The experiment was conducted using Amazon's Mechanical Turk, and the results that we received were compelling:

- H1: The experiment showed that our explanation has a significant effect on the participants' correctness. Overall, the odds of the participants in the treatment group being correct is on average 3.8 times higher than that of the participants in the control group, with 95% confidence interval [2.04, 7.07]. (The fixed-effect logistic regression coefficient estimate is 1.335 and the standard error is 0.317.) Overall, H1 is supported.

- H2: Our explanation has a significant effect on the participants' reliable confidence. Overall, the confidence-weighted score of the participants in the treatment group is on average 1.73 higher than that of the participants in the control group, on the scale of 4 to +4, with 95% confidence interval [1.04, 2.42]. (The fixed-effect linear regression coefficient estimate is 1.727 and the standard error is 0.351.) Overall, H2 is supported.

We are also continuing to work on how to use graceful degradation to respond to security attacks, using DataLog as a reasoning engine. We improved performance of the code by about a factor of about 100-500 through adding richness and realism to the architectural styles, as well as rewriting much of the attack trace generation code. Attack traces now overlay data flows across networks, to evaluate the effects of attacks on data flows (and their associated confidentiality, integrity, and availability security attributes), rather than simply network components, bringing the modeling much closer to the proper way to understand the impact of an attack. We also added richness to a new Functional Perspective, refining and improving on aspects of the DoD Architecture Framework (DoDAF).

In the context of Markov Decision Process (MDP) planning, manually inspecting the solution policy and its value function to gain such understanding is infeasible due to the lack of domain semantics and concepts in which the end-users are interested. They also lack information about which, if any, of the objectives are conflicting in a problem instance, and what compromises had to be made. We investigated an approach for generating automated explanation of a MDP policy that is based on: (i) describing the expected consequences of the policy in terms of domain-specific, human-concept values, and relating those values to the overall expected cost of the policy, and (ii) explaining any tradeoff by contrasting the policy to counterfactual solutions (i.e., alternative policies that were not generated as a solution) on the basis of their human-concept values and the corresponding costs. We demonstrate our approach on MDP problems with two different cost criteria, namely, the expected total-cost and average-cost criteria. Such an approach enhances resilient architectures by helping to explain and have stakeholders explore the decision-making process that goes into automated planning for maintaining system resilience.

Impact on Hard Problems

We are addressing resilience by providing defense plans that are automatically generated as the system runs and accounting for current context, system state, observable properties of the attacker, and potential observable operations of the defense. We are addressing human behavior by providing understandable explanations for automated mitigation plans generated by self-protecting systems that use various models of the software, network, and attack.

Education and Outreach

- We presented our work to the Cylab Partners Conference on October 2, 2019.

Publications

- Ellin Zhao and Roykrong Sukkerd, "Interactive Explanation for Planning-Based Systems," in Proceedings of the *10th ACM/IEEE International Conference on Cyberphysical Systems*, Work-in-Progress track, Montreal, Canada, 16-18 April 2019.

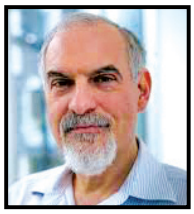
- Cody Kinneer, Ryan Wagner, Fei Fang, Claire Le Goues, and David Garlan, "Modeling Observability in Adaptive Systems to Defend Against Advanced Persistent Threats," in Proceedings of the *17th ACM-IEEE International Conference on Formal Methods and Models for Systems Design (MEMCODE\'19)*, San Diego, USA, 9-11 October 2019.

- Roykrong Sukkerd, David Garlan, and Reid Simmons, "Tradeoff-Focused Contrastive Explanation for MDP planning," Submitted to *AAMAS 2020, the International Conference on Autonomous Agents and Multi-Agent Systems*.

Obsidian: A Language for Secure-by-Construction Blockchain Programs



Jonathan Aldrich



Brad Myers

PI:	Jonathan Aldrich	CO-PI: Brad Myers
SUB-LABLET:	George Mason University	
URL:	https://cps-vo.org/CMU_Obsidian	
HARD PROBLEM:	Scalability and Composability, Policy-Governed Secure Collaboration, Resilient Architectures, Human Behavior	



Goal

Blockchains have been proposed to support transactions on distributed, shared state, but hackers have exploited security vulnerabilities in existing programs. We applied user-centered design in the creation of Obsidian, a new language that uses typestate and linearity to support stronger safety guarantees than current approaches for programming blockchain systems.

Abstract

Programming language designers commonly guess what language designs would be best for their users and create languages accordingly. The outcome of this is languages that are difficult to use and error-prone. In particular, blockchain programs have been plagued by serious bugs. Although techniques from the theory of programming languages can detect many of these kinds of bugs, languages that use these techniques have been too difficult for programmers to use effectively. We have developed Obsidian, which integrates a strong, static type system that detects many of these bugs, using a new user-centered design approach. We have developed formative and summative methods for user-centered design of programming languages and how we have applied them to create Obsidian. This includes a usability study, which demonstrates the effectiveness of our design methods to obtain a usable language. Obsidian addresses the opportunity of directly incorporating models that address the kinds of errors that can occur in distributed systems with shared state and transferable resources. Obsidian uses the technical approaches of typestate (expressing both the types of objects and their state in a way that supports static reasoning) and linearity (to avoid loss or duplication of tracked assets).

This project considers models for secure collaboration and contracts in a decentralized environment among parties that have not established trust. A significant example is blockchain programming, which requires high security but also, in implementations, demonstrates the often-dramatic consequences of defects.

The project research includes both technical and usability assessments of these two ideas. The technical assessment addresses the feasibility of sound and composable static analyses to support these two semantic innovations. The usability assessment focuses on the ability of programmers to use Obsidian effectively to write secure programs with little training. A combined assessment would focus on whether programmers are more likely to write correct, safe code with Obsidian than with Solidity, and with comparable or improved productivity.

Accomplishments

In many developing countries, farming insurance markets have not developed. A severe weather event like a very late hard frost or a drought can therefore devastate these farmers. We are working with the World Bank to develop a parametric insurance platform on the Blockchain with Obsidian to address this need. The platform will serve as an evaluative case study of the expressiveness and effectiveness of the Obsidian language design.

We completed a proof of soundness of the Obsidian language design. Obsidian uses typestate to detect improper state manipulation and uses linear types to detect abuse of assets. The proof shows that Obsidian will correctly detect all possible abuses of these types.

Impact on Hard Problems

Scalability and composability. Obsidian is designed to enable composition of mutually distrusting programs on scalable blockchain platforms. It has a specific focus on preventing specific composition-related vulnerabilities like invalid re-entrancy and protocol violations. This year, we established the soundness of these.

Policy-governed secure collaboration. Obsidian also enables parties that do not fully trust each other to collaborate in a secure way, following a contract that enforces an agreed-upon policy.

Education and Outreach

- Jonathan Aldrich presented "Obsidian: A Language for Smart Contracts Designed for Safety and Usability" at the Lablet Summer Quarterly held at the University of Kansas, 9-10 July.
- We delivered talks on Obsidian at the IBM programming languages day, SPLASH, and at the Israel Institute of Technology (Technion). In addition, we met with leaders of the Etherium Foundation and the Facebook blockchain group about our work.

Publications

- Paulette Koronkevich, "Obsidian in the Rough: A Case Study Evaluation of a New Blockchain Programming Language," in *SPLASH Student Research Competition, 2018*. Awarded 2nd place.
- Michael Coblenz, Jonathan Aldrich, Joshua Sunshine, Brad A. Myers, "Smarter Smart Contract Tools," WETSEB 2019: *2nd International Workshop on Emerging Trends in Software Engineering for Blockchain*, May 27, 2019, Montreal, QC, Canada.
- Gauri Kambhatla, Michael Coblenz, Reed Oei, Joshua Sunshine, Brad Myers, and Jonathan Aldrich, "A Pilot Study of the Safety and Usability of the Obsidian Blockchain Programming Language," the *Tenth Workshop on Evaluation and Usability of Programming Languages and Tools (PLATEAU 2019), at UIST 2019*, October 24, 2019, New Orleans, Louisiana.

Securing Safety-Critical Machine Learning Algorithms



Lujo Bauer



Matt Fredrikson

PI:	Lujo Bauer	CO-PI: Matt Fredrikson
SUB-LABLET:	University of North Carolina	
URL:	https://cps-vo.org/CMU_SecuringSafetyCriticalMLAlgorithms	
HARD PROBLEM:	Security Metrics and Models, Resilient Architectures	

Goal

The goals of this project are: to understand how classifiers can be spoofed, including in ways that are not apparent to human observers; and how the robustness of classifiers can be enhanced, including through explanations of model behavior.

Abstract

Machine-learning algorithms, especially classifiers, are becoming prevalent in safety and security-critical applications. The susceptibility of some types of classifiers to being evaded by adversarial input data has been explored in domains such as spam filtering, but the rapid growth in adoption of machine learning in multiple application domains amplifies the extent and severity of this vulnerability landscape. We propose to: 1) develop predictive metrics that characterize the degree to which a neural-network-based image classifier used in domains such as face recognition can be evaded through attacks that are both practically realizable and inconspicuous; and 2) develop methods that make these classifiers, and the applications that incorporate them, robust to such interference. We will examine how to manipulate images to fool classifiers in various ways, and how to do so in a way that escapes the suspicion of even human onlookers and then develop explanations of model behavior to help identify the presence of a likely attack. We will generalize these explanations to harden models against future attacks.

Accomplishments

We have made significant progress implementing a framework for explaining the predictions made by deep neural networks, and incorporating it into a graphical tool for use by researchers and practitioners. While this tool is useful for our activities on the project, we plan to release it as an open-source project, as well as a more limited interactive web application, for other researchers as well. We have also made progress on leveraging explanations to prune neural network components that contribute to evasion vulnerability. We have developed a model compression technique that results in a smaller model that performs similarly on in-distribution test samples to the original model, and shows promise in correctly classifying adversarial test examples. We are currently working to expand this compression method with a complementary method for efficiently characterizing the evasion vulnerability of individual predictions by structuring the network so that it can be analyzed using a series of linear constraint satisfaction problems.

We have been examining to what extent attackers can fool classifiers in settings in which a successful "attack" must satisfy constraints beyond just fooling a classifier. We developed a generative framework that demonstrates that such attacks are often possible, even when the constraints they need to satisfy cannot be formally specified. We've also made progress in understanding the weaknesses of ML algorithms in another practical setting: as used by anti-virus programs for the detection of malware. This setting is interesting -- aside from its practical relevance -- because previous techniques for creating

evasion attacks usually assumed that the input belonged to a continuous domain, whereas malware binaries are drawn from a discrete domain. We've developed techniques based on binary randomization (as used to harden binaries against return-oriented-programming attacks, for example) that demonstrate how an attacker could alter a malware binary so as to make it avoid detection by malware classifiers.

Finally, we developed a new approach to train *ensembles* of classifiers to better resist attempts to create malicious inputs that would be misclassified. Similar to n-version programming, this approach relies on the assumption that each classifier will make mistakes independently of the others. This assumption typically does not hold for multiple versions of an ML classifier, which are prone to making the same mistakes, and the innovation in our work was in the method for training classifiers to be more purposefully diverse, particularly under adversarial conditions.

Impact on Hard Problems

Both of the Hard Problems this project addresses are tackled in the context of deep neural networks, which are a particularly popular and performant type of machine learning algorithm. This project develops metrics that characterize the degree to which a neural-network-based classifier can be evaded through practically realizable, inconspicuous attacks. The project also develops architectures for neural networks that would make them robust to adversarial examples. The framework for explaining the predictions made by deep neural networks may identify the network-internal factors that cause misclassifications, and we leverage this capability to make progress on the Hard Problems. Finally, as we examine classifiers' robustness to these attacks, we are analyzing the effect that increased robustness may have on other information security metrics, such as those that characterize the confidentiality of training data. Our results have begun to shed new light on the tradeoffs that emerge when certain defensive tactics are employed.

Education and Outreach

- Co-PI Matt Fredrikson is developing a course module in which students apply the attribution framework (developed as part of this project) to convolutional networks. Future plans include extending this module to have students leverage the framework for detecting and explaining adversarial examples.
- PI Lujo Bauer presented work supported by this project at the Summer School on real-world crypto and privacy in Sibenik, Croatia, attended by over 100 PhD students.
- We recruited several new PhD students to work on projects related to adversarial ML starting in the fall of 2019. One PhD student whose primary research was on this topic obtained his PhD in 2019.
- PI Lujo Bauer and Co-PI Matt Fredrikson presented work that was part of this project at the annual CyLab Partners Conference in September, attended by representatives from more than 25 companies and from various government agencies.
- PI Lujo Bauer presented work that was part of this project to Congressional staff on their visit to CMU in September. The work was additionally presented and discussed at a Director, Operational Test and Evaluation (DOT&E) workshop in July, and to Army Research Office (ARO) staff in October.

Publications

- Mahmood Sharif, Sruti Bhagavatula, Lujo Bauer, and Michael K. Reiter, "A General Framework for Adversarial Examples with Objectives," *ACM Transactions on Privacy and Security*, 22(3), June 2019.
- Mahmood Sharif, Lujo Bauer, and Michael K. Reiter. "n-ML: Mitigating Adversarial Examples via Ensembles of Topologically Manipulated Classifiers," arXiv preprint 1912.09059, December 2019.
- Mahmood Sharif, Lucas Keane, Lujo Bauer, Michael K. Reiter, and Saurabh Shintre, "Optimization-Guided Binary Diversification to Mislead Neural Networks for Malware Detection," arXiv preprint 1912.09064, December 2019.
- Samuel Yeom, Irene Giacomelli, Alan Menaged, Matt Fredrikson, and Somesh Jha, "Overfitting, robustness, and malicious algorithms: A study of potential causes of privacy risk in machine learning," *Journal of Computer Security*, 1-36. October 2019.

International Computer Science Institute

The International Computer Science Institute (ICSI) Science of Security and Privacy (SoS) Lablet team is led by Principal Investigator (PI) Serge Egelman. The ICSI Lablet is contributing broadly to the development of privacy science through multiple multi-disciplinary efforts. The overarching goal of this Lablet is to facilitate conducting and disseminating fundamental scientific research on privacy to better understand the implications of data use. Along with Sub-Lablets Cornell Tech and University of California, Berkeley the ICSI researchers are engaged in five research projects.

ICSI project research has been published at the 2019 International World Wide Web (WWW) Conference, the Telecommunications Policy Research Conference, the 2019 ACM Conference on Human Factors in Computing (CHI) conference, ACM Proceedings on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT), the 15th Symposium on Usable Privacy and Security (SOUPS '19), and the 28th USENIX Security Symposium (Distinguished Paper Award).

In addition to presentations at selected conferences, ICSI research has been presented and/or discussed in the following fora:

- Boot Camp for Data Privacy: Foundations and Applications, University of California, Berkeley, Simons Institute for the Theory of Computing
- Simons Institute Theoretically Speaking Series, University of California, Berkeley, Simons Institute for the Theory of Computing,
- Digital Democracies: Artificial Publics, Just Infrastructures, Ethical Learning, Simon Fraser University, Vancouver, May 2019.
- 2019 iConference in Baltimore, Maryland, "Engaging Speculative Practices to Probe Values & Ethics in Sociotechnical Systems"



Serge Egelman

- ACM CSCW 2019 Conference
- World Economic Forum and Stanford University workshop entitled "Redesigning Consent for Better Data Protection"
- NIST Privacy Framework workshop
- FTC PrivacyCon

Details on the five research projects identified below can be found in the following sections.

- Contextual Integrity for Computer Systems
- Designing for Privacy
- Governance for Big Data
- Operationalizing Contextual Integrity
- Scalable Privacy

Contextual Integrity for Computer Systems



Michael Tschantz



Helen Nissenbaum

PI:	Michael Tischantz	CO-PI:	Helen Nissenbaum
SUB-LABLET:	Cornell Tech		
URL:	https://cps-vo.org/ICSI_ContextualIntegrityforComputerSystems		
HARD PROBLEM:	Scalability and Composability, Policy-Governed Secure Collaboration		

The overall goal of the research is to convert the philosophical theory of contextual integrity into terms computer scientists can use. Philosophers and computer scientists have different understandings of context, with philosophers focusing on abstract spheres of life, and computer scientists focusing on the concrete. The goal is to develop models of context and contextual integrity that meet computer scientists on their own truth.

Relevant research questions for this project include accounting for privacy in the design of multi-use computer systems that cut across contexts; modeling the adaptation of contexts to changes in technologies; and determining how contextual integrity relates to differential privacy. The current organizing hypothesis is that contexts are defined by a purpose. The privacy norms of a context promote the purpose, and that purpose restrictions are ubiquitous. There are several possible models including game models, Markov decision process models, partially observable Markov decision process models, and multi-agent influence diagrams. Some of the challenges are that contexts don't exist in a vacuum, contexts might be in competition, privacy is multifaceted, and people often disagree. Potential outcomes are progress on defining privacy, further accountability for big data systems that cut across contexts, and enabling policy governed privacy with respect to collaboration.

Our research will create a formal representation of the contexts found in contextual integrity. Prior work has shown that the term "context" has been interpreted in a wide range of manners. The representation we produce will serve as a reference model for not just comparing different interpretations but also for expressing what Co-PI Helen Nissenbaum, the creator of contextual integrity, sees as the precise form of contexts in her theory. They will also serve as a starting point for adapting contextual integrity to the changing needs of computer science. The current focus is on how a context can be formed by smaller "sub-contexts" composing together. Our working hypothesis is that the "values" of a sub-context may come from the purpose of the super-context.

We have started to translate the concept of context into a formal representation and have been working on a model of context in a formal representation similar to an MDP. After finding that the model was not flexible enough for accommodating disagreements about the state of a system, we developed a new model that is more flexible by adding a layer of interpreted predicates between norms and states. The new model includes the possibility of values coming from global context, to model deontological approaches to ethics.

Our work surfaced that the notion of norms found in the theory of contextual integrity is difficult to precisely pin down, and we will look more closely at the process that legitimates norms in

hopes of providing an operational definition for at least legitimate ones. Our work on sub-contexts has found that a similar concept of sub-goals may be conceptually primary, and they are working toward a model of this prerequisite.

Impact on Hard Problems

The CI framework is being used to abstract real-world communication exchanges into formally defined information flows where privacy policies describe sequences of admissible flows. CI allows us to decouple (1) the syntactic extraction of flows from information exchanges, and (2) the enforcement of privacy policies on these flows. As an alternative to predominant approaches to privacy, which were ineffective against novel information practices enabled by IT, CI was able both to pinpoint sources of disruption and provide grounds for either accepting or rejecting them. Growing challenges from a burgeoning array of networked, sensor-enabled devices (IoT) and data-ravenous machine learning systems, similar in form though magnified in scope, call for renewed attention to theory.

Education and Outreach

- "Achieving Meaningful Privacy in Technology and Design," Boot Camp for Data Privacy: Foundations and Applications, University of California, Berkeley, Simons Institute for the Theory of Computing, January 2019.
- "Applications of Contextual Integrity," Workshop on Data Privacy: Foundations and Applications, University of California, Berkeley, Simons Institute for the Theory of Computing, March 2019.
- "What We Can Learn from AdNauseam About the Threat and Power of Data Obfuscation, and Efforts to Suppress It," Simons Institute Theoretically Speaking Series, University of California, Berkeley, Simons Institute for the Theory of Computing, March 2019 and Digital Democracies: Artificial Publics, Just Infrastructures, Ethical Learning, Simon Fraser University, Vancouver, May 2019.
- "Using MDPs to Model Contextual Integrity," 2nd Annual Symposium on Applications of Contextual Integrity, University of California, Berkeley, August 2019.
- PI Michael Tschantz and Co-PI Helen Nissenbaum have refined their model of context in a formal representation similar to an MDP.

Publications

- Yan Shvartzshnaider, Zvonimir Pavlinovic, Ananth Balashankar, Thomas Wies, Lakshminarayanan Subramanian, Helen Nissenbaum, and Prateek Mittal, "VACCINE: Using Contextual Integrity For Data Leakage Detection," in Proceedings of the *2019 International World Wide Web (WWW) Conference*.
- Helen Nissenbaum (2019) "Contextual Integrity Up and Down the Data Food Chain," Theoretical Inquiries in Law 20:1, 221-256.
- Madelyn Sanfilippo, Yan Shvartzshnaider, Irwin Reyes, Helen Nissenbaum, and Serge Egelman, "Disaster Privacy/Privacy Disaster," *Telecommunications Policy Research Conference 47,* Washington DC, 2019 [Peer reviewed Conference and Proceedings].

Designing for Privacy



Deirdre Mulligan

PI:	Deirdre Mulligan
SUB-LABLET:	University of California at Berkeley
URL:	https://cps-vo.org/ICSI_DesigningforPrivacy
HARD PROBLEM:	Policy-Governed Secure Collaboration, Human Behavior

Goal

Design interventions for privacy can occur at a lot of stages and levels, and the goal of the project is to develop a new toolbox of techniques and help designers understand when best to apply tools.

Abstract

The project focuses on designing for privacy holistically: from "privacy by design" to "privacy with design," i.e., designing with privacy throughout whole life cycle. Privacy is defined in contextual, situational, and relational ways, and its dimensions are theory, protection, harm, provision, and scope. The goal over the next year is to put together design card activities, design workbooks, and privacy design patterns. We also plan to hold privacy design workshops to address engineering practices, methods, and tools, bringing together practitioners, researchers, and policy-makers. One goal for this series of workshops is to examine how current approaches to privacy engineering (e.g., applying Privacy by Design principles) are actually being applied in practice--that is, are there human limitations that are preventing these recommended practices from being used? Another goal is to examine how privacy engineering practices can be improved via policy, both at the organizational level and governmental.

Accomplishments

- Reviewed collaborators' "Won't Somebody Think of the Children?" examining COPPA Compliance at Scale" paper and identified existing privacy patterns that are discussed or could be solutions
- Reviewed FTC privacy actions and identified consent decree requirements and existing privacy patterns that could be solutions
- Completed outlining and began writing a solutions-oriented paper, using Reyes et al. and semi-automated analysis of COPPA violations in smartphone applications and FTC privacy actions as initial corpora, that connects identified privacy problems to existing and new privacy design patterns, and patterns targeted at different ecosystem parties
- Interviewed 16 engineers and designers about privacy and values practices and conducted 15 hours of observations at public meetup events for engineers and designers thinking about privacy and related values
- Continued development on a set of group-based design activities to help practitioners raise questions about privacy

Impact on Hard Problems

Human Behavior: One goal for the series of workshops is to examine how current approaches to privacy engineering (e.g., applying Privacy by Design principles) are actually being applied in practice. That is, are there human limitations that are preventing these recommended practices from being used?

Policy-Governed Secure Collaboration: Another goal is to examine how privacy engineering practices can be improved via policy, both at the organizational level and governmental.

Education and Outreach

- Co-hosted a Session for Interactivity & Engagement at the 2019 iConference in Baltimore, Maryland, "Engaging Speculative Practices to Probe Values & Ethics in Sociotechnical Systems" which included design workshop activities to think about privacy issues related to facial recognition technology.
- Co-organized a workshop at the ACM CSCW 2019 Conference (November 2019) in Austin, TX, for a workshop on privacy and design titled "Ubiquitous Privacy: Research and Design for Mobile and IoT Platforms".
- Participated in a 2-day workshop in San Francisco co-hosted by the World Economic Forum and Stanford University entitled "Redesigning Consent for Better Data Protection".

- Richmond Wong and Deirdre Mulligan, "Bringing Design to the Privacy Table: Broadening "Design" in "Privacy by Design" Through the Lens of HCI," *2019 ACM CHI conference*

Governance for Big Data



Deirdre Mulligan

PI:	Deirdre Mulligan
Sub-Lablet	University of California at Berkeley
URL:	https://cps-vo.org/ICSI_GovernanceforBigData
HARD PROBLEM:	Policy-Governed Secure Collaboration, Human Behavior

Goal

This project aims to synthesize computer science abstractions with governance goals.

Abstract

The risk in governance for big data is that access control does not capture privacy requirements. With respect to sensitive inferences and reidentification, it is difficult to redact sensitive information from rich data sets, and often sensitive data can be reidentified using additional information outside the data set or proxies. It is possible that Machine Learning will find such correlations automatically; binary allow/deny access control fails to capture this well. In limiting sensitive inferences, there are several related issues, including differential privacy, encryption and access control, and fairness issues.

Accomplishments

In addition to continuing our exploration of the literature on privacy governance methods and technologies, we also conducted a preliminary analysis of risks that privacy enhancing technologies provide privacy in name only, examining the extent to which deployments of differential privacy by industry actors in fact protect the privacy of individuals' data, producing a public position paper on the topic. We began a related exploration of provable privacy protection technologies, examining them from as representations which reduce the dimensionality of the input data to limit inference capacity.

With respect to Data Governance, we have collected and mapped out origins and utility and developed a taxonomy to relate 24 flavors of provable privacy. We sought and received an IRB exemption and we have begun interviews.

Impact on Hard Problems

A new data governance approach focuses on accountability and relates more to accounting and auditing. The first step is to develop a design methodology from all different approaches and mechanisms, and then validate the design methodology by working with practitioners and building case studies for generalizable design patterns.

Education and Outreach

- We attended the second NIST privacy framework workshop, in Atlanta, and both observed and interacted with participants as they evaluated and provided feedback on NIST work product and plans

Publications

- Joshua Kroll and Nitin Kohli, "Privacy and Policy in Polystores: A Data Management Research Agenda," in *Proceedings of the Springer Poly 2019/ DMAH 2019*, LNCS 11721

Operationalizing Contextual Integrity



Serge Egelman

PI:	Serge Egelman	CO-PI:	Helen Nissenbaum
SUB-LABLET:	Cornell Tech		
URL:	https://cps-vo.org/ICSI_OperationalizingContextualIntegrity		
HARD PROBLEM:	Scalability and Composability, Policy-Governed Secure Collaboration, Security Metrics and Models, Human Behavior		



Helen Nissenbaum

Goal

Our ultimate goal is to design new privacy controls that are grounded in the theory of contextual integrity so that they can automatically infer contextual norms and handle data-sharing and disclosure on a per-use basis.

Abstract

This project centers around work on mobile device apps that is the basis for what we plan to do in the future, addressing privacy as contextual integrity. Inappropriate data flows violate contextual information norms; data flows occurring within specific contextual information norms are modeled using as a data subject, data sender, data recipient, information type, and transmission principle (constraints). In questioning what this means for user-centered design, it is suggested that an app should only provide notice when reasonable privacy expectations are expected to be violated. The next steps are to determine what parameters are actually important to users:

- Phase 1: Factorial vignette studies (interviews, surveys; randomly generated scenarios based on controlled parameters)
- Phase 2: Observational studies (instrument phones, detect parameters and resulting behaviors)

We are working on improving infrastructure to allow us to study privacy behaviors *in situ*, long-term project planning to examine new ways of applying the theory of contextual integrity to privacy controls for emergent technologies (e.g., in-home IoT devices), and constructing educational materials based on our research findings for use in the classroom.

Accomplishments

In addition to the publications and engagement events, we are currently designing new studies to better understand users' privacy perceptions surrounding in-home voice assistants, and voice capture in general. The goal is to gather data that can be used to predict privacy-sensitive events based on contextual data.

We deployed a study to examine contextual norms around in-home audio monitoring, which is likely to proliferate as new devices appear on the market. We are recruiting users of both the Google Home and Amazon Echo to answer questions about previously-recorded audio from their devices. Both manufacturers make audio recordings accessible to device owners through a web portal, and so our study involves using a browser extension to randomly present these clips to users,

and then have them answer questions about the circumstances surrounding the recordings. We're interested in whether they were aware that the recordings were made, how sensitive the content was, as well as participants' preferences for various data retention and sharing policies.

In another set of studies, we're examining existing audio corpora, and then using crowdworkers to identify sensitive conversations, that we can then label and use to train a classifier. The goal is to design devices that can predict when they should not be recording or sharing data. We deployed this study for several hundred participants, and are now in the process of reviewing the data.

Impact on Hard Problems

- Scalability and Composability: Ultimately, our goal is to be able to design systems that function on contextual integrity's principles, by automatically applying inferred privacy norms from one context and applying them to future contexts.
- Policy-Governed Secure Collaboration: One goal of this project is to examine how policies surrounding the acceptable use of personal data can be adapted to support the theory of contextual integrity.
- Security Metrics and Models: We seek to build models of human behavior by studying it in both the laboratory and the field. These models will inform the design of future privacy controls.
- Human Behavior: We are designing human subjects studies to examine how privacy perceptions change as a function of contextual privacy norms. Our goal is to design and develop future privacy controls that have high usability because their design principles are informed by empirical research.

Education and Outreach

- We organized a workshop, open to the public, to present nascent research that makes use of the Conextual Integrity framework. The "Symposium on Applications of Contextual Integrity" was to have been hosted at Berkeley on August 19-20, and feature more than 20 different speakers, but had to be cancelled.
- We hosted a tutorial at the Symposium on Usable Privacy and Security (SOUPS), in which we explained the CI framework and showed other researchers how they can apply it in their own privacy research.

Publications

- Madiha Tabassum, Tomasz Kosinski, Alisa Frik, Nathan Malkin, Primal Wijesekera, Serge Egelman, and Heather Lipford, "Investigating Users' Preferences and Expectations for Always-Listening Voice Assistants," in *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT)*, 2020.
- Alisa Frik, Leysan Nurgalieva, Julia Bernd, Joyce Lee, Florian Schaub, and Serge Egelman, "Privacy and Security Threat Models and Mitigation Strategies of Older Adults," in *Proceedings of the 15th Symposium on Usable Privacy and Security (SOUPS '19),* 2019, Berkeley, CA, USA.
- Nathan Malkin, Joe Deatrick, Allen Tong, Primal Wijesekera, David Wagner, and Serge Egelman, "Privacy Attitudes of Smart Speaker Users," in *Proceedings on Privacy Enhancing Technologies (PoPets), 2019 (4).*
- Lysan Nurgalieva, Alica Frik, Francesco Ceschel, Serge Egelman, and Maurizio Marchese, "Information Design in An Aged Care Context: Views of Older Adults on Information Sharing in a Care Triad," in *The 13th International Conference on Pervasive Computing Technologies for Healthcare)Pervasive Health '19)*, May 20-23, 2019, Trento, Italy. ACM, New York, NY,USA.

Scalable Privacy Analysis



Serge Egelman

PI:	Serge Egelman	CO-PI:	Narseo Vallina-Rodriquez
URL:	https://cps-vo.org/ICSI_ScalablePrivacyAnalysis		
HARD PROBLEM:	Scalability and Composability, Policy-Governed Secure Collaboration, Security Metrics and Models		



Narseo Vallina-Rodriquez

Goal

We have constructed a toolchain that allows us to automatically perform dynamic analysis on mobile apps to monitor what sensitive personal information they attempt to access, and then to whom they transmit it. This is allowing us to perform large-scale studies of the privacy behaviors of the mobile app ecosystem, as well as devise new methods of protecting user privacy.

Abstract

Governments and private organizations codify expectations of privacy into enforceable policy. These policies have taken such forms as legislation, contracts, and best practices, among others. Common to these rules are definitions of what constitutes private information and which uses of that information are appropriate or inappropriate. Additionally, policies might place restrictions on what pieces of data may be collected, for what purposes it may be used, how long that data may be retained for yet-unspecified future applications, and under which circumstances (if any) are disclosure and dissemination to other parties permitted.

Different motivations drive different policies. There are procedures and restrictions meant to maintain strategic advantages for holders of sensitive information. The United States government, for instance, routinely classifies information based on the amount of harm to national interests its disclosure would bring. Other policies on data usage seek to protect vulnerable populations by establishing rules limiting how information from those individuals is collected and used: the Family Educational Rights and Privacy Act (FERPA) requires appropriate consent before an individual's educational records are disclosed; the Health Insurance Portability and Accountability Act (HIPAA) regulates the use of Protected Health Information (PHI) by defining what is considered PHI and how individual patients should be de-identified in records prior to aggregation for research purposes; and the Children's Online Privacy Protection Act (COPPA) prohibits the collection of personal information (e.g., contact information and audio/visual recordings) by online services from users under 13 years of age.

The problem is that the constraints for data usage stated in policies—be they stated privacy practices, regulation, or laws—cannot easily be compared against the technologies that they govern. To that end, we propose a framework to automatically compare policy against practice. Broadly, this involves identifying the relevant data usage policies and practices in a given domain, then measuring the real-world exchanges of data restricted by those rules. The results of such a method will then be used to measure and predict the harms brought onto the data's subjects and holders in the event of its unauthorized usage. In doing so, we will be able to infer which specific protected pieces of information, individual prohibited operations on that data, and aggregations thereof pose the highest risks compared to other items covered by the policy. This will shed light on the relationship between the unwanted collection of data, its usage and dissemination, and resulting negative consequences.

Accomplishments

During the past year, we used our app testing infrastructure to automatically detect security vulnerabilities in the Android operating system. We published our results at the USENIX Security Symposium, which demonstrated how we used our infrastructure to automatically detect tens of thousands of apps that were actively exploiting vulnerabilities in the Android permissions system. We were awarded the Distinguished Paper Award for this work, and Google provided us a bug bounty and announced changes to Android Q based on our findings.

In a second line of inquiry, we performed a study to evaluate the privacy protections afforded by paid versions of free mobile apps. In doing this side-by-side comparison, we observed that almost half of the paid "premium" app versions collected the same personal data and shared it with the same third parties as their free counterparts. Worse, we show that consumers have no way of knowing when this is occurring.

Impact on Hard Problems

Human Behavior: Our primary goal is to understand the mobile privacy landscape, so that we can better understand how to provide end-users with usable privacy controls. This specifically involves understanding how mobile apps collect sensitive data and the shortcomings of existing privacy mechanisms in practice.

Policy-Governed Secure Collaboration: Another goal is to assess the effectiveness of privacy-related policies at scale, so that more effective policies—based on empirical data—can be proposed.

Education and Outreach

- PI Egelman presented this research at the FTC's PrivacyCon event, which was covered by multiple media outlets
- We reported several security vulnerabilities to Google based on our mobile app analysis findings. Google is providing us with a bounty for one of the vulnerabilities. This vulnerability is actively being exploited by multiple ad SDKs, which we've reported to the FTC, and expect to follow up with them. Similarly, we received a bounty from Facebook for reporting massive misuse of their SDK.

Publications

- Joel Reardon, Alvaro Feal, Primal Wijesekera, Amit Elazari Bar On, Narseo Vallina-Rodriguez, and Serge Egelman, "50 Ways to Leak Your Data: An Exploration of Apps' Circumvention of the Android Permissions System," in *Proceedings of the 28th USENIX Security Symposium, 2019,* **Distinguished Paper Award**
- Ehimare Okoyomon, Nikita Samarin, Primal Wijesekera, Amit Elazari Bar On, Narseo Vallina-Rodriguez, Irwin Reyes, Álvaro Feal, and Serge Egelman, "On the Ridiculousness of Notice and Consent: Contradictions in App Privacy Policies," *Workshop on Technology and Consumer Protection (ConPro '19)*, 2019.
- Catherine Han, Irwin Reyes, Amit Elazari Bar On, Joel Reardon, Álvaro Feal, Kenneth A. Bamberger, Serge Egelman, and Narseo Vallina-Rodriguez, "Do You Get What You Pay For? Comparing The Privacy Behaviors of Free vs. Paid Apps," *Workshop on Technology and Consumer Protection (ConPro '19),* 2019.

North Carolina State University

The North Carolina State University (NCSU) Science of Security (SoS) Lablet team is led by Principal Investigator (PI) Laurie Williams and Co-PI Munindar Singh. The team of faculty and PhD student researchers from NCSU and Sub-Lablets Purdue University, Rochester Institute of Technology, and University of Alabama are engaged in five research projects.



Laurie Williams



Munindar Singh

Results of project research in 2019 have led to papers accepted at the IEEE International Conference on Cloud Engineering (IC2E), the HICSS Symposium on Cybersecurity Big Data Analytics, the 41st International Conference on Software Engineering (ICSE), the 42nd International Conference on Software Engineering (ICSE), the 13th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM), the 34th Conference on Artificial Intelligence (AAAI), in Proceedings on Privacy Enhancing Technologies (PoPETs), and the 2019 Network and Distributed System Security Symposium (NDSS).

NCSU researchers' work dealing with the Collegiate Penetration Testing Competition (CPTC) has been presented in multiple fora, and one of the researchers worked on the monitoring teams at both the regional and national level. The national CPTC data set has been released via their papers and can now be used in the classroom. NCSU researchers identified the underlying design weaknesses of 4G and 5G cellular paging protocols and proposed approaches to address these weaknesses. Those findings were provided to various software vendors and the GSM Association (GSMA) leading GSMA to list the PI in their Mobile Security Research Hall of Fame. NCSU researchers conducted a workshop at the 2019 International Software Engineering Research Network (ISERN) meeting to gather feedback on Version 2 of the paper review rubric being developed under the methodology for security research effort, and began reaching out to potential interviewees to gauge interest in the rubric. Members of the NCSU Lablet routinely bring up Science of Security at the conferences and meetings they attend.

Details on the research projects identified below can be found in the following sections.

- Coordinated Machine Learning-Based Vulnerability and Security Patching for Resilient Virtual Computing Infrastructure
- Development of Methodology Guidelines for Security Research
- Principles of Secure Bootstrapping for IoT
- Predicting the Difficulty of Compromise through How Attackers Discover Vulnerabilities
- Reasoning about Accidental and Malicious Misuse via Formal Methods

NC STATE UNIVERSITY

Coordinated Machine Learning-Based Vulnerability and Security Patching for Resilient Virtual Computing Infrastructure



Xiaohui (Helen) Gu

PI:	Xiaohui (Helen) Gu
URL:	https://cps-vo.org/NCSU_CoordinatedML-BasedVulnerability
HARD PROBLEM:	Resilient Architectures

Goal

Our research aims to assist administrators of virtualized computing infrastructures in making services more resilient to security attacks. We do that through applying machine learning to reduce both security and functionality risks in software patching by continually monitoring patched and unpatched software to discover vulnerabilities and triggering proper security updates.

Abstract

The existing approach to making services more resilient to security attacks is static security analysis and scheduled patching. In our experiments, this approach fails to detect 90% of vulnerabilities, displays high false alarms, and shows memory inflation caused by unnecessary security patching. This project is runtime vulnerability detection using online machine learning methods and just-in-time security patching. Just-in-time security patching includes applying patches intentionally after attacks are detected, enforcing update validation, making intelligent decisions on update vice rebuild, and adhering to system operational constraints.

Accomplishments

Containers have become increasingly popular for deploying applications in cloud computing infrastructures. However, our previous study showed that containers are prone to various security attacks. We are continuing our research on real time container vulnerability discovery, refining our approaches for aggregated learning and application classification. We refined our dynamic vulnerability exploit detection system (changing the algorithm from SOM to autoencoder) and collected results over a set of real world exploits. Our initial results using a limited set of exploits show that our exploit detection has 100% detection rate and lower than 1% false alarms. We implemented our aggregated behavior learning framework that leverages learning data from different containers running the same application together. After conducting dynamic container vulnerability exploit detection experiments over 31 vulnerability exploits (e.g., return a shell, execute arbitrary code, privilege escalation) detected in 23 commonly used server applications (e.g., Apache, CouchDB, Bash, ElasticSearch, JBoss, OpenSSH), we compared the detection accuracy of aggregated behavior learning versus individual behavior learning. Our experiments show that aggregated learning can improve the accuracy and detection lead time in 70% of tested exploits. We further conducted an initial study on a container classification scheme that can recognize containers running the same application using runtime monitoring data. The experiments show that we can achieve over 80% classification accuracy. We saw improved detection results while training using a larger number of containers. We further refined our application classification algorithm using random forest learning methods and can achieve over 96% accuracy now. We also completed attack type classification implementation by extracting top frequently used system calls during the attack period and conducted extensive experiments. Our attack signature can achieve 90% classification accuracy over all tested 31 vulnerability exploits. We also started to implement the on-demand targeted patching (OPatch) system. Our initial experiments show that OPatch can reduce memory overhead

incurred by patching by up to 84% and disk overhead by up to 40% compared to the existing whole software upgrade approach.

Impact on Hard Problems

Our approach in applying machine learning for security patching seeks to make services in virtualized computing infrastructures more resilient to security attacks.

Education and Outreach

- One of our PhD students attended IC2E to present our container vulnerability exploit detection work.

Publications

- Olufogorehan Tunde-Onadele, Jingzhu He, Ting Dai, and Xiaohui Gu, "A Study on Container Vulnerability Exploit Detection," in
- *Proceedings of IEEE International Conference on Cloud Engineering (IC2E)*, Prague, Czech Republic, June, 2019.

Development of Methodology guidelines for Security Research



Jeffrey Carver

PI:	Jeffery Carver
SUB-LABLET:	University of Alabama
URL:	https://cps-vo.org/NCSU_MethodologyGuidelinesforSecurityResearch

Goal

The goal of this project is to aid the security research community in conducting and reporting methodologically sound science through development, refinement, and use of community-based security research guidelines. We proposed the characterization of the security literature based upon those guidelines

Abstract

This research project is aimed at providing support to researchers interested in the quality of scientific reporting in the cybersecurity community by developing a rubric to provide insight into the scientific rigor of the information included in a research report (i.e. a journal or conference paper). By providing a rubric so that researchers report the most important information relative to scientific rigor, this project will help ensure that other researchers are more easily able to replicate published cybersecurity research. It will also help the readers of these publications better analyze their importance and usefulness in the current environment. Lastly, by helping to ensure that all crucial information is present in papers, this project will support theory building that can provide a foundation for additional research in the future. To create this rubric, we are interviewing experts from the cybersecurity community. The goal of these interactions is to determine what type of information is most important for judging scientific rigor in different facets of the cyber security community. To ensure that the rubric is widely usable and accepted, it is important that we included experts from different parts of the cyber security landscape who can provide different perspectives. In addition to interviewing researchers who are connected to the Lablets, we will also seek to interview researchers outside the Lablets who can provide different perspectives. This approach not only accommodates the diversity of the field itself but allows for a larger portion of the research community to have input into the contents of the rubric. By gathering input from a wide swath of researchers representing different perspectives, our rubric should be more widely accepted in the larger cyber security research community.

Accomplishments

1. *Development of the interview process*

Based on our prior experience, we have developed a series of questions that will allow us to elicit important information from the interviewees regarding scientific rigor in cyber security publications.

2. *Identification of interview participants*

We have interacted with our NSA champion to begin identifying the list of people from within the Lablets to interview. Our goal in building this list of people is to include interviewees who represent different parts of the cyber security domain. Building from the initial list provided, we have expanded the list based on needs. We have also begun building a list of cyber security experts who are outside the Lablet community in preparation for our second round of interviews.

3. *Obtain IRB approval*

We have obtained IRB approval to conduct the research.

4. *Schedule and conduct interview*

We were able to schedule two in-person interviews during the January Quarterly Meeting at NCSU. We used these two interviews both to get helpful information for the rubric, but also to pilot test the interview process and questions. The results

of these interviews have already begun to provide interesting insights. We are currently in the process of scheduling 5-10 more interviews from people who have already expressed interest. We will conduct these interviews via Skype/Zoom/Telephone.

5. Initial Results

In our initial set of interviews, we have already identified some important themes that we expect will expand as we talk to more people. First, both interviewees identify the problem statement as the most critical aspect of a paper when trying to determine its scientific rigor. The problem statement must be both crisp and informative. If the reader is not clear on exactly what the problem is and how to know if it is solved, then the remainder of the paper is less useful. Second, both interviewees mentioned that there are some types of papers for which the contribution is about a novel idea rather than about a scientifically rigorous study. The eventual rubric will need to ensure that it has a way both to identify and to evaluate this type of paper.

6. Paper illustrating good examples

Based on some of the work done earlier under Science of Security, we are in the process of assembling a paper that provides examples from published cyber security literature to illustrate various aspects of scientific rigor and how they can be done well in a research paper. This paper will help people in the cyber security community see good examples from literature with which they are familiar.

Impact on Science of Security

This project addresses the challenges of research methods and community. The rubric we construct will provide a basis upon which to judge the scientific rigor in the reporting of cyber security research. This rubric will have community input to ensure it captures all relevant perspectives and is applicable to different types of research methodologies and problems.

Education and Outreach

- One doctoral student has worked on this project.

Publications

- None thus far

Predicting the Difficulty of Compromise through How Attackers Discover Vulnerabilities



Andy Meneely

PI:	Andy Meneely	CO-PI:	Laurie Williams
SUB-LABLET:	Rochester Institute of Technology		
URL:	https://cps-vo.org/NCSU_PredictingDifficultyofCompromise		
HARD PROBLEM:	Security Metrics and Models		



Laurie Williams

Goal

Our goal is to provide actionable feedback on the discoverability of a vulnerability. This feedback is useful for in-process software risk assessment, incident response, and the vulnerabilities equities process. Our approach is to combine the attack surface metaphor and attacker behavior to estimate how attackers will approach discovering a vulnerability. The researchers want to develop metrics that are useful and improve the metric formulation based on qualitative and quantitative feedback.

Abstract

This project focuses on the attack surface based on the notion that pathways into the system enable attackers to discover vulnerabilities. This knowledge is important to software developers, architects, system administrators, and users. A literature review to classify attack surface definitions led to six clusters of definitions which differ significantly (methods, avenues, flows, features, barriers, and vulnerabilities). The methodology used to discover the attack surface (mining stacktraces from thousands of crash reports) and what the attack surface meant within the context of metric actionability, will lead to evolving the models for risky walk and deploying a human-in-the-loop study. Attacker behavior data is gathered from the National Collegiate Penetration Testing Competition (CPTC) from years 2018 and 2019.

Accomplishments

A national collegiate penetration testing competition allowed us to collect a massive data set. Nine teams from schools across the United States performed coordinated sets of network and device penetration attempts during the 2018 National Collegiate Penetration Testing Competition (CPTC). During the ten-hour attack window, the teams generated an aggregate of more than 300GB of alert logs across duplicate networks consisting of 252 virtual machines in total. We captured and made available full images of 99 virtual machines (as .vmdk) and approximately 200GB of alert log data (as JSON) for the six teams who consented to allow their data to be published in this academic research. The inclusion of virtual machine instances along with network alert data is a novel contribution to the research community.

This national collegiate penetration testing competition data set enables us to provide a fine-grained history of vulnerability discovery and exploitation. With this data, we can enrich our models of the attack surface which will in turn lead to more robust metrics of difficulty to compromise. Given that this data is from a competition where teams were assigned the same systems and evaluated on their attacks, the difficulty to compromise will come from the correlation of competition data with the alert and virtual machine data.

In completing the collection and analysis of the CPTC data set, we manually annotated and curated the data from the CPTC data set. We have over 400 events logged with 79 vulnerabilities reported. We have classified each event as part of the MITRE ATT&CK framework, with reviewer agreement kappa scores over 80%. We have begun our attacker behavior model analysis. As part of our submission to the ESEM 2019 New Ideas and Emerging Results track (ESEM NIER), we honed our annotation process by curating one team's entire timeline. Then, to scale up

our approach, we hired a CPTC-trained competitor to capture the timelines and map the events to MITRE ATT&CK. The resulting data set will be a fine-grained log of events that map the techniques attackers use to break into systems in the competition environment. The next step for this data set will be to create a probabilistic model that can estimate the probability of discovery for a given vulnerability based upon attacker behavior. A single team had 47 relevant events with approximately 17 vulnerabilities found, which we filtered from millions of events from the SPLUNK monitoring system.

We also collected data from the 2019 CPTC competition, which included over 7TB of logs ranging from command histories, process information, network metadata, and a variety of other fine-grained measurements from the Splunk incident response system. Our techniques for gathering data and creating timelines were optimized ahead of this competition through our research assistants volunteering to contribute to the competition's monitoring team. We were able to create example infrastructures and tested out our monitoring tools to ensure reliable data collection. With this data, we can re-create the timelines and deliver them faster to the research public.

Defects in Infrastructure as Code (IaC) scripts can have serious consequences for organizations who adopt DevOps. While developing IaC scripts, practitioners may inadvertently introduce security smells. Security smells are recurring coding patterns that are indicative of security weakness and can potentially lead to security breaches. The goal of this work is to help practitioners avoid insecure coding practices while developing IaC scripts through an empirical study of security smells in IaC scripts. We expanded the scale and depth of our previously-reported security smell work by increasing the number of languages to three: Ansible, Chef, and Puppet (prior reported results were for the Puppet language only). We identify nine security smells for IaC scripts.

Security testers often perform a series of operations to exploit vulnerabilities in software source code. We are conducting a study to aid them in exploiting vulnerabilities in software by synthesizing attack mechanisms used in open source software. By characterizing these operations, we will identify patterns of tools and techniques used to exploit a vulnerability. This information will help security test detect vulnerabilities before code changes are integrated and deployed. We are currently investigating security reports archived in open source software bug reports, such as the Bugzilla reports, hosted by Mozilla to characterize what steps are executed to exploit a vulnerability. We are using the Structured Threat Information eXpression (STIX) format to record and standardize the steps needed to exploit a vulnerability.

We have begun our analysis of stacktraces in Red Hat data to discover attack surface insights.

Impact on Hard Problems

We are developing a new set of metrics for measuring exploitability using the attack surface. These metrics are based on the behavior observed by penetration testers in a competition environment. The intrusion detection data collected from the CTPC have provided us with detailed timelines of how attackers find, exploit, and pivot with vulnerabilities. When studying how they work with the known attack surface, we will develop metrics that show which vulnerabilities are at highest risk based on the current deployment.

Education and Outreach

- CO-PI Laurie Williams presented the paper dealing with the CTPC at ICSE 2019, and the paper also won an ACM Distinguished Paper award.
- One of our researchers worked on the monitoring team at the regional CPTC as well as the national CPTC.
- The national collegiate penetration testing competition data set has now been released via our publication, which can be used in the classroom.
- We have begun consulting with RIT's Department of Computing Security in working on a new Foundations of Computing Security that would be required for all first-year students in the Golisano College of Computing and Information Sciences (about a thousand students per year). The work from this grant helps shape that course by providing specific examples from competition for students to analyze.
- PI Andy Meneely presented our ESEM NIER paper in Porto de Galinhas, Brazil

Publications

- Justin Pelletier, Nuthan Munaiah, Shau-Hsuan Su, S. Jay Yang, and Andrew Meneely, "A Cybersecurity Dataset Derived from the National Collegiate Penetration Testing Competition," *HICSS Symposium on Cybersecurity Big Data Analytics*. To appear.
- Rahman Akond, Chris Parnin, and Laurie Williams, "The Seven Sins: Security Smells in Infrastructure as Code Scripts," Proceedings of the *41st International Conference on Software Engineering*. To appear
- Nuthan Munaiah, Akond Rahman, Justin Pelletier, Laurie Williams, and Andrew Meneely, "Characterizing Attacker Behavior in a Cybersecurity Penetration Testing Competition," in proceedings of the *13th ACM/ IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)* (New Ideas and Emerging Results track), Porto de Galinhas, Sep 2019, pages 1--6.

Principles of Secure Bootstrapping for IoT



Ninghui Li

PI:	Ninghui Li
SUB-LABLET:	Purdue University
URL:	https://cps-vo.org/NCSU_PrinciplesofSecureBootstrapping
HARD PROBLEM:	Policy-Governed Secure Collaboration

Goal

This research is motivated by the fact that IoT devices need trust and secure communication—trust between devices and trust between device and users. Constraints, however, limit options, and deployment scenarios determine resource availability, including power supply, computing resources, and serviceability.

The research goal is to develop a lexicon and principles to model the different IoT security bootstrapping scenarios and tools to help developers. The success criteria include being able to see the developed lexicon and develop the most important IoT bootstrapping tool.

Our research plan for modeling IoT bootstrapping scenarios is as follows:

- Determine how it works today in different application domains
- Develop a conceptual framework and vocabulary
- Analyze device interactions from the perspective of a single device
- Analyze combinations of adversary model, capability, resource, protocols, and security goals
- Develop a tool to aid developers

Building on our previous work, we identified the underlying design weaknesses of 4G and 5G cellular paging protocols, and proposed efficient and backward-compatible approaches to address these weaknesses. We also demonstrated the deployment feasibility of our enhanced paging protocol by implementing it on an open-source cellular protocol library and commodity hardware. Our evaluation demonstrates that the enhanced protocol can thwart attacks without incurring substantial overhead.

Our work analyzing the security and privacy issues in the Zigbee protocol continued. In order to design a systematic analysis of the protocol, we first extracted the Finite State Machine (FSM) of the ZigBee end-device, the controller, and the router from the ZigBee specification. We then analyzed the FSM and found a number of potential design vulnerabilities and corresponding exploitations. To address these problems, we developed a new design for the initial joining procedure. We are currently integrating our defense in the open-source implementation of Zigbee stack so that we can evaluate it.

We studied IoT implementations of cryptographic primitives that are critical in secure bootstrapping. We analyzed 15 recent open-source implementations using symbolic execution and found semantic flaws in 6 of them. Further analysis of these flaws showed that 4 implementations are susceptible to new variants of the Bleichenbacher lowexponent RSA signature forgery, which enable an attacker to forge signature on a certificate without breaking the RSA algorithm. One implementation suffers from potential denial of service attacks with purposefully crafted signatures.

Impact on Hard Problems

Policy-Governed Secure Collaboration aims to develop the science underlying methods to express and enforce normative requirements and policies for handling data with differing usage needs and among users in different authority domains. Our research that examines the need for trust in IoT systems deals directly with the challenges associated with this Hard Problem.

Education and Outreach

- Isaac Lammers, a high school student who worked on this project, won multiple awards for his research in this project. His project presentation, titled "WiFi Deauthorization: Pentesting Network Security" won Intel Excellence in Computer Science award.
- We have notified various software vendors and the GSM Association about our findings. GSMA listed us in their Mobile Security Research Hall of Fame at https://www.gsma.com/aboutus/workinggroups/working-groups/fraud-security-group/mobile-security-research-hall-fame

Publications

- Ankush Singla, Syed Rafiul Hussain, Omar Chowdhury, Elisa Bertino, and Ninghui Li, "Protecting the 4G and 5G Cellular Paging Protocols against Security and Privacy Attacks," in *Proceedings on Privacy Enhancing Technologies (PoPETs)*.
- Sze Yiu Chau, Moosa Yahyazadeh, Omar Chowdhury, Aniket Kate, and Ninghui Li, "Analyzing Semantic Correctness with Symbolic Execution: A Case Study on PKCS#1 v1.5 Signature Verification," in proceedings of *2019 Network and Distributed System Security Symposium (NDSS)*.
- Syed Rafiul Hussain, Mitziu Echeverria, Omar Chowdhury, Ninghui Li, and Elisa Bertino, "Privacy Attacks to the 4G and 5G Cellular Paging Protocols Using Side Channel Information," in *Proceedings of 2019 Network and Distributed System Security Symposium (NDSS)*.

Reasoning about Accidental and Malicious Misuse via Formal Methods



Munindar Singh



William Enck



Laurie Williams

PI:	Munindar Singh	CO-PIs:	William Enck , Laurie Williams
URL:	https://cps-vo.org/NCSU_ReasoningAccidentalandMaliciousMisuse		
HARD PROBLEM:	Policy-Governed Secure Collaboration		

Goal

This project seeks to aid security analysts in identifying and protecting against accidental and malicious actions by users or software through automated reasoning on unified representations of user expectations and software implementations to identify misuses sensitive to usage and machine context.

Abstract

This research project deals with accidental and malicious misuse case discovery in sociotechnical systems. System misuse is conceptually a violation of a stakeholder's expectation of how a system should operate. Whereas existing tools make security decisions using context of usage, including environment, time, and execution state, they lack an ability to reason about the underlying stakeholder expectations, which are often crucial to identifying misuses. Our vision is that if existing tools making security decisions could reason about expectations, they could automatically prevent, discover, and mitigate misuse. Unfortunately, automatic extraction of stakeholder expectations remains ineffective.

The foregoing leads us to identify the following research questions: What are the key components of stakeholders' expectations and how may they be represented computationally? How would we identify the relevant stakeholder expectations? In what ways can we employ reasoning about expectations to inform the specification of sociotechnical systems to promote security?

Accomplishments

We have studied these research questions through case studies on mobile applications as a basis for studying accidental and malicious misuse in a practical setting. First, we investigated misuse of credit card information by mobile applications, as defined by Payment Card Industry Data Security Standard (PCI-DSS) standards. In studying the PCI-DSS standard, we identified six restrictions (prohibition norms) that apply to end-user software and codified these restrictions into static program analysis checks. We also built an initial tool called Cardpliance that identifies PCI-DSS misuse violations in Android applications. We have applied Cardpliance to a set of real applications that are known to ask the user for credit card information and found instances of PCI-DSS violations. We also found that most applications that ask the user for credit card information are handling the information correctly.

Second, based on our understanding that user-reported textual artifacts express user expectations, we developed a tool, CASPAR, that understands stories of problems from app reviews. We investigated user reviews of 151 popular apps covering different genres on Apple's App Store. From these app reviews, we extracted 85,099 event pairs describing user-app interactions wherein the apps behave differently from expectations after users' actions. We developed and applied natural language

techniques for event inference on these event pairs to help developers more efficiently address these misbehaviors. Some of these misbehaviors, such as misusing location information, may raise serious security concerns, including enabling misuses such as spying on other people.

Impact on Hard Problems

This project addresses the Hard Problem of policy-governed secure collaboration. Specifically, Carpliance represents a computational implementation of policy checking, and CASPAR is a computational approach to extract violations of stakeholder expectations, thereby making implicit expectations explicit.

Education and Outreach

- Three doctoral students were supported on this project.
- During Summer 2019, we hosted a high school student who learned about violations of stakeholder expectations and their identification from app review

Publications

- None thus far

University of Illinois Urbana-Champaign

The University of Illinois at Urbana-Champaign (UIUC) Science of Security (SoS) Lablet team is led by Principal Investigator (PI) Sayan Mitra and Co-PIs David Nicol and William Sanders. UIUC and its Sub-Lablets, Illinois Institute of Technology and the University of Texas at Austin are engaged in five research projects that leverage UIUC expertise in resiliency, which in this context means a system's demonstrable ability to maintain security properties even during ongoing cyber attacks. The Lablet's work draws on several fundamental areas of computing research as well as ideas from other the mathematics and engineering disciplines.

UIUC research has been published in the 30th International Symposium on Software Reliability Engineering (ISSRE 2019), the European Dependable Computing Conference, ACM Transactions on Modeling and Computer Simulation, the Journal of Modern Power Systems and Clean Energy (MPCE), 2019 IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), Dependable Systems and Networks (DSN), 2019 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (PADS), 2019 International Conference on Omni-Layer Intelligent Systems (COINS), 17th USENIX Symposium on Networked Systems Design and Implementation, 2018 Winter Simulation Conference (WSC), in the book titled Design Automation for Cyber-Physical Systems, and the International Symposium on Automated Technology for Verification and Analysis (ATVA).

In addition to paper presentations at selected conferences mentioned above and ongoing workshops, seminars and talks within the UIUC community, UIUC research has been presented at the 2019 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm), the 10th Anniversary of the Distinguished Colloquium Series of University of Maryland on Verification for safe autonomy: Challenges and recent developments, the Internet of Things (IoT) Systems Research Center at the University of Wisconsin Madison, and the ECE Symposium Lecture, University of Minnesota.


Sayan Mitra


William Sanders


David Nicol

Details on the five research projects identified below can be found in the following sections.

- A Human Agent-Focused Approach to Security Modeling
- An Automated Synthesis Framework for Network Securitand Resilience
- Monitoring, Fusion, and Response for Cyber Resilience
- Resilient Control of Cyber-Physical Systems with Distributed Learning
- Uncertainty in Security Analysis

Special Note: William H. Sanders, the Herman M. Dieckamp Endowed Chair in Engineering at the University of Illinois at Urbana-Champaign and the interim director of the Discovery Partners Institute (DPI) in the University of Illinois System, has been appointed the next dean of the College of Engineering at Carnegie Mellon University, effective January 2020.

ILLINOIS

A Human Agent-Focused Approach to Security Modeling



William Sanders

PI:	William Sanders
URL:	https://cps-vo.org/UIUC_Human-Agent-FocusedApproach-SecurityModeling
HARD PROBLEM:	Human Behavior

Goal

The aim of this project is to make fundamental advances in scientifically-motivated techniques to aid risk assessment for computer security through the development of a general-purpose, easy-to-use formalism. This formalism will allow for realistic modeling of cyber systems and all human agents that interact with the system with the ultimate goal of generating quantitative results that will help system architects make better design decisions.

Abstract

Our hypothesis is that models that incorporate all human agents who interact with the system will produce insightful metrics. System architects can leverage the results to build more resilient systems that are able to achieve their mission objectives despite attacks.

We are currently conducting a literature review with the goal of constructing a high-quality case study to exercise the human-centric cyber security modeling formalism that we are developing. Our case study will focus on comparing the security and usability of different password policies (e.g. password length, time until password expires, etc.) which a hypothetical institution may enact. Our case study will construct submodels of the institutions, its employees and customers, and the adversaries. We shall compose these submodels and study the interaction to give insight into the relative strengths and weaknesses of the password policies. We will validate our model by using previously-conducted studies of human behavior with regard to passwords.

Accomplishments

This project was initiated in March 2019. We have written a position paper on why quantitative cyber security modeling is important and how it should be done. We believe that, once published, the paper will help introduce those who are unfamiliar with modeling to the advantages of quantitative modeling approaches for security, teach them how to do it and common pitfalls to avoid. As a result of our literature review, we have identified a state-of-the-art approach.

Impact on Hard Problems

The Hard Problem of Human Behavior focuses on how to handle the unpredictability and complexity of human actors in cybersecurity. These actors include malicious attackers, system users, and software/system developers. Our hypothesis is that models that incorporate all human agents who interact with the system will produce insightful metrics. System architects can leverage the results to build more resilient systems that are able to achieve their mission objectives despite attacks.

Education and Outreach

- No initiatives thus far

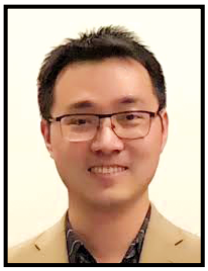
Publications

- No publications thus far

An Automated Synthesis Framework for Network Security and Resilience



Matt Caesar



Dong Jin

PI:	Matt Caesar CO-PI: Dong (Kevin) Jin
Sub-Lablet:	Illinois Institute of Technology
URL:	https://cps-vo.org/UIUC_AutomatedSynthesisFramework
HARD PROBLEM:	Resilient Architectures, Scalability and Composability, Policy-Governed Secure Collaboration, Security Metrics and Models



Goal

We propose to develop the analysis methodology needed to support scientific reasoning about the resilience and security of networks, with a particular focus on network control and information/data flow.

Abstract

The core of this vision is an Automated Synthesis Framework (ASF), which will automatically derive network state and repairs from a set of specified correctness requirements and security policies. ASF consists of a set of techniques for performing and integrating security and resilience analyses applied at different layers (i.e., data forwarding, network control, programming language, and application software) in a real-time and automated fashion. The ASF approach is exciting because developing it adds to the theoretical underpinnings of SoS, while using it supports the practice of SoS.

Accomplishments

As we have done since this project started, we continue to transfer our technology to industry through interactions with Veriflow, a startup company commercializing verification technology that came out of our project's SoS Lablet funding. In September 2019, Veriflow was sold to VMWare, and our technology is slated to be incorporated into VMWare NSX, a very widely used virtualization platform in industry.

In continuing our investigation of automated synthesis of network control to preserve desired security policies and network invariants, we have designed a list of approximately 30 important and useful invariants to showcase the functionality of our system as well as to test it in practical use. We have completed development of our parsing infrastructure, and performed a large-scale evaluation of our approach on real network data.

Our development of a simulation/emulation-based platform for Cyber-Physical System (CPS) resilience and security evaluation remains under way. The platform combines physical computing and networking hardware for the cyber presence while allowing for offline simulation and computation of the physical world. The platform consists of a distributed virtual time kernel module for efficient synchronization between simulation and network emulation/hardware. The kernel module enables processes and their clocks to be paused, resumed, and dilated across embedded Linux devices through the use of hardware interrupts.

We continue to explore self-healing network management to address the resilient architecture hard problem and apply the methods to applications in cyber-physical energy systems, in particular, Phasor Measurement Unit (PMU) networks in the electric grid. We developed an efficient algorithm that considers both power system observability and communication network characteristics. The algorithm also assigns weights to the end-devices according to selected power system metrics before performing the optimization. We are refining our evaluation methods based on feedback from power system engineers as well as the experimental results.

Impact on Hard Problems

Resilient Architectures. Our platform has been integrated into a real production system, VMWare NSX, and thus provides a fundamental advance in industry's ability to construct and maintain highly-resilient cloud architectures with provable guarantees.

Scalability and Composability. Our unique emulation/simulation-integrated testing platform enables scalable evaluation methodologies, and the experimental results indicate the approach scales to large operational environments while providing formal guarantees on correctness of the cyber domain.

Security Metrics and Models. We have designed a set of security metrics and invariants for predicting/enforcing certain network behaviors so that operators can prioritize effort towards the portion of the system that indicate the highest risk.

Education and Outreach

- PI Matt Caesar was selected to become the Director of Education for ACM SIGCOMM.
- PI Matt Caesar has created a new class on Internet of Things. The class contains extensive coverage of security in this important domain. The class is slated for public release to tens of thousands of students on Coursera's Massive Online Open Course (MOOC) platform.
- PI Matt Caesar continues to refine his Networking Laboratory class, targeting release for Fall 2020.
- PI Matt Caesar is currently constructing an online platform for working with IoT devices in the cloud, targeting future release to K-12 institutions across the United States.
- Co-PI Kevin Jin served as the general chair of the 2019 ACM SIGSIM-PADS Conference, and served as a track coordinator of the "Simulation and Cyber Security" track in the Winter Simulation Conference 2019.
- Co-PI Kevin Jin and Dr. Kyle Hale (Illinois Institute of Technology) have developed a new graduate-level cyber security class "CSP544 System and Network Security" for Spring 2020 at Illinois Institute of Technology.
- Co-PI Kevin Jin has been appointed as the Director of the new Master of Cybersecurity Program in the College of Science at Illinois Institute of Technology.
- Co-PI Kevin Jin and Dr. Chen Chen (Argonne National Lab) delivered a tutorial "Electric Power System Resilience" at the 2019 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm) in October 2019.
- Co-PI Kevin Jin and Dr. Philippe J. Giabbanelli (Miami University) organized a Ph.D. colloquium as part of the ACM SIGSIM-PADS conference in June 2019.
- Co-PI Kevin Jin gave a full-day tutorial on "Cyber Security and Resilience of Cyber-Physical Systems" in the Internet of Things (IoT) Systems Research Center at the University of Wisconsin Madison, June 2019

Publications

- Shangyu Xie, Han Wang, Shengbin Wang, Haibing Lu, Yuan Hong, Dong Ji, and Qi Liu, "Discovering Communities for Microgrids with Spatial-Temporal Net Energy," *Journal of Modern Power Systems and Clean Energy (MPCE)*, July 2019
- Jiaqi Yan, Guanhua Yan, and Dong Jin, "Classifying Malware Represented as Control Flow Graphs using Deep Graph Convolutional Neural Network," *2019 IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, June 2019.
- Mohammad Noureddine, Amanda Hsu, Matthew Caesar, Fadi Zaraket, and William Sanders, "P4 AIG: Circuit-Level Verification of P4 Programs," *2019 IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, June 2019.
- Xiaoliang Wu, Jiaqi Yan, and Dong Jin, "Virtual-Time-Accelerated Emulation for Blockchain Network and Application Evaluation," *2019 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (PADS)*, June 2019
- Christopher Hannon, Jiaqi Yan, Dong Jin, and Yuan-An Liu, "A Distributed Virtual Time System on Embedded Linux for Evaluating Cyber-Physical Systems," *2019 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (PADS)*, June 2019, **Best Paper Award**
- Christopher Hannon and Dong Jin, "Bitcoin Payment-Channels for Resource Limited IoT Devices," *2019 International Conference on Omni-Layer Intelligent Systems (COINS)*, May 2019.
- Santhosh Prabhu, Ali Kheradmand, Brighten Godfrey and Matthew Caesar, "Plankton: Scalable Network Configuration Verification through Model Checking," *17th USENIX Symposium on Networked Systems Design and Implementation*, Santa Clara, CA, February 25-27, 2020.
- Christopher Hannon, Nandakishore Santhi, Stephan Eldenbenz, Jason Liu, and Dong Jin, "Just-In-Time Parallel Simulation," *2018 Winter Simulation Conference (WSC)*, Gothernburg, Sweden, December 9-12, 2018.

Monitoring, Fusion, and Response for Cyber Resilience



William Sanders

PI:	William Sanders
URL:	https://cps-vo.org/UIUC_MonitoringFusionandResponse
HARD PROBLEM:	Policy-Governed Secure Collaboration, Resilient Architectures

Goal

The objectve of this project is to develop a scientifically optimized method to deploy sensors, aggregate data and find correlations so that attackers cannot lurk in the shadows of the network. These sensors would also be part of the protected infrastructure.

Abstract

We believe that diversity and redundancy can help us prevent an attacker from hiding all of his or her traces. Therefore, we will strategically deploy diverse security monitors and build a set of techniques to combine information originating at the monitors. We have shown that we can formulate monitor deployment as a constrained optimization problem wherein the objective function is the utility of monitors in detecting intrusions. In this project, we will develop methods to select and place diverse monitors at different architectural levels in the system and evaluate the trustworthiness of the data generated by the monitors. We will build event aggregation and correlation algorithms to achieve inferences for intrusion detection. Those algorithms will combine the events and alerts generated by the deployed monitors with important system-related information, including information on the system architecture, users, and vulnerabilities. Since the rule-based detection systems fail to detect novel attacks, we will adapt and extend existing anomaly detection methods. We will build on our previous SoS-funded work that resulted in the development of the special-purpose intrusion detection methods.

Accomplishments

To support coordinated analysis of heterogeneous monitor data (which spans numerical metrics to unstructured, textual log data) that is present in large-scale distributed systems, such as enterprise and cloud systems, we have developed a framework to semi-automatically process monitor data from multiple levels of said systems into a manageable set of meaningful time series features for further intrusion or incident analysis. Based on an analysis of how incident response teams in industry utilize monitor data, we have come up with a taxonomy of monitor data fields and devised an approach in which we can take monitor data for which fields have been annotated using our taxonomy, automatically unstack and aggregate them into meaningful time series features, and group together redundant features across all monitors in the system. We have evaluated our approach on experimental PaaS cloud data from an industry partner containing eight different monitor types.

We have also begun work on an approach to explicitly protect monitoring infrastructures against monitor compromise through defensive monitor placement, in a manner that does not require changing intrusion and incident analysis mechanisms already in place. We are developing a methodology to examine a system and network monitoring infrastructure in the context of incident and intrusion detection to quantify susceptibility to monitor compromise and suggest changes to the monitoring that can improve the resiliency of the monitoring infrastructure. As part of this work, we are methodically examining the motives an attacker might have to compromise monitors and the types of compromise that can occur, and considering the effects of each type of compromise on incident detection ability.

Impact on Hard Problems

Policy-Governed Secure Collaboration - We analyzed the issues surrounding the Software-Defined Networking (SDN) architecture from an accountability standpoint, considering various principals involved (e.g., controller software, network applications, administrators, end users, organizations), mechanisms for assurance about past network state (e.g., data provenance, replicated data stores, roots of trust), thoughts on judging and assessing standards for accountability (e.g., legal, contractual, regulatory), and mechanisms for decentralized enforcement (e.g., blockchain-based smart contracts). We motivated the need for accountability though a network application use case, and we argued that an assured understanding of the past for attribution can help lead to taking better responses for resiliency.

Resilient Architectures - Experience suggests that even heavily defended systems can be breached by attackers given enough time, resources and talent. We propose the concept of a Response and Recovery Engine (RRE) so that a system could "tolerate" an intrusion and provide a base level of service. RRE incorporates modules to monitor current state of a system, detect intrusions, and respond to achieve resilience-specific goals. Our work focuses on a few example attacks. These attacks include lateral movement within a network as part of an Advanced Persistent Threat (APT) and application-level distributed denial of service attacks (DDoS).

Education and Outreach

- No initiatives thus far.

Publications

- Uttam Thakore, Harigovind Ramasamy, and William Sanders, "Coordinated Analysis of Heterogeneous Monitor Data in Enterprise Clouds for Incident Response," *30th International Symposium on Software Reliability Engineering (ISSRE 2019)*.
- Carmen Cheh, Uttam Thakore, Binbin Chen, William Temple, and William Sanders, "Leveraging Physical Access Logs to Identify Tailgating: Limitations and Solutions," *European Dependable Computing Conference*.
- Carmen Cheh, Uttam Thakore, Ahmed Fawaz, Binbin Chen, William Temple, and William Sanders, "Data-Driven Model-Based Detection of Malicious Insiders via Physical Access Logs," (Extended paper), *ACM Transactions on Modeling and Computer Simulation*.

Resilient Control of Cyber-Physical Systems with Distributed Learning



Sayan Mitra



Geir Dullerud



Sanjay Shakkotai

PI:	Sayan Mitra CO-PIs: Geir Dullerud, Sanjay Shakkottal
SUB-LABLET:	University of Texas at Austin
URL:	https://cps-vo.ort/UIUC_ResilientControlofCPSDistributedLearning
HARD PROBLEM:	Security Metrics and Models, Resilient Architectures

Goal

The goal of the project is to develop rigorous monitoring and verification techniques for resiliency and risk-reduction in autonomous systems and Cyber-Physical Systems (CPS) (e.g., autonomous vehicles, drones, and manufacturing systems). Monitoring can catch problems in deployed systems and avert major failures. Verification solves a similar problem at design time and helps catch bugs and vulnerabilities. We are looking beyond the currently available model-based methods (for monitoring and verification) and developing a detailed understanding of sample complexity (i.e., how many test samples do we need to draw to answer a risk / resiliency question) for these problems. Thus, the project would expand the range of applicability as well as guide optimal resource usage for verification.

Abstract

This project was motivated by the complex interaction of dynamics and decision-making given that the integration of hundreds of components expose CPS to I/O attacks and component compromises. This project addresses resiliency and risk-reduction in CPS through rigorous monitoring and verification, going beyond model-based approaches. In going beyond model-based approaches, the researchers exploit models when it makes sense but otherwise get (gracefully degrading) guarantees with blackbox executables. They also focus on the marginal benefits of models/executable fidelity in security and risk reduction.

Accomplishments

We have formulated a new direction of scientific enquiry into safety and security analysis of systems. The approach relies on distributed and sample-efficient optimization techniques that have been developed in the context of the *Multi-armed bandit problem*. We have shown how these optimization algorithms can be used effectively for statistical model checking of Markov decision processes. We have built a suite of benchmarks related to online safety analysis of autonomous and semi-autonomous vehicles. Our initial results are very promising as the data usage and the running time of our algorithms can be several orders of magnitude better than state-of-the-art model checking approaches such as Storm and Prism.

This research is a part of a bigger effort from our team in

developing algorithms and software tools for Verification and Validation (V&V) of Cyber-Physical Systems. Our group has developed the C2E2 and DryVR tools for model checking CPS. These are two of the leading verification tools today for nonlinear and blackbox CPS and have been used to verify key properties of automotive and aerospace control systems.

Impact on Hard Problems

Metrics: How much data is necessary to achieve a certain level of confidence regarding a safety/security claim

Resiliency: Effective verification of safety and security properties of autonomous and cyber-physical systems

Education and Outreach

- PI Sayan Mitra gave an invited lecture at the 10th Anniversary of the Distinguished Colloquium Series of University of Maryland on "Verification for safe autonomy: Challenges and recent developments," May 3rd, 2019.
- Saturday Engineering for Everyone (SEE) Talk on "Future of Safe Autonomous Vehicles" given by PI Sayan Mitra on March 2nd. This outreach event is hosted by the Department of Electrical and Computer Engineering and attended by members of the broad Urbana-Champaign community.
- ECE Symposium Lecture, University of Minnesota, "Statistical Validation and Principle-Based Simulation of Complex Cyber-Controlled Systems," Co-PI Geir Dullerud, May 5th, 2019.
- CES Symposium Lecture, University of Texas, Austin, "Statistical Validation and Principle-Based Simulation of Complex Cyber-Controlled Systems," Co-PI Geir Dullerud, December 3rd, 2019.
- ECE CSP Seminar, University of Michigan at Ann Arbor, "Hyper-parameter Tuning for ML Models: A Monte-Carlo Tree Search (MCTS) Approach," lecture given by Co-PI Sanjay Shakkottai, October 24, 2019.
- EE Seminar, IIT Bombay, "Hyper-parameter Tuning for ML Models: A Monte-Carlo Tree Search (MCTS) Approach," lecture given by Sanjay Shakkottai, December 17, 2019.
- Invited Speaker, Formal Methods in Mathematics Workshop, Mathematics, Carnegie Mellon University, "Learning and Statistical Validation of Complex Cyber-Physical Systems," Co-PI Geir Dullerud, January 7th, 2020 on logic in engineering systems.
- PI Mitra's new course, Principles of Safe Autonomy, at UIUC came to a successful conclusion. This unique class takes a deep dive into the seminal topics in object recognition, learning, localization, decision-making, path planning, control, and safety verification. Twenty-five students from ECE and CS completed the course and participated in a competition sponsored by Galois Inc. The course team designed six new programming assignments involving topics such as lane detection, road-sign recognition with deep neural networks, localization with particle filters, decision-making with reinforcement learning, path planning with rapidly expanding random trees, and safety verification using simulation-driven proofs. The next edition of this course will be offered in Spring 2020 and students will get an opportunity to deploy their code on a real autonomous vehicle.
- Chuchu Fan from PI Sayan Mitra's group finished her PhD in November of 2019 and will be joining MIT's Aerospace department as a tenure track Assistant Professor in Fall of 2020.
- Collaborations with AFRL, UTC Aerospace, and Boeing on the research are ongoing.

Publications

- Ronshee Chawla, Abishek Sankararaman, Ayalvadi J. Ganesh and Sanjay Shakkottai, "The Gossiping Insert-Eliminate Algorithm for Multi-Agent Bandits," in *23rd International Artificial Intelligence and Statistics (AISTATS 2020)*, Palermo, Sicily, June 2020, to appear.
- Hussein Sibai, Navid Mokhlesi and Sayan Mitra, "Using symmetry transformations in equivariant dynamical systems for their safety verification," in proceedings of the *Seventeenth International Symposium on Automated Technology for Verification and Analysis (ATVA),* October 28-31, 2019, Taipei City, Taiwan. **Nominated for best paper award.**
- Chuchu Fan and Sayan Mitra, "Data-driven safety verification of complex cyber-physical systems," A chapter in the book titled Design Automation for Cyber-Physical Systems, edited by Mohammad Abdullah Al Faruquqe and Arquimedes Canedo, pages 107-143, Springer, 2019.
- Yu Wang, Nima Roohi, Mahesh Viswanathan, Matt West, and Geir Dullerud, "Statistical verification of PCTL using antithetic and stratified sampling," Formal Methods in Systems Design, volume 54, pages,145-163, Springer, 2019.

Uncertainty in Security Analysis



David Nicol

PI:	David Nicol
URL:	https://cps-vo.org/UIUC_UncertaintyinSecurityAnalysis
HARD PROBLEM:	Security Metrics and Models, Resilient Architectures

Goal

The goal of this project is to develop a mathematical basis for describing and analyzing the ability of an adversary to laterally traverse networks in the presence of uncertainty about connections and uncertainty about exploitable vulnerabilities. We will use this basis to develop algorithms for quantified risk analysis of cyber-physical systems.

Abstract

Cyber-security vulnerabilities in Cyber-Physical Systems (CPS) allow an adversary to remotely reach and damage physical infrastructure. Following the initial point of entry, the adversary may move laterally through the computer network using connections that are allowed by the access control but which give access to services with exploitable vulnerabilities. Using lateral movement, the adversary may eventually have control of monitors and actuators in the CPS, corrupt data being reported and/or issue malicious control commands the consequences of which may inflect significant damage. Analyses of the risk of such attacks are known, under the assumption that all vulnerabilities and all connections in the cyber-system are known perfectly. They aren't. We are interested in developing the mathematical basis for describing the ability of the adversary to reach actuators in the presence of uncertainty with respect to the connections and the vulnerabilities which enable lateral movement.

Edges derived from topological analysis may be thought to have "exploitation probabilities" which quantify with a single probability the possibility of the adversary traversing that edge in a lateral movement. An edge probability models the possibility of an adversary on one host A being able to connect to another host B and exploit a vulnerability there, enabling the adversary to launch further attacks from B. In a previous study, we used expressions of Boolean random variables to describe these probabilities, in order to be able to escape the otherwise necessary assumption of independence among edges. Since edges quantify the likelihood of an adversary exploiting a vulnerability, distinct edges which describe the same vulnerability will not have independent probabilities. Using Boolean expressions enables us to describe those correlations we know must exist. Our current investigation generalized this model. Point probabilities implicitly assert certainty *in the probability*. There are different reasons why an exploitation probability may be non-zero, and some uncertainty in knowing just what causes the variability in connection. The new work replaces edge probabilities with edge probability distributions, which allows us greater flexibility in expression the certain or uncertainty of the edge probability. Using the beta distribution, one set of parameters can create a very spiked distribution centered on the mean, while another set of parameters can create a distribution with the same mean, but with the probability mass distributed so as greater variance is captured. However, while beta distributions are closed under some operations they are not closed under others, and so the new work considers how to compute the parameters of approximated betas with good accuracy.

Accomplishments

In investigating the potential computational benefit of the Beta approximation result, we showed that using a small number of samples (less than a few thousands), computing the parameters of an approximating Beta yields a significantly better estimate of the reliability distribution than constructing the empirical distribution.

We used the proposed model to study the reliability of two realistic systems, a distributed system with redundant deployment and a gas distribution network.

We completed the evaluation of the simulation-based experiments. Numerical results from Monte Carlo simulation of an approximation scheme and from two case studies strongly support the observations made above, especially for non-corner cases where the model parameters do not take extreme values.

We are preparing for the next phase of the project which is to embed the developed security model in a risk assessment framework. More specifically, we are systematically surveying the literature on risk assessment in SCADA systems to identify the common approaches, their abilities and limitations. The lessons learned will help us build our own risk assessment framework.

Impact on Hard Problems

This research intersects the predictive security *metric* problem since we are attempting to predict uncertainty associated with a system model. It also intersects with resilience as a system's resilience will be established by analysis of some model and decisions (e.g., how significant the breach may be, whether to interdict and where, where to focus recovery activity) will be made as a result. Those decisions will be better informed when some notion of uncertainty is built into the model predictions, or accompanies those model predictions.

Education and Outreach

- No initiatives thus far

Publications

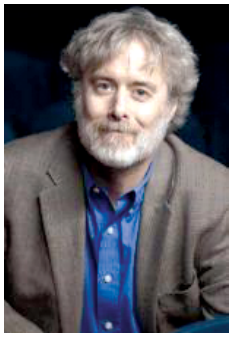
- Hoang Hai Nguyen, Kartik Palani, and David Nicol, "Extensions of Network Reliability Analysis," *49th IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2019)*, Portland, OR, June 24-27, 2019.

University of Kansas

The University of Kansas (KU) Science of Security and Privacy Lablet and its Sub-Lablet the University of Tennessee are making interdisciplinary contributions to security science, synthesizing knowledge and innovation from computer science, electrical engineering, psychology, sociology, and philosophy. The Lablet's work focuses on the foundational nature of resiliency, defining and establishing trust, understanding privacy in IoT architectures, understanding and preventing side-channel attacks, and developing techniques for secure, native binary execution. In all areas the Lablet seeks foundational solutions rooted in formal mathematical analysis and empirical scientific study. The interface between analytical and experimental research promises a broad basis for understanding security problems and solutions. Applications are drawn primarily from Cyber Physical Systems (CPS) and Internet of Things (IoT) where proliferation and rapid change present increasingly difficult security problems.

The KU Lablet continued five projects in 2019. The first project, **Cloud-Assisted IOT Systems Privacy**, is developing a method to enable cloud-assisted, privacy-preserving machine learning classification over encrypted data for IoT devices. Under the **Formal Approaches to the Ontology and Epistemology of Resilience** project researchers are developing an epistemology and ontology for framing resilience. The third project, **Side-Channel Attack Resilience**, is focused on reducing micro-architectural side-channels by introducing new OS abstractions while minimally modifying micro-architecture and OS. The **Secure Native Binary Executions** project seeks to design a mechanism to both analyze and quantify the level of security provided and performance penalty imposed by different solutions to various security risks affecting native binaries. Finally, the **Scalable Trust Semantics and Infrastructure** project is developing a formal basis for remote attestation that supports scalability to distributed, heterogeneous systems and demonstrating the sufficiency and soundness of remote attestation processes.

This year KU research has been published in the International Information Security and Privacy Conference (IFIP SEC), Privacy Enhancing Technologies Symposium (PETS), the ACM Journal of Data and Information Quality (JDIQ), IEEE Conference on Communications and Network Security, Security and Privacy in Communication Networks (SecureComm), ACM Conference on Computer and Communications Security (CCS), the APA Newsletter on Philosophy and Computers, Emerging Ubiquitous Systems and Pervasive Networks (EUSPN), International Workshop on Resilient Networks Design and Modeling (RNDM), Foundations of Science, Philosophy and Technology, Minds and Machines, Stanford Encyclopedia of Philosophy, Hot Topics in the Science of Security (HoTSoS), Principles of Security and Trust (POST), IEEE Real Time and Embedded Technology and Applications Symposium (RTAS), and the IEEE/ACM Design Automation Conference (DAC).



Perry Alexander

The KU Lablet is housed in The Information and Telecommunications Technology Center (ITTC), an interdisciplinary research center focused on all aspects of information and its applications. ITTC researchers working on Lablet projects are drawn from across the University including:

- The Department of Electrical Engineering and Computer Science
- The Department of Philosophy
- The Department of Psychology
- The Department of Sociology

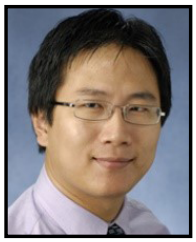
The KU Lablet is led by Perry Alexander who serves as Principal Investigator. He is assisted by Patricia Bergman who provides research and administrative support. Lead project PIs include Bo Luo, Fengjun Li, John Symons, Heechul Yun, Garrett Morris, and Prasad Kulkarni. The KU Sub-Lablet at the University of Tennessee is led by Michael Jantz. Approximately 12 graduate students and 7 faculty are supported by the program assisted by several research staff and undergraduate students. The Lablet team meets regularly with an industrial advisory board that provides input on research directions, facilitates outreach and helps establish a regional, midwestern cyber security community.

KU

Cloud Assisted IoT Systems Privacy



Fengjun Li



Bo Luo

PI:	Fengjun Li	CO-PI: Bo Luo
URL:	https://cps-vo.org/KU_Cloud-AssistedIoTSystemsPrivacy	
HARD PROBLEM:	Scalability and Composability, Security Metrics and Models, Resilient Architectures	

Goal

The goal of this project is to develop principles and methods to model privacy requirements, threats, and protection mechanisms in cloud-assisted IoT systems. The work aims to address the hard problems of resilient architectures, security metrics, and scalability and composability.

Abstract

The key to realizing the smart functionalities envisioned through the Internet of Things (IoT) is to securely and efficiently communicate, maintain, and analyze the tremendous data generated by IoT devices. Therefore, integrating IoT with the cloud platform to utilize its computing and big data analysis capabilities becomes critically important, since IoT devices are computational units with strict performance and energy constraints. However, when data is transferred among connected devices or to the cloud, new security and privacy issues arise. In this project, we investigate the privacy threats in the cloud-assisted IoT systems, in which heterogeneous and distributed data are collected, integrated and analyzed by different IoT applications. The goal of the project is to develop a privacy threat analysis framework for modeling privacy threats in the cloud-assisted IoT systems and provide a holistic solution towards privacy protection.

Accomplishments

The number of IoT devices is expected to reach 125 billion by 2030. These devices collect a variety of data that may contain privacy-sensitive information of the users. However, it is difficult to quantitatively assess privacy leakage of a text snippet. To address this problem, we are among the first to develop a context-aware, text-based quantitative model for private information assessment. We developed a computational framework using natural language processing and deep neural networks to train prediction models, which can be used to measure privacy scores of texts from a social network dataset. These models serve as the foundation for developing user alerting mechanisms to warn users when they attempt to disseminate sensitive information online.

To support big data analytics over IoT-generated data while protecting IoT data privacy, we developed two privacy-preserving learning frameworks for collaborative learning tasks in cloud-assisted IoT systems. First, in the scenarios where clients (e.g., IoT devices) hold horizontally partitioned data (i.e., data with the same distribution or feature space), we developed a privacy-preserving incremental learning protocol. It allows clients to train support vector machines (SVM) using a linear, polynomial kernel function and securely update the SVM model by exchanging the kernel matrix using homomorphic encryption. As IoT devices are continuously generating new data, this privacy-preserving incremental learning scheme can actively update the learning model by integrating the new data into the quadratic program and modifying the kernel parameters when necessary, and, therefore, is critical to predication tasks in IoT applications.

We also designed a blockchain-based framework to support privacy-preserving and accountable federated learning tasks in IoT applications. As an emerging collaborative learning technique, Federated Learning (FL) allows learning of high-quality, personalized models from data across distributed sources, without transmitting the raw data from the devices to anywhere else. This facilitates large-scale collaborative IoT applications by not only leveraging computational resources on IoT devices but also improving user privacy protection. To

prevent privacy leakage, model updates (e.g., updated stochastic gradient descents) are securely exchanged among the clients using homomorphic encryption, multi-party computation, and differential privacy schemes. As the IoT devices are loosely federated in the FL tasks, some of them may be malicious by cheating or colluding in model update. We developed a blockchain-based framework to provide an accountable federated learning service by leveraging the immutability and decentralized trust properties of the blockchain to provide provenance of model update.

Impact on Hard Problems

Successful completion of this project will result in: 1) a systematic methodology to model privacy threats in data communication, storage, and analysis processes in IoT applications; 2) a privacy threats analysis framework with an extensive catalogue of application-specific privacy needs and privacy-specific threat categorization; and 3) a privacy protection framework that maps existing Privacy Enhancing Technologies (PETs) to the identified privacy needs and threats of IoT applications to simplify the selection of sound privacy protection countermeasures.

Education and Outreach

- We continued the bi-weekly Privacy Study Group in KU. Regular attendees include Lablet PIs and graduate students, students from computer science and philosophy departments, and researchers from Kansas State University.
- Co-PI Bo Luo served as a Grand Award Judge of the 2019 Intel International Science and Engineering Fair, one of the world's largest STEM competitions, May 12-17, 2019.
- Co-PI Bo Luo presented "Arcana: Enabling Private Posts on Public Microblog Platforms" at the IFIP SEC Conference in Lisbon, Portugal, June 25, 2019.
- Co-PI Bo Luo gave an invited talk on "From Identity to Content: Privacy Protection in Social Networks" at the International Workshop on Cyber Security and Data Privacy in Hangzhou, China, July 6, 2019.
- PI Fengjun Li gave an invited talk on "Protecting Vulnerable IoT Devices via Multipath Onion IoT Gateways" at the International Workshop on Cyber Security and Data Privacy, Hangzhou, China, July 6, 2019.
- Co-PI Bo Luo presented "#DontTweetThis: Scoring Private Information in Social Networks" at 2019 PETS, Stockholm, Sweden, July 15, 2019.
- PI Fengjun Li gave an invited talk on "Cloud-Assisted Secure Computing for IoT Applications" at the University of South Florida, October 23, 2019.
- PI Fengjun Li presented "Ticket Transparency: Accountable Single Sign-On with Privacy-Preserving Public Logs" at SecureComm, Orlando, October 24, 2019.
- PI Fengjun Li served the panelist role in the panel on "Unique Privacy and Security Challenges in IoT Devices" at SecureComm, Orlando, October 24, 2019.
- PI Fengjun Li helped to organize the 2019 Central Area Networking and Security Workshop in Iowa State University and served as the Travel Award Chair, October 26, 2019.
- PI Fengjun Li presented "A Reliable and Accountable Privacy-Preserving Federated Learning Framework using the Blockchain" at ACM CCS, London, UK, Nov. 12, 2019.

Publications

- Anirudh Narasimman, Qiaozhi Wang, Fengjun Li, Dongwon Lee, and Bo Luo, "Arcana: Enabling Private Posts on Public Microblog Platforms," *34th International Information Security and Privacy Conference (IFIP SEC)*, Lisbon, Portugal, June 25-27, 2019.
- Qiaozhi Wang, Hao Xue, Fengjun Li, Dongwon Lee, and Bo Luo, "#DontTweetThis: Scoring Private Information in Social Networks," *19th Privacy Enhancing Technologies Symposium (PETS)*, vol 4, Stockholm, Sweden, July 15-16, 2019.
- Hao Xue, Qiaozhi Wang, Bo Luo, Hyunjin Seo, and Fengjun Li, "Content-Aware Trust Propagation Towards Online Review Spam Detection," in *ACM Journal of Data and Information Quality (JDIQ) - Special Issue on Combating Digital Misinformation and Disinformation*, Volume 11, Issue 3, 2019.
- Dawei Chu, Jingqiang Lin, Fengjun Li, Xiaokun Zhang, Qiongxiao Wang and Guangqi Liu, "Ticket Transparency: Accountable Single Sign-On with Privacy-Preserving Public Logs," *15th EAI International Conference on Security and Privacy in Communication Networks (SecureComm)*, Orlando, October 23-25, 2019.
- Bingyu Li, Jingqiang Lin, Fengjun Li, Qiongxiao Wang, Qi Li, Jiwu Jin, and Congli Wang, "Certificate Transparency in the Wild: Exploring the Reliability of Monitor," *ACM CCS*, London, UK, Nov 11-15, 2019. (full paper)
- Xuhua Bao, Xiaokun Zhang, Jingqiang Lin, Dawei Chu, Qiongxiao Wang, and Fengjun Li, "Towards the Trust-Enhancements of Single Sign-On Services," *IEEE Conference on Dependable and Secure Computing (DSC)*, Hangzhou, China, Nov. 18-20, 2019.
- Dawei Chu, Kaijie Zhu, Quanwei Cai, Jingqiang Lin, Fengjun Li, Le Guan, and Lingchen Zhang, "Secure Cryptography Infrastructures in the Cloud," *IEEE Global Communications Conference (GLOBECOM)*, Waikoloa, Hi, USA, Dec. 9-13, 2019.

Formal Approaches to the Ontology and Epistemology of Resilience



John Symons

PI:	John Symons
URL:	https://cps-vo.org/KU_FormalApproachesOntologyofResilience
HARD PROBLEM:	Resilient Architectures

Goal

Successful completion of this research effort will result in principled and formally tractable ways to think about the differences between:

- Conditions for the individuation of systems
- Conditions for the identification of systems
- Properties that contribute to the persistence of systems
- Properties that contribute to the functional reliability of systems

Abstract

Security Science requires reflection on its foundational concepts. Our contention is that in order to make informed decisions about trade-offs with respect to resilient properties of systems, we must first precisely characterize the differences between the mechanisms underlying valuable functions, those functions themselves, and the conditions underlying the persistence of the systems in question.

When we say that a system persists, we can mean a variety of things. If we consider an electrical power system or a communications network, for example, our initial evaluation of persistence might involve deciding whether or not the system continues to function: Is the grid continuing to deliver power where it is needed? Is it still possible to send and receive messages reliably through the communications network? This is a functional account of the individuation of systems. The functional account is foundational to contemporary thinking in the science of security. While it is an intuitively sensible and pragmatically grounded way of thinking about systems, it does not shed light on the question of resilience. Functions are also difficult to capture in a purely network theoretic strategy for reasons that this research group will explore and explain.

In order to understand why some systems are resilient and others are not we propose to apply existing work in philosophy of science and metaphysics.

Accomplishments

Our accomplishments over the past year are reflected in the education and outreach efforts and the publications shown below.

Impact on Hard Problems

Security Science has focused on network-based measures of resilience. This is a valuable formal approach, but its range of application is narrower than the general problem requires. In order to make progress on these questions, a broader theoretical approach is required, and we will need to call on a range of other formal and informal methods.

Education and Outreach

- Ramon Alvarado and John Symons "Epistemic Injustice in Machine Learning Technology," The Society of Learning Algorithms, Conference series: Science and Art of Simulation (SAS) High-Performance Computing Center, Stuttgart, Germany
- Ramon Alvarado and John Symons "Epistemic Injustice in the Development and Deployment of Data Science," 5th International Conference on the History and Philosophy of Computing, Bergamo, Italy
- Ramon Alvarado and John Symons "Epistemic Injustice in Data Science," Surveillance Technology and Bias Panel at the Institute for Digital Research in the Humanities 9th annual Digital Humanities Forum, University of Kansas
- John Symons, "Resilience and the Limitations of Graph Theoretical Network Analysis," Science of Security Summer Quarterly
- Amir Modarresi, John Symons, and James Sterbenz, "A Multidimensional Multilevel Model for Smart Environment," 6th Annual Symposium on Hot Topics in the Science of Security (HotSoS)
- We hosted a conference, "Computational Methods and The Future of Science" at the University of Kansas.
- The graduate seminar on Emergence and Resilience enrolled students from engineering, linguistics, psychology, medicine, music, and philosophy. It was a very highly rated course, and in addition to graduate student participation, faculty from the computer science and psychology departments also attended the seminar

Publications

- John Symons, "Should physical computation be understood mechanistically?" *The APA Newsletter on Philosophy and Computers*
- Amir Modarresi and John Symons, "Modeling and Graph Analysis for Enhancing Resilience in Smart Homes," *EUSPN 2019*, Coimbra, Portugal
- Amir Modarresi and John Symons, "Modeling Technological Interdependency in IoT - A Multidimensional and Multilayer Network Model for Smart Environments," *RNDM*, 2019
- John Symons and Jack Horner, "Why there is no general solution to the problem of software verification," *Foundations of Science*
- John Symons and Jack Horner, "Understanding Error Rates in Software Engineering: Conceptual, Empirical, and Experimental Approaches," *Philosophy and Technology*
- John Symons and Ramon Alvorado, "Epistemic Entitlements and the Practice of Computer Simulations," *Minds and Machines*
- Rasmus Rendsvig and John Symons, "Epistemic Logic," *Stanford Encyclopedia of Philosophy*

Scalable Trust Semantics and Infrastructure



Perry Alexander



Garrett Morris

PI:	Perry Alexander	CO-PI: Garrett Morris
URL:	https://cps-vo.org/KU_ScalableTrustSemanticsandInfrastructure	
HARD PROBLEM:	Scalability and Composability, Policy-Governed Secure Collabortion	

Goal

Remote attestation has enormous potential for establishing trust in highly distributed IoT and cyber-physical systems. However, significant work remains to define a unifying semantics of remote attestation that precisely defines guarantees that scales to large, heterogeneous systems Successful completion of this project will result in a science of trust and remote attestation, and prototype infrastructure for building remote attestation systems.

Abstract

Remote attestation provides boot and run-time capabilities for attesting to system behavior and establishing trust. When using remote attestation an appraiser requests evidence from a target that responds by performing measurement to gather evidence and adding cryptographic signatures to assure integrity and authenticity. The appraiser receives evidence and appraises the it to determine if the target is who it claims to be and is behaving as expected.

Accomplishments

KU researchers developed an initial specification of an attestation monad that describes the formal semantics of attestation. The attestation monad is defined as a state monad with exceptions that provides a stateful environment for executing Copland phrases. The attestation monad serves as the basic building block for constructing various elements of remote attestation systems. Specified in Coq, the attestation monad has a well-defined formal semantics that defines requirements for specifying and verifying systems.

Researchers began implementation and verification of an Attestation Virtual Machine (AVM) that serves as a compilation target for Copland protocols. Current attestation managers interpret the protocol directing making it difficult to associate access control policies with each attestation service. The virtual machine resembles a JVM implementing an instruction set for attestation. We are currently verifying the compiler from Copland to the AVM in Coq and verifying the AVM captures Copland event semantics. Of particular interest is interfacing with the maat attestation system developed by our collaborators at JHUAPL under NSA direction. This will allow us to take full advantage of the measurement and run-time management infrastructure they develop.

Our work on negotiating attestation protocols continued with significant progress related to extending ISAKMP to protocol negotiation and establishing ordering and selection criteria for protocols. ISAKMP is an IETF standard for establishing a security association between communicating systems. It establishes common cryptography algorithms and sets up keys to initiate secure communication. ISAKMP can be extended to include negotiation of attestation protocols by mapping an attestation request to a collection of candidate protocols by the attestation target, selection of a "best" protocol by the appraise, and establishing common criteria for encryption, hashing and signing. By layering our negotiation process on ISAKMP we avoid re-building the negotiation infrastructure.

With our collaborators at JHUAPL, MITRE, and NSA, we have started defining an ontology of attestation protocol shapes representing common attestation activities. This ontology is critical to scaling attestation to the enterprise as it provides building blocks for constructing systems at a higher level of abstraction. Shapes include layering attestation, delegation attestation and appraisal, Kerberos-style attestation, and caching.

We now have Attestation Managers running under Windows, Linux, seL4 and MacOS enabling scaling to heterogeneous infrastructure. The Windows implementation uses a different code base to support interfacing to the TPM 2.0, but the Linux, seL4 and MacOS implementations all share a common CakeML codebase. Our goal is formally verifying the CakeML implementation as we extend it to include negotiation and a broader set of measurement techniques.

Impact on Hard Problems

- Semantics of trust: Definitions of trust and metrics for soundness of evaluation and appraisal
- Semantics of measurement, attestation and appraisal: Metrics for soundness and sufficiency of evidence, semantic mechanisms for identity and attestation, formal definitions of evidence, and meta-evidence appraisal
- Systematic mechanisms for establishing roots of trust: Metrics for evaluating roots of trust and general mechanisms for establishing roots of trust on cyber-physical systems
- Attestation protocol representation and semantics: Formal, executable representations for attestation protocols and tools for static analysis
- Implementing and scaling trust infrastructure: Hierarchical frameworks for trust infrastructure including virtualized TPM implementations, trust aggregation and trust as a service

Education and Outreach

- PI Perry Alexander, "Remote Attestation and Trust," invited presentation at University Blockchain Research Initiative Connect (UBRICon'19), University of California at Berkeley, Berkeley, CA.
- PI Perry Alexander, "Trust and Proof," invited panel presentation at Blockchain KC Conference, Kansas City, MO, April 24, 2019.
- PI Perry Alexander, "A Science of Attestation," keynote presentation at Cyberattacks, Intellectual Property, and University Open Research: The Perfect Storm}, ICCAE Spring Colloquium, University of Kansas, Lawrence, KS, April 23, 2019.
- Collaborators presented papers on the Copland semantics for attestation protocols at HoTSoS'19 and POST'19.

Publications

- Adam Petz and Perry Alexander, "A Copland Attestation Manager," *Hot Topics in Science of Security (HoTSoS'19)*, Nashville, TN.
- John Ramsdell, Paul Rowe, Perry Alexander, Sarah Helble, Pete Loscocco, Aaron Pendergrass, and Adam Petz, "Orchestrating Layered Attestations," *Principles of Security and Trust (POST'19)*, Prague, Czech Republic.

Secure Native Binary Execution



Prasad Kulkarni

PI:	Prasad Kulkarni
SUB-LABLET:	University of Tennessee
URL:	https://cps-vo.org/KU_SecureNativeBinaryExecution
HARD PROBLEM:	Security Metrics and Models, Scalability and Composability

Goal

This goal of this research is to build tools and techniques that will allow users to know the security level of their packaged software and enable them to add security to it. Our overall project aim is to provide greater control to the end-user to actively assess and secure the software they use.

Abstract

Typically, securing software is the responsibility of the software developer. The customer or end-user of the software does not control or direct the steps taken by the developer to employ best practice coding styles or mechanisms to ensure software security and robustness. Current systems and tools also do not provide the end-user with an ability to determine the level of security in the software they use. At the same time, any flaws or security vulnerabilities ultimately affect the end-user of the software.

Our research goal is to develop a high-performance framework for client-side security assessment and enforcement for binary software. Our research is developing new tools and techniques to: (a) assess the security level of binary executables, and (b) enhance the security level of binary software, when and as desired by the user to protect the binary against various classes of security issues. Our approach combines static and dynamic techniques to achieve efficiency, effectiveness, and accuracy.

Accomplishments

With respect to developing the framework to update any binary with additional security checks, we completed our initial implementation of the SoftBound technique using static analysis information gathered from Ghidra and dynamic instrumentation using Pin. The SoftBound technique can provide complete spatial memory safety for code written in unsafe languages, like C and C++. The original implementation of SoftBound relied on high-level program information available in the source code and additional checks generated and inserted by the compiler. Instead, our implementation gathers the static information through Ghidra and inserts the checks at run-time using Pin, and does not require the source code. While our first implementation maintains program symbol information in the binary executable (-g flag in GCC) to provide best case static analysis data, we will remove this limitation in later iterations of this work. We are studying the data from our implementation to understand the following issues: a) How accurate and complete is the static analysis conducted by state-of-the-art reverse engineering tools, like Ghidra, with and without program symbol information? b) What is the comparative overhead of our approach? In the future, we will continue to study and understand the reasons when Ghidra is unable to adequately reverse engineer necessary program information, develop techniques to resolve these issues so our defense works in all cases, and reduce run-time overhead of binary-level client-side defense techniques.

In addressing the framework to determine the security level of a given binary executable, we are exploring a fundamental question: how can we assess if a given binary program is susceptible to certain classes of attacks? We are exploring mechanisms based on signature/code-pattern detection in the static binary code and orchestrating attacks (using the Pin instrumentation engine)

at execution-time that would be caught by a defense, if it was implemented.

Work on the framework to reduce run-time overhead for our technique is being led by our Sub-Lablet, the University of Tennessee. The research team has been developing a custom binary instrumentation tool to reduce the execution time overhead of analyzing and performing security checks in programs where source code is not available. They have developed a prototype version of this tool and have used it to instrument the entire set of memory instructions in a small set of custom programs. They are currently validating its correctness and performance on a larger set of standard benchmark programs (SPEC CPU 2017). When the implementation is complete and has been validated, they will adapt it for use with the dynamic security instrumentation that has been developed at KU, and evaluate the performance of this approach compared to instrumentation with the standard Pin framework.

Impact on Hard Problems

Security Metrics and Models: Our research develops quantitative measures to determine the security level of a binary executable. Researchers have classified security vulnerabilities and attacks into classes, with some attacks easier to launch or more dangerous than others. A particular security solution can eliminate susceptibility to one or multiple classes of attacks. The security level of a given binary executable will depend on the number, ratio, and types of attack classes eliminated due to the security measures present in the hardened binary.

Scalability and Composability: To address the important issue of scalability, our solution employs both static and dynamic components. Static analysis of the binary happens offline and before execution, allowing it more time to extract the necessary information. The dynamic component uses this static information to minimize run-time overhead. Program properties that either will take too long or cannot be determined along all static program paths will be resolved only along the executed program paths at run-time.

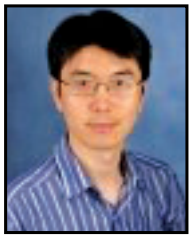
Education and Outreach

- A paper on which this work was based (Surya Tej Nimmakayala and Prasad Kulkarni, "Improving Startup Performance in Dynamic Binary Translators,") was published in the 27th EuroMicro Conference on Parallel, Distributed and Network-Based Processing, Pavia, Italy, Feb 2019.

Publications

- None thus far

Side-Channel Attack Resilience



Heechul Yun

PI:	Heechul Yun
URL:	https://cps-vo.org/KU_Side-ChannelAttackResistance
HARD PROBLEM:	Resilient Architectures

Goal

Successful completion of this project will result in empirical studies on micro-architectural side-channels in safety-critical CPS and criticality-aware OS and architecture prototypes for side-channel attack resistant CPS.

Abstract

Cyber-Physical Systems (CPS)--cars, airplanes, power plants, etc.--are increasingly dependent on powerful and complex hardware for higher intelligence and functionalities. However, this complex hardware may also introduce new attack vectors--hardware side-channels--which can be exploited by attackers to steal sensitive information, to disrupt timing of time-critical functions that interact with the physical plants, or to break memory protection mechanisms in modern computers. Because these attacks target hardware, even logically safe and secure software such as a formally verified OS, could still be vulnerable. Given the safety-critical nature of CPS, hardware side-channels should be thoroughly analyzed and prevented in CPS. This project focuses on micro-architectural side channels in embedded multicore computing hardware, and aims to develop fundamental OS and architecture designs that minimize or completely eliminate the possibility of potential hardware-level side-channel attacks. In this project, we aim to fundamentally reduce or completely eradicate these micro-architectural side-channels by introducing new OS abstractions and minimally modifying micro-architecture and OS.

Accomplishments

In a paper entitled "Denial-of-Service Attacks on Shared Cache in Multicore: Analysis and Prevention," we demonstrated the feasibility and severity of micro-architectural DoS attacks on shared caches in widely used contemporary COTS multicore processors. The paper won an *Outstanding Paper Award* at IEEE Real Time and Embedded Technology Applications Symposium (RTAS 2019). We released the code to reproduce the results at https://github.com/mbechtel2/CacheDOS.

We have developed a comprehensive OS-level scheduling framework, which we call RT-Gang, to mitigate timing related micro-architectural DoS attacks on multicore platforms. The work was also published at IEEE RTAS 2019 and the code was released as open-source at https://github.com/CSL-KU/RT-Gang. We are currently extending the capability of the RT-Gang framework to improve real-time schedulability and isolation guarantees.

We integrated a quad-core RISC-V SoC and an NVDLA DNN accelerator on Amazon FPGA cloud environment and plan to use this platform for our micro-architectural side-channel research in the future. The integration and some preliminary results were published at the EMC^2 workshop. The integrated RISC-V SoC testbed, which we call FireSim-NVDLA, has received significant attention from industry practitioners and academic researchers. For example, Nvidia revealed in their official developer blog (https://devblogs.nvidia.com/nvdla/) that it uses FireSim-NVDLA to evaluate the software release for their open-source deep neural network hardware accelerator called Nvidia Deep Learning Accelerator (NVDLA).

SpectreGuard, a software/hardware collaborative defense mechanism against Spectre attacks, we successfully developed, leverages software-provided information to mitigate Spectre attacks at low hardware and performance cost. The work was published at ACM/IEEE DAC 2019 and the code was released as open-source at https://github.com/CSL-KU/SpectreGuard.

We have also successfully developed a small hardware unit, which we call Bandwidth Regulation Unit (BRU), that regulated

memory traffic at the on-chip interconnect level within the RISC-V multicore. The work was accepted for publication at IEEE RTAS 2020. We are currently working on to prepare a camera-ready version of the paper and open-source release of the BRU.

Impact on Hard Problems

Our paper referenced above ("Denial-of-Service Attacks on Shared Cache in Multicore: Analysis and Prevention") was one of three papers that won the Outstanding Paper Award at IEEE RTAS 2019, a premiere conference in the embedded/real-time systems area. This award is an indication of the importance and impact of our work, and more generally of micro-architectural attacks on safety-critical CPS, which we pioneered in a series of papers (including our RTAS'16 paper, which won the best paper award).

Furthermore, we have developed a series of resilient OS and hardware architecture prototypes, by extending open-source OS (Linux) and hardware (RISC-V SoC), that can defend against micro-architectural attacks with minimal impacts on performance. We believe that these realistic prototypes have real potential to influence the broader computer industry.

Education and Outreach

- Farzad Farshchi, Qijing Huang, and Heechul Yun. "Integrating NVIDIA Deep Learning Accelerator (NVDLA) with RISC-V SoC on FireSim." *Workshop on Energy Efficient Machine Learning and Cognitive Computing for Embedded Applications (EMC^2)*, February 2019.
- Heechul Yun, "Micro-Architectural Attacks on Cyber-Physical Systems," *Real-Time Scheduling Open Problems Seminar (RTSOPS)*, July 2019.
- Waqar Ali, Michael Bechtel, and Heechul Yun, "Analyzable and Practical Real-Time Gang Scheduling on Multicore Using RT-Gang." *Workshop on Operating Systems Platforms for Embedded Real-Time applications (OSPERT)*, July 2019.

Publications

- Michael Garrett Bechtel and Heechul Yun, "Denial-of-Service Attacks on Shared Cache in Multicore: Analysis and Prevention," *IEEE International Conference on Real-Time and Embedded Technology and Application Symposium* (RTAS), April 2019. (**Outstanding Paper Award**)
- Waqar Ali and Heechul Yun, "RT-Gang: Real-Time Gang Scheduling Framework for Safety-Critical Systems." *IEEE Intl. Conference on Real-Time and Embedded Technology and Applications Symposium* (RTAS), April 2019.
- Jacob Michael Fustos, Farzad Farshchi, and Heechul Yun, "SpectreGuard: An Efficient Data-centric Defense Mechanism against Spectre Attacks," *ACM/IEEE Design Automation Conference* (DAC), June 2019.
- Renato Mancuso, Heechul Yun, and Isabelle Puaut, "Impact of DM-LRU on WCET: a Static Analysis Approach," *Euromicro Conference on Real-Time Systems* (ECRTS), July 2019.

Vanderbilt University

The Vanderbilt University (VU) Science of Security (SoS) Lablet team is led by Principal Investigator (PI) Xenofon Koutsoukos. The team of faculty, postdoctoral and PhD student researchers from VU and Sub-Lablets Massachusetts Institute of Technology and the University of California, Berkeley are engaged in four research projects focusing on Cyber-Physical Systems (CPS).

The VU Lablet aims at developing the principles governing secure and resilient CPS in adversarial environments and using these principles for system design and management. Systems approaches require a mix of methods and tools. Our projects build upon our strengths on system and game theory, formal methods, data science, incentive engineering, and social science. Under these projects we are committed to developing integrated solutions that increase our understanding of complex interrelationships, anticipate future conditions, and support decision and policy making. In particular, we are seeking intellectual advances in which underlying theories are integrated and abstracted to develop explanatory models. These explanatory models derived from the underlying theoretical foundations lead to testable hypotheses. Hypotheses are tested using simulation and experimentation testbeds to gain greater understanding of CPS attacks and defenses. Based on collected evidence supporting or falsifying the hypotheses, new insights are obtained allowing the explanatory models to be refined or updated.

Results of the project research in 2019 have led to papers accepted by the International Conference on Cyber Conflict, the SSRN Electronic Journal, the IEEE Symposium on Technologies for Homeland Security, the 7th Workshop on Modeling and Simulation of Cyber-Physical Energy Systems, the Proceedings of Advances in Neural Information Processing, the Conference on Decision and Game Theory for Security, and the Resilience Week 2019 Symposium.


Xenofon Koutsoukos

Moreover, VU Lablet researchers have given talks on their projects at the NITRD Privacy R&D Interagency Working group and the Cybersecurity Research Group at Fujitsu System Integration Laboratories Ltd.

Details on the four research projects identified below can be found in the following sections.

- Analytics for CPS Cybersecurity
- Foundations of CPS Resilience
- Mixed Initiative and Collaborative Learning in Adversarial Environments
- Multi-Model Test Bed for the Simulation-Based Evaluation of Resilience

VANDERBILT®

Analytics for Cyber-Physical System Cybersecurity



Nazli Choucri

PI:	Nazli Choucri
SUB-LABLET:	Massachusetts Institute of Technology
URL:	https://cps-vo.org/VU_AnalyticsforCPS-Cybersecurity
HARD PROBLEM:	Policy governed Secure Collaboration

Goal

The overarching purpose of this project is to develop analytical methods to support the national strategy for cybersecurity, as outlined in Presidential Executive Orders and National Defense Authorization Acts. Operationally, we seek to develop analytics for cyber security policies and guidelines targeted specifically to (a) overcome the limitations of text-based guidelines, (b) extract the knowledge embedded in policy guidelines, (c) use this knowledge for risk and contingency analysis, and (d) assist the user community, analysts, and operators in implementation. Our task is to construct new tools and methods that are applicable to policy directives, regulations, and guidelines for diverse issue areas. The tools we provide will enable users to explore mission-related properties, concerns, or contingencies.

Abstract

Mounting concerns about safety and security have resulted in an intricate ecosystem of guidelines, compliance measures, directives and policy reports for cybersecurity of all critical infrastructure. By definition, such guidelines and policies are written in linear sequential text -- word form that makes them difficult to integrate or track feedback and delays, or to understand the policy-technology-security interactions, thus limiting their relevance for science of security. The challenges are to develop a structured system model from text-based policy guidelines and directives in order to identify major policy-defined system-wide parameters, situate vulnerabilities, map security requirements to security objectives, and advance research on how multiple system features respond to diverse policy controls to strengthen the security of fundamentals in cyber physical systems.

We draw on major reports presented by the National Institute for Standards and Technology (NIST) as the source of our data. While some efforts have already been made to mine NIST materials, few exploit the value of multi-methods for knowledge mining and analytical tools to support user understanding, analysis, and eventually action. Our approach learns from, and transcends, the above efforts by developing a platform for multi-methods cybersecurity analytics based entirely on the contents of policy documents. The case application focuses on cybersecurity of smart grid for electric power systems.

Accomplishments

Thus far we have been able to align the project vision and mission to National Cybersecurity Policy and identify the policy-relevant ecosystem. Focusing on national cybersecurity policies for securing CPS, we developed the research design. First we identified core policy documents for smart grid CPS. Then we extracted data from the documents and created a Dependency Structure Matrix (DSM) of the "as-is" smart grid NIST reference model. We also completed the design and operational strategy of the data extraction and linkage method, including (1) identifying and extracting the value-added of policy documents and guidelines, and (2) developing the process to move from "policy-as-text" to "text-as-data" for required steps of "data-as-metrics" in constructing the Platform for Cyber Analytics. In addition, we pre-tested the operational framework.

A review of rules and methods we developed for extracting data from key documents and creating the linked database allowed us to create (i) initial exploratory tools for analysis of system information, and (ii) a core DSM of the CPS system by identifying the first level information and relationship dependencies. The dependency matrix will be (a) examined closely and validated, (b) further transformed as needed into clusters and partitions of structure and process in order to (c) explore properties that reveal interconnections and "hidden features". It is also the basis upon which added policy imperatives - also in text form - are incorporated later on in expanded DSM forms.

We have also focused on important research steps: * making correctives; * replicating structured DSM model; * developing operational extensions; * testing applications of new method; * validating of initial structural model; * exploring automation issue.

Impact on Hard Problems

Our major contribution to the specific hard problem we address is dealing with "text-as-data" then "data-as metrics" in a complex cyber-physical where threats to operations serve as driving motivations for policy responses. The research outputs of this core project include. but are not limited to: (a) methods to examine the implications of cyber security directives and guidelines directly applicable to the system in question, (b) information about risk and relative vulnerability pathways throughout the whole or parts of system-network as embedded in the policy documents and guidelines, (c) insights from contingency investigations, that is, "what...if...", (d) framework information management within the organization, and (e) ways to facilitate information flows bearing on decision-making for cybersecurity.

Education and Outreach

- The validation study of "text-as-data" on cyber operations during war (*Tallinn Manual 2.0* [13]) (which we conducted earlier as proof of concept for the research methodology) has now been prepared for circulation and publication.
- PI Nazli Choucri is teaching a graduate-level class on *International Relations Theory in the Cyber Age,* drawing on policy-relevant materials. The course draws from student participants from different parts of MIT.
- PI Nazli Choucri presented the project overview, and its relevance to the United States at the 2019 Summer Science and Security and Privacy Lablet Quarterly Meeting, hosted by the University of Kansas in Lawrence, KS, July 9 and 10, 2019.
- The work was presented at NITRD Privacy R&D Interagency Working group, on August 2, 2019.

Publications

- Thomas Klemas, Rebecca Lively, and Nazli Choucri, "Cyber Acquisition: Policy Changes to Drive Innovation in Response to Accelerating Threats in Cyberspace," *2018 International Conference on Cyber Conflict U.S.* (CyCon U.S.), November 14-15, 2018, Washington DC.
- Nazli Choucri and Gaurav Agarwal, "International Law for Cyber Operations: Complexity and Transparency," SSRN Electronic Journal.
- Nazli Choucri and Gaurav Agarwal, "Imperatives for Securing the Long-Chain of Cyber-Physical Global Communication Infrastructure," *IEEE Symposium on Technologies for Homeland Security*.
- Mathieu Dahan, Saurabh Amin, and Patrick Jaillet, "Probability Distributions on Partially Ordered Sets and Network Security Games," https://arxiv.org/abs/1811.08516

Foundations of a CPS Resilience



Xenofon Koutsoukos

PI:	Xenofon Koutsoukos
URL:	https://cps-vo.org/VU_FoundationsforCPS-Resilience
HARD PROBLEM:	Resilient Architectures

Goal

The goals of this project are to develop the principles and methods for designing and analyzing resilient CPS architectures that deliver required service in the face of compromised components. A fundamental challenge is to understand the basic tenets of CPS resilience and how they can be used in developing resilient architectures.

Abstract

As CPS become more prevalent in critical application domains, ensuring security and resilience in the face of cyber-attacks is becoming an issue of paramount importance. Cyber-attacks against critical infrastructures, smart water-distribution and transportation systems for example, pose serious threats to public health and safety. Owing to the severity of these threats, a variety of security techniques are available. However, no single technique can address the whole spectrum of cyber-attacks that may be launched by a determined and resourceful attacker. In light of this, we consider a multi-pronged approach for designing secure and resilient CPS, which integrates redundancy, diversity, and hardening techniques for designing passive resilience methods that are inherently robust against attacks and active resilience methods that allow responding to attacks. We also introduce a framework for quantifying cyber-security risks and optimizing the system design by determining security investments in redundancy, diversity, and hardening. To demonstrate the applicability of our framework, we use a modeling and simulation integration platform for experimentation and evaluation of resilient CPS using CPS application domains such as power, transportation, and water distribution systems.

Accomplishments

A Game-theoretic Approach for Selecting Optimal time-Dependent Thresholds for Anomaly Detection: Adversaries may cause significant damage to smart infrastructure using malicious attacks. To detect and mitigate malicious attacks before they can cause physical damage to smart infrastructure, operators can deploy Anomaly Detection Systems (ADS), which can alarm operators to suspicious activities. However, detection thresholds of ADS need to be configured properly, as an oversensitive detector raises a prohibitively large number of false alarms, while an undersensitive detector may miss actual attacks. Using a game-theoretic approach, we formulate the problem of computing optimal detection thresholds which minimize both the number of false alarms and the probability of missing actual attacks as a two-player Stackelberg security game.

Diversity and Trust to Increase Structural Robustness in Networks: Our goal is to improve the structural robustness in networks using the notions of diversity and trustiness. Diversity means that nodes in a network are of different types and have many variants. Trustiness means that a small subset of nodes is immune to failures and attacks. We show that by combining diversity and trustiness within the network, we can significantly limit the attacker's ability to change the underlying network structure by strategically removing nodes.

Integrated Data Space Randomization and Control Reconfiguration for Securing Cyber-Physical Systems: Non-control data attacks have become widely popular for circumventing authentication mechanisms in websites, servers,

and personal computers. In the context of CPS, attacks can be executed against not only authentication but also safety. Moving Target Defense (MTD) techniques such as Data Space Randomization (DSR) can be effective for protecting against various types of memory corruption attacks, including non-control data attacks. This work addresses the problem of maintaining system stability and security properties of a CPS in the face of non-control data attacks by developing a DSR approach for randomizing binaries at runtime, creating a variable redundancy-based detection algorithm for identifying variable integrity violations, and integrating a control reconfiguration architecture for maintaining safe and reliable operation.

Impact on Hard Problems

PS are ubiquitous in critical application domains which necessitates that systems demonstrate resiliency under cyber-attacks. Our proposed approach integrates redundancy, diversity, and hardening methods for designing both passive resilience methods that are inherently robust against attacks and active resilience methods that allow responding to attacks.

Education and Outreach

- Our research was presented at conferences such as the 2019 Hot Topics in the Science of Security Symposium (HotSoS), CPS Week 2019, and GameSec 2019.
- We offered two summer camps for high-school students and one summer camp for teachers on CPS security based on Roboscape, a collaborative, networked robotics environment that makes key ideas in computer science accessible to groups of learners in informal learning spaces and K12 classrooms.

Publications

- Himanshu Neema, Harsh Vardhan, Carlos Barreto, and Xenofon Koutsoukos, "Web-based Platform for Evaluation of Resilient and Transactive Smart-Grids," *7th Workshop on Modeling and Simulation of Cyber-Physical Energy Systems*, Montreal, Canada, April 15, 2019.
- Bradley Potteiger, Zhenkai Zhang, and Xenofon Koutsoukos, "Integrated Data Space Randomization and Control Reconfiguration for Securing Cyber-Physical Systems," *Science of Security Symposium (HotSoS 2019)*, Nashville, TN, April 2-3, 2019.
- Amin Ghafouri, Aron Laszka, Waseem Abbas, Yevgeniy Vorobeychik, and Xenofon Koutsoukos, "A Game-Theoretic Approach for Selecting Optimal Time-Dependent Thresholds for Anomaly Detection. Autonomous Agents and Multi-Agent Systems," 33:430-456, 2019.
- Akos Ledeczi, Miklos Maroti, Hamid Zare, Bernard Yett, Nicole Hutchins, Brian Broll, Peter Vogyesi, Michael B. Smith, Timothy Darrah, Mary Metelko, Xenofon Koutsoukos, and Gautam Biswas, "Teaching Cybersecurity with Networked Robots," *SIGCSE 2019*. Minneapolis, MN, Feb. 27 - March 2, 2019.
- Waseem Abbas, Aron Laszka, and Xenofon Koutsoukos, "Diversity and Trust to Increase Structural Robustness in Networks," *2019 American Control Conference*, Philadelphia, PA. July 10-12, 2019.
- Aron Laszka, Xenofon Koutsoukos, and Yevgeniy Vorobeychik, "Towards High-Resolution Multi-Stage Security Games,"Cliff Wang and Zhuo Lu (eds.), Proactive and Dynamic Network Defense, Advances in Information Security 74, pp. 139-161. Springer, Cham, 2019.
- Carlos Barreto and Xenofon Koutsoukos. "Attacks on Electricity Markets," *57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, September 24-27, 2019.
- Aron Laszka, Waseem Abbas, Yevgeniy Vorobeychik, and Xenofon Koutsoukos, "Detection and mitigation of attacks on transportation networks as a multi stage security game,"Computers & Security,¬†Volume 87, Nov. 2019.
- Carlos Barreto and Xenofon Koutsoukos. "Design of Load Forecast Systems Resilient Against Cyber-Attacks," *2019 Conference on Decision and Game Theory for Security (GameSec 2019)*, Stocholm, Sweden, October 30 - November 1, 2019.
- Xingyu Zhou, Yi Li, Carlos Barreto, Jiani Li, Peter Volgyesi, Himanshu Neema, and Xenofon Koutsoukos, "Resilient Forecasting of Grid Loads under Stealthy Adversarial Attacks," *Resilience Week 2019*, San Antonio, TX, November 4-7, 2019.

Mixed Initiative and Collaborative Learning



Claire Tomlin



Shankar Sastry

PI:	Claire Tomlin	CO-PI:	Shankar Sastry
SUB-LABLET:	University of California, Berkeley		
URL:	https://cps-vo.org/VU_MixedInitiativeCollaborativeLearning		
HARD PROBLEM:	Human Behavior 		

Goal

One of the goals of the research is to characterize the limiting behavior of machine learning algorithms deployed in competitive settings.

Abstract

This research project focuses on a game theoretic approach to learning dynamic behavior safely through reachable sets, probabilistically safe planning around people, and safe policy gradient reinforcement learning. An understanding of the behaviors (convergence, optimality, etc.) of these algorithms in such settings is sorely lacking. We look at disturbance (attempts to force system into unsafe region) and control (attempts to stay safe) as well as fundamental issues with gradient play in games, since machine learning algorithms are increasingly being implemented in competitive settings. The fundamental issues with gradient play revolve around whether gradient descent is the right backbone for learning algorithms in games, and open questions include whether there are classes of algorithms that preclude limit cycles in their dynamics and how fast machine learning algorithms can escape saddles in competitive settings.

Accomplishments

Our primary work on this project has dealt with the brittleness of machine learning (especially deep learning) algorithms when used for intrusion detection or for the detection of Advanced Persistent Threats. We have addressed the following concerns:

Step size matters in deep learning: Training a neural network with the gradient descent algorithm gives rise to a discrete-time nonlinear dynamical system. Consequently, behaviors that are typically observed in these systems emerge during training, such as convergence to an orbit but not to a fixed point or dependence of convergence on the initialization. To elucidate the effects of the step size on training of neural networks, we study the gradient descent algorithm as a discrete-time dynamical system, and by analyzing the Lyapunov stability of different solutions, we show the relationship between the step size of the algorithm and the solutions that can be obtained with this algorithm. We showed that choice of step size was critical in determining when deep learning algorithms converged or when they resulted in limit cycles. This has implications for their use in intrusion detection. The results provide an explanation for several phenomena observed in practice, including the deterioration in the training error with increased depth, the hardness of estimating linear mappings with large singular values, and the distinct performance of deep residual networks.

Cross-entropy loss and low-rank features have responsibility for adversarial examples: Because state-of-the-art neural networks are vulnerable to adversarial examples, they can easily misclassify inputs that are imperceptibly different than their training and test data. In this work, we establish that the use of cross-entropy loss function and the low-rank features of the training data have responsibility for the existence of these inputs. Based on this observation, we suggest that addressing adversarial examples requires rethinking the use of cross-entropy loss function and looking for an alternative that is more suited for minimization with low-rank features. In this direction, we present a training scheme called differential training, which uses a loss function

defined on the differences between the features of points from opposite classes. We show that differential training can ensure a large margin between the decision boundary of the neural network and the points in the training dataset. This larger margin increases the amount of perturbation needed to flip the prediction of the classifier and makes it harder to find an adversarial example with small perturbations. We test differential training on a binary classification task with CIFAR-10 dataset and demonstrate that it radically reduces the ratio of images for which an adversarial example could be found -- not only in the training dataset, but in the test dataset as well.

Key to the success of robot motion planning in environments with humans is the ability to design plans that are robust to human behavior. In order to safely operate around humans, robots can employ predictive models of human motion, but these models cannot capture the full complexity of human behavior and necessarily introduce simplifying assumptions. As a result, predictions may degrade whenever the observed human behavior departs from the assumed structure. In our research we observe how "rational" human actions that appear under a particular model can be viewed as an indicator of that model's ability to describe the human's current motion. By reasoning about this model confidence in a real-time Bayesian framework, we show that the robot can very quickly modulate its predictions to become more uncertain when the model performs poorly, and we leverage these confidence-aware human motion predictions to generate assured autonomous robot motion. Our new analysis combines worst-case tracking error guarantees for the physical robot with probabilistic time-varying human predictions, yielding a quantitative, probabilistic safety certificate. Current safety analysis tools enable autonomous systems to reason about safety given full information about the state of the environment a priori. However, these tools do not scale well to scenarios where the environment is being sensed in real time, such as during navigation tasks. We propose a novel, real-time safety analysis method based on Hamilton-Jacobi reachability that provides strong safety guarantees despite environment uncertainty. We demonstrate our approach in simulation and in hardware to provide safety guarantees around a state-of-the-art vision-based, learning-based planner.

Impact on Hard Problems

We have been developing a framework for incorporating **human behavior** into **resilient** robot motion planning. We have been using reachability analysis to develop **scalable**, online safety updates of these motion plans. We have been developing **scalable**, analyzable methods for learning unknown dynamics within the framework of feedback linearization.

Education and Outreach

- We have thought about how machine learning is taught; particularly about the relationship between low rank features of the data and machine learning. Professor Yi Ma taught a class on this topic in Fall 2019.
- PI Claire Tomlin ran the 5th installment of Berkeley Girls in Engineering (GiE), a program held at UC Berkeley for middle school students, in Summer 2019. The program runs for 4 weeks, with 30 students participating per week, for a total of 120 students each summer. The week-long day camp includes modules across all types of engineering, with hands-on experiments, to teach about bioengineering, robotics, material science, computer science, water treatment, concrete design, and a range of other engineering topics. Students team up to complete a poster about an engineering problem and how they would solve it, which was presented at the end of the week to the camp and family members.
- We are developing a new course in systems theory at Berkeley, to be taken by upper level undergraduates and first and second year graduate students, on a rapprochement between control theory and reinforcement learning. The course will focus on a modern viewpoint on modeling, analysis, and control design, leveraging tools and successes from both systems and control theory and machine learning. The first version of this course will be taught by Shankar Sastry in Spring 2020.

Publications

- Kamil Nar and Shankar Sastry, "Step Size Matters in Deep Learning," in *Proceedings of Advances in Neural Information Processing*, December 2018.
- Kamil Nar, Orcan Ocal, Shankar Sastry, and Kannan Ramachandran, "Cross-Entropy Loss and Low-Rank Features Have Responsibility for Adversarial Examples," https://arxiv.org/abs/1901.0836 January 2019.
- Andrea Bajcsy, Somil Bansal, Eli Bronstein, Varun Tolani, and Claire Tomlin, "An Efficient Reachability-Based Framework for Provably Safe Autonomous Navigation in Unknown Environments," *Conference on Decision and Control (CDC)*, 2019.
- Andrea Bajcsy, Sylvia Herbert, David Fridovich-Keil, Jaime Fisac, Sampada Deglurkar, Anca Dragan, and Claire Tomlin, "A Scalable Framework for Real-Time Multi-Robot, Multi-Human Collision Avoidance," *International Conference on Robotics and Automation (ICRA)*, 2019.
- David Fridovich-Keil, Andrea Bajcsy, Jaime Fisac, Sylvia Herbert, Steven Wang, Anca Dragan, and Claire Tomlin, "Confidence-aware motion prediction for real-time collision avoidance," *International Journal of Robotics Research (IJRR)*, 2019.

Multi-Model Testbed for the Simulation-Based Evaluation of Resilience



Peter Volgyesi



Himanshu Neema

PI:	Peter Volgyesi	CO-PI:	Himanshu Neema
URL:	https://cps-vo.org/VU_Multi-modelTestBed-EvaluationofResilience		
HARD PROBLEM:	Security Metric and Models, Resilient Architectures		

Goal

The goal of the Multi-Model Testbed is to provide a collaborative design tool for evaluating various cyber attack/defense strategies and their effects on the physical infrastructure. The web-based, cloud-hosted environment integrates state-of-the-art simulation engines for the different CPS domains and presents interesting research challenges as ready-to-use scenarios. Input data, model parameters, and simulation results are archived, versioned with a strong emphasis on repeatability and provenance.

Abstract

We have developed the SURE platform, a modeling and simulation integration testbed for evaluation of resilience for complex CPS. Our previous efforts resulted in a web-based collaborative design environment for attack-defense scenarios supported by a cloud-deployed simulation engine for executing and evaluating the scenarios. The goal of this project is to significantly extend these design and simulation capabilities for better understanding the security and resilience aspects of CPS systems. These improvements include the first-class support for the design of experiments (exploring different parameters and/or strategies), target alternative CPS domains (connected vehicles, railway systems, and smart grids), incorporating models of human behavior, and executing multistage games. We also integrate state-of-the-art machine learning libraries and workflows to support security research with AI-assisted CPS applications.

To achieve these goals, we introduced significant changes to the SURE testbed architecture, replacing HLA-based C2 Windtunnel federated simulation engine with a more lightweight integration approach within WebGME and DeepForge.

Accomplishments

Our current testbed efforts are focused on developing a fully integrated workflow in DeepForge targeting the smart grid CPS domain. This work had two major goals: 1) A complete set of prediction, attack and detection models have been developed for load forecasting applications; and 2) Several building blocks of these models--most notably for gradient-based deep neural network attacks--are generalized to form the basis of a future library of reusable components to create SoS experiments involving learning enabled components. Our Testbed development effort is now focusing on the initial integration of two existing design studios**:** (1) DeepForge, our collaborative deep neural network experimentation platform with TensorFlow/Keras backend support and (2) GridLAB-D Design Studio, for configuring and executing smart power grid simulation models through a web-based interface.

Impact on Hard Problems

As part of our framework's metrics-driven evaluation capability, we are developing a library of cyber-attacks (as DeepForge workflows/pipelines) and neural network models. The testbed provides a strictly versioned store for these models along with input datasets (nor part of the visual model) and generated results, thus all experiments can be traced, re-executed and cloned for testing alternative algorithms or parameters. The current modeling strategy – which allows for custom Python-based processing blocks to be implemented through the web interface – provides a more open and pragmatic approach for experienced researchers to develop re-usable components for the testbed.

Education and Outreach

- We have established a collaboration and technical exchange with the Cybersecurity Research Group at Fujitsu System Integration Laboratories Ltd. This group uses WebGME, DeepForge and technology elements of our SURE testbed to develop their Cyber Range product. We had a collaborative workshop in Tokyo, Japan between November 24-27, 2019

Publications

- Akos Ledeczi, Miklos Maroti, Peter Volgyesi, Brian Broll, Hamid Zare, Bernard Yett, Timothy Darrah, Mary Metelko, Nicole Hutchins, Michael Smith, Gautam Biswas and Xenofon Koutsoukos, "Teaching Cybersecurity with Networked Robots*," Special Interest Group on Computer Science Education (SIGCSE) Technical Symposium 2019*, February 27th - March 2nd, Minneapolis, Minnesota.
- Bradley Potteiger, Zhenkai Zhang and Xenofon Koutsoukos, "Integrated Data Space Randomization and Control Reconfiguration for Securing Cyber-Physical Systems," *6th Annual Hot Topics in the Science of Security, HotSoS 2019,* April 2-3, Nashville, TN.
- Carlos Barreto and Xenofon Koutsoukos, "Attacks on Electricity Markets," *6th Annual Hot Topics in the Science of Security, HotSoS 2019*, April 2-3, Nashville, TN. (*work-in-progress presentation*).
- Xingyu Zhou, Yi Li, Carlos Barreto, Jiani Li, Peter Volgyesi, Himanshu Neema, and Xenofon Koutsoukos, "Evaluating Resilience of Grid Load Predictions under Stealthy Adversarial Attacks," *Resilience Week 2019 Symposium*, San Antonio, TX, November 4-7, 2019. Received **Best Paper Award**
- Carlos Barreto, Alvaro Cárdenas, and Nicanor Quijano, "Design of Load Forecast Systems Resilient Against Cyber-Attacks," in *International Conference on Decision and Game Theory for Security (GameSec),* Springer, 2019.
- Carlos Barreto and Xenofon Koutsoukos, "Attacks on electricity markets," in *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2019.

Science of Security

Quarterly Meetings



Winter 2019 Lablet Quarterly

The Winter 2019 Science of Security and Privacy (SoS) Lablet Quarterly meeting was held at the International Computer Science Institute (ICSI) in Berkeley, California on January 10-11, 2019. Hosted by Serge Egelman, Principal Investigator (PI) at ICSI, the meeting was attended by over thirty representatives from NSA and the six Lablets as well as some of the Sub-Lablets and industry. This Quarterly meeting focused on privacy and Machine Learning. The agenda included a presentation on NSA's privacy research, an industry presentation on practical data governance, and Lablet project presentations that dealt with aspects of privacy.

Invited Talks

After welcomes from Serge Egelman and Adam Tagert, (SoS Technical Director), Ray Brown from NSA's Laboratory for Advanced Cybersecurity Research spoke on NSA's Privacy Research Problems. His stated long-term research interest and goal is to enable people to communicate with computers informally, using natural language. He described three areas being prototyped that are aimed at preventing more intrusions:

- Chat Bot (Topic Modeling): Can this technique complement or improve data discovery currently done with tools like elastic search or "google-like" techniques?
- Can secure multi-party communication techniques provide minimally intrusive means of discovering information in order to prevent more intrusive measures?
- Development of a Common ML Studio.

He then identified the following research areas that are of interest:

- What is the impact of General Data Protection Regulation (GDPR) on data sets used for training models?
- What is the impact of GDPR on models that "remember"

- training data?
- Can ML help disrupt online tracking?
- Can we create digital canaries that are geographically focused on discovering biasing attempts within social media or news?

He concluded by noting that they have been successful in developing an application that can detect when and how things are changing, which will allow someone to detect and mitigate effects of the change.



Ray Brown presented NSA's Privacy Research Problems

Chris Hoofnagle, who holds dual appointments at the School of Law and School of Information, University of California (UC) Berkeley, gave a presentation on the Federal Trade Commission (FTC) and Cyber, in which he described the history of the FTC and how it came to address cybersecurity, how the FTC regulates cyber, and some of its current challenges. In addressing how the FTC regulates cyber, he noted that the FTC interprets its unfairness authority to require "reasonable" security among all companies that collect personal information; the FTC bolsters "consent" requirements in surveillance matters; and the FTC erodes intermediary, software immunity. When he spoke of FTC challenges, he said they could also be interpreted as research questions. He cited the Commission's small size in questioning how it could promote security in a reasonable way. He also posed the following additional research questions:

- What "cybersecurity" is promoted by a consumer protection-oriented agency, and is this the kind of security we most need?
- Should "reasonable security" be dependent on the FTC?
- How can the FTC signal its expectations to businesses and afford them due process without developing formal rules, which can be sclerotic?

David Marcos, Senior Program Manager for Microsoft Cloud and Enterprise, spoke on "Practical Data Governance at Microsoft" describing how Microsoft seeks to address the requirements of GDPR. He addressed the challenges associated with building out privacy at scale and described in detail some of Microsoft's goals in that endeavor. Microsoft's challenges include dealing with a tremendous amount of legacy software and building privacy at scale. Among their goals are establishing a common language and applying it consistently across policy and engineering practices; at-scale solutions to manage data as it is collected, copied, and moved; structured process integrated into engineering tools; and monitoring to help drive engineering accountability. Mr. Marcos made the point that the Cloud is data-driven and, consequently, at-scale data governance must also be data-driven to keep pace with Cloud growth; at-scale data governance is accomplished via tagging of data. He noted that engineers need practical tooling tied to data tagging in order to implement simplified, at-scale data governance. Finally, he said that data governance at-scale must be quantifiable to be able to evaluate health and status, tied to tooling for end-to-end tracking, and scaled according to business needs.

Lablet Project Presentations

Following the introductory talks which provided external perspectives and context, the remaining presentations related to Lablet projects dealing with privacy.

In his presentation "Data Governance Requirements," Joshua Kroll, UC Berkeley, described work that is yet to be done, and characterized the drivers of data governance as regulatory, public pressure, consumer demands, and global reach (different legal regimes in different jurisdictions). The goal of this project is to develop a methodology for data governance including developing a validated model for translating amorphous or contested values into actionable requirements. This study has a two-pronged approach: 1) A taxonomy of privacy needs and desires—a field study of privacy practices in organizations and how they are realized; and 2) a taxonomy of privacy tools such as encryption, auditing, and tagging. These prongs will establish what questions the methodology must answer and the tools. With respect to data governance practices, the study seeks to establish how privacy goals, policies, and disclosures are currently set; what sort of mechanisms are used to assure both insiders and outsiders that goals and policies are being met; how subjects and organizations cope with tradeoffs; who has input into the data governance development process; and the gate points through which control and internal governance is implemented. Dr. Kroll enumerated the research questions as follows:

- How do organizations conceptualize their privacy goals?
- What privacy guarantees would organizations desire to make, and to whom, unencumbered by technical or organizational limitations?
- How do organizations assure themselves that they are upholding the commitments?

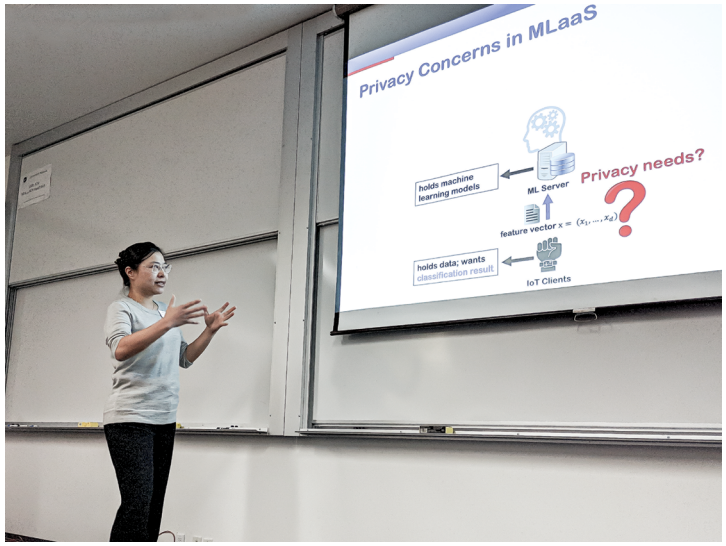
He described the progress to date as having developed and refined the approach and protocol for the field portion, the completion of an initial evaluation of candidate field study sites, and the development of an initial taxonomy of technical and non-technical tools.

Will Enck, North Carolina State University (NCSU), gave a presentation entitled "PolicyLint: Investigating Internal Privacy Policy Contradictions on Google Play" in which he described PolicyLint as a text analysis tool designed to aid humans in identifying contradictory collection and sharing statements in

software privacy practices. He pointed out that privacy policies define what apps are legally allowed to do with user information and identified several apps that are violating their privacy policies. He spoke of behavior-to-policy compliance and creating analysis tools to compare behavior in mobile apps to statements in privacy policies. Dr. Enck described the research methodology for the privacy study and presented the study findings which centered on the fact that policies directly contradict themselves. This generated a lively discussion among the attendees about the development of policy regulations and prior studies that addressed similar issues.

Fengjun Li, University of Kansas, spoke on "Privacy-Preserving Classifications for IoT" which centered on privacy concerns associated with Machine Learning as a service (MLaaS). Dr. Li noted that the future success of IoT is based on the data it will generate and how it will be able to be processed. In addressing MLaaS as the way to make big data analysis easy, she said that it requires infrastructure investment, integration and development, data modeling, and deep ML expertise. One privacy concern is that the client should learn only the classification result and nothing about the model. In developing privacy-preserving MLaaS for IoT, Dr. Li proposes shifting costly computation and communication out of the IoT devices and outsourcing some computations to a trusted or semi-trusted component. The goals are no privacy leakage in the outsourcing and a generic solution for most commonly used ML algorithms. This work, she stated, is a cloud-assisted design and a practical scheme with affordable computation and communications costs that will meet the privacy needs of multiple stakeholders, provide lightweight security solutions tailored for IoT applications, and reduce the computation and communication load at the device. She concluded by saying that the solution to use outsourcing aligns with emerging IoT infrastructures as well as existing functionalities of the application.

The final talk of the first day was "Uncertainty in Network Reliability Analysis" presented by Frank Nguyen, University of Illinois at Urbana-Champaign (UIUC). This presentation dealt with network reliability, specifically, the probability that the network remains connected under random and independent edge failures. In addressing network reliability applied to security, he laid out the approach and identified the challenges as how to model correlation between link existences and now to express uncertainty about link existence probabilities. He summarized the work under way and noted that the findings are currently under review.



Helen Nissenbaum, who is working with ICSI at Cornell Tech, spoke on "Privacy as Contextual Integrity". She described the fundamentals of Contextual Integrity (CI), and stated that privacy requires that actions, practices, and policies conform with legitimate contextual informational norms. In addressing contextual integrity fundamentals, she cited four essential claims and what each assertion was and was not:

- Flow: It *is* privacy as an appropriate flow of information; it *is not* privacy as secrecy/no collection/no flow, privacy as data minimization, or privacy as leakage;
- Appropriate flow: It *is* flow conforming with (entrenched) contextual norms; it *is not* procedural;
- Contextual information (privacy) norms refer to five independent parameters (actors: subject, sender, recipient; information types; transmission principles); such norms are *not* subject control: notice and consent; public/private sensitive; and general (vs. contextual) norms;
- Establishing the legitimacy of contextual norms considers the interests and preferences of the affected parties (individual); the ethical and political principles and values (societal); and the contextual functions, purposes, and values (societal); they are *not* only interests of data subject (harm to the individual); always contrary to societal values; and always contrary to utility

She stated that issues that are problems for the theory are disruptive technology, environmental change, scientific discoveries, and social change. She concluded by noting that privacy requires that actions, practices, and policies conform with legitimate contextual information norms, and that privacy is a value that serves societal and contextual ends.

Michael Tschantz, ICSI, gave a presentation entitled "Modeling Privacy" which also dealt with contextual integrity but focused on clarifying the topic for computer scientists using the language of mathematics. He spoke about the old model for contextual integrity being built on four sets of entities (agents, messages, information types, and roles), adding that a context in this model is simply a set of roles. The new model of contextual integrity that he proposes adds purposes and values that drive or constrain actions to the four sets of entities noted above. In this model, the context's defining feature is its purpose. He cautioned that this should not be taken too literally, and that the idea was to achieve your purpose as well as you can without violating any of your values. He continued his presentation with a discussion of differential privacy as a causal property and represented his example as a program.

Serge Egelman, ICSI, continued the contextual integrity theme in his presentation "User Privacy Preferences and IoT," and described the goal of this research as determining what privacy laws should be put in place for smart devices. He defined the

research approach as "how are we applying CI?" and described the steps as observing a data flow, measuring expectations, deconstructing the data flow into CI parameters, and inferring how expectations change based on the parameters. He presented some of the research that had already been carried out on this topic and then discussed an upcoming study to be undertaken on "Smart Speaker" users. Dr. Egelman also raised the question of what sort of privacy protections should future "always listening" devices provide, and addressed both the methodology and potential metrics for such research.

Nicolas Christin spoke about the privacy research being done at the Security Behavior Observatory (SBO) at Carnegie Mellon University (CMU), and focused on the work being done in predictive security analytics. He described the motivation for predictive security analytics as what needs to be done to recover after a successful attack with the research question being whether we can proactively sort out how to prevent most attacks from happening without resorting to a "fixing everything right now" strategy. Since most security attacks are economically motivated and the attackers have little interest in the targets themselves but rather care about volume, that fact leads to predictable behavior. Predictive security analytics is needed to protect both the infrastructure and the user. Dr. Christin discussed the data source and the methodology used in the research as well as their findings, adding that they will be using the SBO to validate their study.

Lujo Bauer, also of CMU, presented "ML-generated Attacks on ML" which provided an update on ongoing work dealing with fooling ML generators. He summarized by noting that "attacks" may not be meaningful until we can fix the context; even in a practical context, real attacks exist; hard-to-formalize constraints can be captured by a DNN; and similar principles about constrained context apply to other domains. Dr. Bauer then gave a presentation on understanding privacy valuations which discussed research that compared hypothetical and realistic privacy valuations. He defined privacy valuation as the willingness to sell and the selling price for personal information, and then described the study that sought to answer whether privacy valuations can be predicted.

The final presentation "Frame Semantics for Text Understanding" was given by Collin Baker, ICSI, and dealt with the FrameNet project, an effort to produce a lexicon of English that is both human and machine readable, and its implications for privacy. He described how semantic frames are represented in FrameNet, applications of Frame Semantic Annotation/Role Labeling, current projects, and future research.

The complete agenda and selected presentations are available for viewing on the Science of Security Virtual Organization website at https://cps-vo.org/LabletQTRLY/2019/ICSI.

Spring 2019 Quarterly

In lieu of a Spring Quarterly, the SoS community met in April at HotSoS 2019 hosted by the VU Lablet in Nashville, TN.

See Section 3, page 90, for full HotSoS 2019 details.

Summer 2019 Lablet Quarterly



The Summer 2019 Science of Security and Privacy (SoS) Lablet Quarterly meeting was held at the University of Kansas (KU) in Lawrence, Kansas on 9-10 July 2019. This quarterly meeting focused on Cyber-Physical Systems. It was hosted by Perry Alexander, Principal Investigator (PI) at KU. The attendees were representatives from NSA, other Lablets Sub-Lablets and industry. In addition to presentations on Lablet research projects, the agenda included presentations from industry, a panel discussion, and two NSA speakers.

Invited Presentations

Dr. Deborah Frincke, NSA Director of Research, began the meeting with a presentation on how to make innovation successful, and encouraged the attendees to ask her questions. She noted that lessons-learned from NSA research projects point to risk taking, teaming, budgeting, structure, and training as critical components for innovation to work. She raised the issue of research on Machine Learning, saying that while we need to understand the life cycle of ML and focus on the future, we also can't forget to fix the problems of the present. Continuing on the topic of ML, she said we need a better understanding of what it means to have a finding by ML algorithms, noting that ML testing is not necessarily the best way to test for attacks. She was asked about what the focus on innovation means for partnerships between NSA and industry, including realizing the importance of discovery-based research. She said that transitions between agencies are key as is the commitment to multiple approaches to achieve objectives. Her final question on what she saw as some of the most complex problems today led her to cite quantum resistance and quantum computing being misunderstood, as well as a general lack of understanding of the lifecycle of ML.

Mathew Kemph, Director of Intelligence Community Support at In-Q-Tel (IQT), the strategic investor for the US intelligence and defense communities that identifies startups with the potential for high impact on national security, discussed the work of the organization. IQT works closely with the identified companies to deliver new capabilities needed to boost the US government's technological edge. He described the firm's challenge of connecting the three very different worlds of venture capital, government, and startups in order to deliver capabilities, and said that IQT is focused on global investing for mission impact, noting that 10% of their investments are based outside the US. He pointed out that while IQT is funded by government partners and driven by mission, all of their work is unclassified. He described the four pillars of IQT's value proposition, and talked about the IQT labs and their research in geospatial analytics, cybersecurity, advanced analytics, biotechnology, and data visualization. He noted that IQT has provided more than 850 pilots for evaluation by government partners and facilitated more than 300 technology transfers into government agencies.

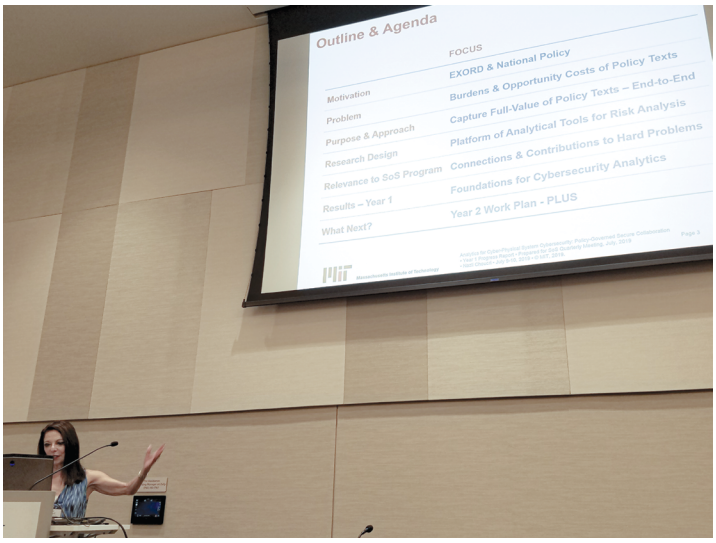
Maurice (Mo) Bland, who has experience at both NSA and the ODNI in cybersecurity operations, led off the second day of the Quarterly meeting. He expressed interest in emerging technologies and how they are going to be used, and is seeking to build partnerships with academia and industry to determine how academic research can be applied to IC challenges both today and in the future.

Panel Presentation

The Kansas Lablet has formed an Advisory Board (KUAB) whose members were recruited from a cross-section of industry and government (finance, healthcare, infrastructure, security, and manufacturing) that deal in CPS domains. The KU Lablet is the only SoS Lablet with an advisory board. Because of this uniqueness, a panel was held to discuss the board. Three KUAB members, from Collins Aerospace, Nokia, and Sprint, participated in a panel session where they provided their perspectives on their industries' challenges and responses dealing with security and privacy of CPS. David Hardin, Collins Aerospace, said that his company is dealing with CPS threats from both a tactical and strategic perspective. Tactically, they are characterizing the threats and attack surfaces for their products, performing penetration testing, finding and fixing vulnerabilities as they are reported, and characterizing and minimizing the trusted computing base. Their strategic approaches include investing in formal methods-based automated test generations, integrating formal methods into the development process, and researching advanced languages and operating systems. Doug Hohulin of Nokia presented Nokia's areas of expertise and cited their track record of innovation. He discussed the convergence of transportation, communications, and public safety policies as well as IoT security solutions and best practices. The third panelist, Lyle Paczkowski of Sprint, provided his perspective of CPS security challenges noting that the heterogeneity of CPS components and the diversity of CPS systems have made it very difficult to study the problem with one generalized mode.



KU Advisory Board members addressed CPS security during a panel discussion.



Nazli Choucri presented the results from the first year of her project

Lablet Project Presentations

The first research presentation was from Ninghui Li of Purdue University, a Sub-Lablet of the North Carolina State University (NCSU) Lablet. Li spoke on the project "Principles of Secure Bootstrapping for IoTs," addressing the motivations for the project, current research activities, and research directions that include identifying critical flaws in crypto design and implementations. The motivation for the research is that IoT devices need trust and secure communication—trust between devices; trust between devices and users; and trust between devices and the environment. Constraints—including power supply, resources, and serviceability—limit options, and deployment scenarios determine resource availability. He presented the results of two studies: Privacy attacks on the 4G and 5G cellular paging protocols using side channel information and a Zigbee security analysis.

Nazli Choucri of MIT, a sub-Lablet of the Vanderbilt University (VU) Lablet, presented an overview and Year 1 report on "Analytics for Cybersecurity of CPS". The project is focused on introducing analytics for CPS cybersecurity to enhance value of guidelines and directives. Year 1 results included identifying the cybersecurity policy ecosystem and developing data extraction and linkage methods. Their research approach is a testbed that focuses on NIST reports, data guidelines, and directives. The expected product is a platform of tools for analytics of cybersecurity, applications for which include support for the NIST Cyber Security Framework. She presented the research design and noted how this work contributes the addressing the SoS Hard Problem of policy-governed secure collaboration at the enterprise level and contributing to the Science of Security--the research provides replicable methods for cybersecurity policy analytics of a system and an enterprise; demonstrates multi-methods for cybersecurity policy of complex systems; and connects to diverse communities.

Heechul Yun of KU spoke on "Micro-Architectural Attacks on Cyber-Physical Systems" and described the activities that had been undertaken to address the issue. The goal of the project is to develop a microarchitectural attack-resistant computing infrastructure for secure Cyber Physical Systems. He presented the work that he and colleagues had done on SpectreGuard, a data-centric software/hardware defense mechanism against Spectre attacks in which the software tells the hardware what data (not code) needs protection and the hardware selectively protects the identified data from Spectre attacks. He also presented work done on denial of service attacks and on real-time task scheduling framework. His research has developed fundamental computing infrastructure technologies to enable safe, secure, and intelligent CPS.

The final research presentation on Day 1 was done by Xenofon Koutsoukos, Lead-PI at the VU Lablet, entitled "A Game-Theoretic Framework for Security and Resilience of CPS". The main goal of the research is to develop a model that enables finding optimal resilient configurations in the face of strategic attacks. Their research has enabled them to develop a game theoretic approach to detect attacks in CPS, featuring time dependent attack detection, detection and mitigation as a multi-stage security game, adversarial regression for sensor attack detection, and adversarial Gaussian process regression. Current and future work focuses on a game theoretic approach for randomized feature selection for anomaly detection and prediction using power and transportation systems in their case studies.

Nathan Malkin of University of California, Berkeley, a Sub-Lablet of the International Computer Sciences Institute (ICSI) Lablet spoke on "Smart Speakers Today and Tomorrow: Privacy Attitudes and Research Opportunities". The goal of the project is to develop privacy controls for always-listening devices using contextual integrity to evaluate flows and blocking inappropriate ones. He presented the results of a study that examined users' attitudes about the data derived from smart speaker interactions. He also discussed another study being undertaken to investigate

users' preferences and expectations for always-listening voice assistants, and a new one that would collect contextual cues for always-listening devices using the experience sampling model.

Himanshu Neema, VU, spoke on "Multi-Model Testbed for Deep Learning-Driven Resilient CPS" and described designing



Nathan Malkin, ICSI, spoke on privacy issues related to smart speakers.

resilient CPS with integrated Machine Learning components. The testbed is designed to produce repeatable experiments with quantifiable metrics using integrated machine learning components, and drawing upon multiple CPS domains that will provide for research collaboration and dissemination. Dr. Neema spoke on recent research work the first of which, railway network simulations with HIL (Hardware In-the-Loop), had as its goal the development of a simulation testbed with HIL that enables enhanced analysis of the resilience of railway networks against cyber and/or physical attacks. He also provided demos of two other research activities—web-based collaborative modeling and experiments on power distribution grids with market attacks; and resilient forecasting of grid loads under stealthy adversarial attacks.

John Symons, KU, gave a presentation entitled "Resilient Institutions and Normativity: Privacy, Autonomy, and the Cyber-Nonphysical," and noted that the research being done under this project proposes new research approaches and technologies with special focus on norms, institutions, and the protection of individual autonomy. He raised the topic of the Cyber Non-Physical by saying that building resilient systems and societies involves recognizing that the social and cultural domain is often more important than the technological, and that new technologies have made social infrastructure part of the attack surface in ways that we were all struggling to understand. He noted that a variety of norms, social and linguistic practices, agent interactions and games presuppose common knowledge, and the formal strategy for thinking about common knowledge is quite different from that employed in network science. This work contributes to the establishment of an interdisciplinary science of security by bringing an informed but critical attitude towards limitations of common formalisms and technologies and being attentive to neglected aspects of resilience.



John Symons, KU, gave a presentation dealing with the Cyber-NonPhysical.

Jonathan Aldrich, Co-PI at the Carnegie Mellon University (CMU) Lablet, presented "Obsidian: A Secure, Usable Language for Smart Contracts" and described progress under this project. He spoke about the language design process for Obsidian that took an interdisciplinary approach to first understand the user's problem, then apply iterative design, and finally validate through case studies and randomized controlled trials. He presented some case study results and concluded by noting that while it wasn't obvious that they were successful in making Obsidian usable, an interdisciplinary approach to language design is promising.

Sayan Mitra, Lead-PI at the University of Illinois at Urbana-Champaign (UIUC) Lablet, gave the final presentation, "Verification of Cyber-Physical Systems Through Learning". The project addresses resiliency and risk-reduction in CPS through verification, going beyond purely model-based approaches with machine learning. He suggests exploiting models when it makes sense and focusing on the marginal benefits of models' executable fidelity to security. He discussed DryVR, a new model for hybrid systems, as well as Statistical Model Checking (SMC) which is a data-driven approach for verifying requirements of stochastic systems specified in a stochastic temporal logic

The complete agenda and selected presentations are available for viewing on the Science of Security Virtual Organization website. https://cps-vo.org/LabletQTRLY/2019/KU.





Fall 2019 Lablet Quarterly

The Science of Security and Privacy (SoS) 2019 Fall Quarterly Lablet Meeting was held at the Discovery Partners Institute (DPI) in downtown Chicago on November 5th and 6th, 2019. This quarterly meeting was hosted by The University of Illinois Urbana-Champaign (UIUC). NSA researchers traveled to the event with the SoS team to join Principal Investigators from the six Lablets: Carnegie Mellon University (CMU), the International Computer Sciences Institute (ICSI), North Carolina State University (NCSU), UIUC, University of Kansas (KU), and Vanderbilt University (VU). This quarterly meeting included Lablet program reviews, technical presentations, an SoS-themed review presentation, and briefings presented by NSA leadership.

The theme of the Fall Quarterly Meeting was Resiliency which was a natural progression from the previous Summer Quarterly meeting's topic of Cyber Physical Systems at KU. At the summer Quarterly, NSA leadership described their collective efforts to leverage cybersecurity partnerships and expertise to protect the Nation's critical infrastructure, ensuring that the SoS initiative continues to connect research with Agency mission. The Fall Quarterly continued the emphasis on leveraging SoS Lablet research to support the Agency's mission.

Invited Talks

Following a welcome from Sayan Mitra, UIUC PI and an overview of the SoS initiative by the SoS team, Grant Wagner, Technical Director of the Cybersecurity Directorate Future and Standards Office at NSA, spoke on NSA's new Cybersecurity Directorate which had been stood up just a month earlier. He gave a brief overview of the Directorate and described the Directorate's goal of leveraging the Agency's current Lablet research partnerships to develop an unclassified cybersecurity collaboration center at NSA. He acknowledged the enormous amount of talent available in the private sector dealing with cybersecurity challenges and explained how his team will be leading the effort to leverage unclassified cybersecurity research at industrial scale.

Joy Garrett, NSA's Analysis and Production Investment Portfolio Manager, briefed the attendees on the business perspective and complexities of NSA's overarching SIGINT and Cyber Missions. She described her background and how that experience has given her a more informed and well-rounded perspective to manage the incredibly complex and ever-changing budget. Ms. Garrett explained the process and timeline the budget goes through and how she's already planning for Agency budget priorities through 2025. In response to the question "How do you identify the anticipatory challenges that far in advance?" she said that she looks to the NSA Research Directorate as the gateway to address what threats NSA should be anticipating. Ms. Garrett emphasized how critical it is that academia partner with NSA Research to identify these emerging threats and the eventual transfer of the needed cybersecurity tools.

Panel Discussion

A panel discussion on resiliency was led by SoS Technical lead, Adam Tagert. The panelists John Symons (KU), David Garlan (CMU), and Helen Gu (NCSU) discussed how systems resilience research has evolved over the past decade and the ever-growing need for more benchmark data and testbeds. The

focus of resiliency in academia and the tremendous benefits by interacting with interdisciplinary fields was reiterated by several panel members.

Lablet Project Presentations

Shankar Sastry, University of California, Berkeley, a Sub-Lablet of VU, gave a presentation entitled "Mixed Initiative and Collaborative Learning in Adversarial Environments". One of the goals of the research is to characterize the limiting behavior of machine learning algorithms deployed in competitive settings. The research includes learning dynamic behavior safely through reachable sets, probabilistically safe planning around people, and safe policy gradient reinforcement learning. The primary work on this project has dealt with the brittleness of machine learning (especially deep learning) algorithms when used for intrusion detection or for the detection of Advanced Persistent Threats and has the two major concerns: 1) Step size matters in deep learning; and 2) Cross-entropy loss and low-rank features have responsibility for adversarial examples.

Helen Gu, NCSU, presented her research on "Online Container Vulnerability Exploit Detection and Security Patching" which dealt with securing containers such as ones like Docker Containers. Containers are a popular applications deployment method with easy installation and low overhead. Security vulnerabilities widely exist in both official and community containers. Her research showed that a combination of static and dynamic analysis can identify in-progress cyber-attacks on 86% of vulnerabilities in their dataset. This research was published in the paper "A Study on Container Vulnerability Exploit Detection" and was presented at the 2019 IEEE International Conference on Cloud Engineering (IC2E).

Geir Dullerud, UIUC, spoke on "Verifying Performance Specifications on Hybrid Systems and Markov Decision Processes via Reinforcement Learning." He discussed a new reinforcement learning method and presented case studies.

Serge Egelman, ICSI PI, presented "Do you get what you pay for? Comparing the privacy practices of free vs. paid mobile apps" in which he discussed the trails of testing a corpus of over 6,000 phone apps. Egelman and his team have been successful in proving that the privacy benefits of the paid version of apps are negligible at best. Their research has led to the creation of a tool on Natural Language Processing to assist others in understanding the privacy expectations of phone apps.

Fengjun Li, KU, spoke on "Protecting Privacy in Online Social Networks: From Identity to Content" which was based on a paper presented at the 19th Privacy Enhancing Technologies Symposium (PETS), in Stockholm in July, 2019. Their research led to the development of PrivScore, which she described as the first attempt to quantitatively assess the level of privacy/ sensitiveness of unstructured text. The research team developed context-free, context-aware and personalized privacy scoring models, and experiment results show that PrivScore is consistent with users' perceptions.

Mohammad Nourredine, UIUC, spoke on "Achieving Intrusion Tolerance through Automatic Response and Improved Forensics" in which he presented two research questions: 1) How can we preserve acceptable operation during a state-exhaustion DDoS attack? and 2) If an exploit is detected, how can we preserve the system's identity? He presented the results of experiments and case studies and their conclusions that addressed each of the research questions.

Kevin Jin, Illinois Institute of Technology, a UIUC Sub-Lablet, gave a presentation entitled "Combining Simulation and Emulation for Evaluation of Secure and Resilient Cyber Physical Systems". He identified the problem statement for the research as developing a scalable and high-fidelity testbed for evaluating cyber effects on the physical system, described the testbed, and presented case studies using the research team's evaluation methodology.

The final presentation, "Improving Explainability with Contrastive Alternatives," was given by David Garlan, CMU. This research builds upon his SoS project work which seeks to show how probabilistic models and model checkers can be used both to synthesize complex plans that involve a combination of human and automated actions. Work would also provide human-understandable explanations of mitigation plans proposed or carried out by the system.

A full agenda and selected presentations are available on the Science of Security Virtual Organization website at https://cps-vo.org/SoSLmtg/UIUC/2019.

Promoting Rigorous Scientific Principles

The Science of Security and Privacy Initiative (SoS) promoted rigorous scientific principles through their sponsorship of fundamental research at the Lablet universities as well as through the sponsorship of two major competitions, the Annual Best Scientific Cybersecurity Paper Competition and awards at the Intel International Science and Engineering Fair (ISEF). The SoS initiative also sponsored a Best Paper Award at the Hot Topics in the Science of Security Symposium (HotSoS 2019).



Attendees at the 2019 Intel International Science and Engineering Fair (ISEF) try out the Enigma machine at the NSA booth.

The winner of the 7th Annual Paper Competition was "Evaluating Fuzz Testing" which was presented at ACM SIGSAC Conference on Computer and Communications Security (CCS '18) in Toronto. The research team investigated the evaluation process of fuzz testing tools. This paper is already having tremendous impact on fuzzing research, setting the standard for how evaluations should be done. It is bringing scientific principles, such as rigorous conclusions and reproducibility to an area in need of scientific understanding.

The 2019 competition also awarded two Honorable Mentions. The first paper, "Continuous Formal Verification of Amazon s2n" was presented at the International Conference on Computer Aided Verification, 2018. This paper discussed the researchers' advancements in the application of formal methods, and was selected for an honorable mention because it addresses one of the enduring challenges of formal methods. The researchers made a foundational research advancement and also provided a solution.

The second paper to receive an Honorable Mention, "Meltdown: Reading Kernel Memory from User Space," was presented at the 2018 USENIX Security Symposium. The paper detailed an attack to exploit a side effect of a computer processors' out-of-order execution performance feature. Meltdown was selected because it is an excellent example of researchers challenging and disproving assumptions of the security community.

This was the fifth consecutive year that SoS has sponsored awards at ISEF, and this year there were seven winners in the SoS category: one First Place, three Second Places, and three Honorable Mentions. The first-place winner, Suha Hussain, had been awarded a second place in the same category in 2018, and the second-place winner, Divya Amirtharaj, had also been awarded a second place prize in SoS in 2018. Other ISEF award categories were Mathematics, Cyber Pioneer, and Material Science. The first-place winner of the Mathematics category, Daniel Santiago, had received Honorable Mention in the SoS category in 2018. The Cyber Pioneer winner, Kevin Meng, had placed first in the Material Science category in 2018. A total of 17 prizes were awarded by the Research Directorate for ISEF 2019.

The winner of the HotSoS Best Paper competition "Integrated Data Space Randomization and Control Reconfiguration for Securing Cyber-Physical Systems" was selected by the program committee. This paper was written by a Vanderbilt Lablet researcher. The winning paper was automatically submitted to the 8th Annual Best Scientific Cybersecurity Paper Competition.

Details on the Best Scientific Cybersecurity Paper Competition and the Intel ISEF can be found in the following pages; details on the winning paper at HotSoS can be found in Section 3.



The winning paper of the 7th Annual Best Scientific Cybersecurity Paper Competition was *Evaluating Fuzz Testing* by George Klees, Andrew Ruef, Benji Cooper, Shiyi Wei, and Michael Hicks. The researchers were from the University of Maryland and the University of Texas at Dallas. This paper was presented at ACM SIGSAC Conference on Computer and Communications Security (CCS '18) in Toronto.

The research team investigated the evaluation process of fuzz testing tools. Fuzz testing tools help evaluate quality of software code by inputting large amounts of random data and monitoring the program for unexpected behavior. The researchers asked an important question, how to determine which tool is best. They began with a survey of the experimental evaluation contained within 32 fuzz testing papers. Their analysis identified problems in each of these evaluations so the research team developed their own extensive evaluation process and then re-evaluated the tools. The standardized evaluations showed that existing ad-hoc evaluation methodologies can lead to wrong or misleading assessments. From this research, the team developed guidance on evaluating tools. The paper was selected because it embodies the attributes of outstanding science and the criteria of the competition: rigorous research, generalizable results, and clarity of presentation. This paper is a step forward in bringing scientific understanding to the security community. It is grounded to current understanding by its methodological survey of evaluation practices, then advances the science through quantitative analysis and proposes conclusions that apply broadly in the fuzzing community. This paper is already having tremendous impact on fuzzing research, setting the standard for how evaluations should be done. It is bringing scientific principles, such as rigorous conclusions and reproducibility, to an area in need of scientific understanding.

Two additional papers were selected for an Honorable Mention recognition.

The first paper, *Continuous Formal Verification of Amazon s2n* by Andrew Chudnov, Stephen Magill, and Eric Mullen, was presented at the International Conference on Computer Aided Verification, 2018. This paper discussed the researchers' advancements in the application of formal methods. Formal methods are security techniques where mathematical approaches are used to verify that software is written correctly.



These techniques are often applied to security critical code. The researchers reduced the manual effort required in revaluating code after modifications. This makes it practical to continuously verify the functionality of code during the development and maintenance stages of the software lifecycle, greatly increasing the use of formal methods and their security guarantees. This paper was selected for an honorable mention because this paper's impact on security assurance is very high as it addresses one of the enduring challenges of formal methods. The researchers make a foundational research advancement and also provide a solution. This paper describes the first integration of full formal specification and verification into a modern continuous integration development environment. The researchers enable formal methods approaches to be useful and practical in real life, a testament to the researchers' breakthroughs. Their breakthrough work will enable developers in many fields to incorporate formal methods into their software development processes.

The second paper to receive an Honorable Mention was *Meltdown: Reading Kernel Memory from User Space* by Moritz Lipp, Michael Schwarz, Anders Fogh, Daniel Genkin, and Werner Haas. This paper was presented at the 2018 USENIX Security Symposium. The paper detailed an attack to exploit a side effect of a computer processors' out-of-order execution performance feature. The vulnerability allowed an attacker's program to read almost any memory on a computer system including sensitive information such as password and cryptographic keys. Meltdown received worldwide attention when the information became public in 2018. This paper was selected because it is an excellent example of researchers challenging and disproving assumptions of the security community. The Meltdown exploit proved that the community had oversimplified the model of a processor. Instead of reasoning at the architecture level, the community had to look at the micro-architecture level. This paper publicly opened up this new area of security research. The paper competition also recognizes the work the authors did to responsibly disclose this far-reaching vulnerability

About the Competition

The National Security Agency (NSA) established the Annual Best Scientific Cybersecurity Paper Competition in 2012 to encourage the development of the scientific foundations of cybersecurity. These needed foundations in cybersecurity are constructed from a broad set of disciplines and thus the paper competition is open to all fields that contributes to security science.

The process of the competition begins with the paper nominations. Papers published in peer-reviewed journals, magazines, or technical conferences the previous year, 2018, are eligible to be nominated by the public. Self-nominations

(authors and co-authors) are not eligible for award. Those that nominate a paper provide a nomination statement. The nomination statement describes the scientific contribution of the paper and explain why this paper merits the award. For the 7th competition, 37 papers were nominated. These papers were reviewed by a set of NSA experts and a committee of Distinguished Experts on the basis of scientific merit, significance of the work reported, and the degree to which the paper exemplifies how to perform and report scientific research in cybersecurity. The Distinguished Experts are leaders in cybersecurity research who volunteer their time to evaluate papers. The Experts' individual assessments are provided to the NSA Research Director who decides the winning paper and possible honorable mentions.

The following individuals served as Distinguished Experts for the 7th annual competition:

- Dr. Whitfield Diffie, Cybersecurity Advisor
- Dr. Dan Geer, In-Q-Tel
- Dr. Eric Grosse, Cybersecurity Advisor
- Dr. John Launchbury, Galois, Inc.
- Prof. Stefan Savage, University of California, San Diego
- Mr. Phil Venables, Goldman Sachs
- Dr. Arun Vishwanath, Cybersecurity Advisor
- Prof. David Wagner, University of California at Berkeley
- Ms. Mary Ellen Zurko, MIT Lincoln Laboratory

Awards Ceremony



Winners for "Evaluating Fuzz Testing": Dr. Andrew Ruef, Prof. Michael Hicks, Mr. George Klees, Dr. Deborah Frincke, NSA Research Director, Dr. Shiyi Wei

On 10 October, the National Security Agency recognized the winners of the Science of Security 7th Annual Best Scientific Cybersecurity Paper Competition at a ceremony which included a presentation of the winning paper and a panel discussion. Also visiting NSA were the authors from the two Honorable Mentions who were recognized for their awards as well.

Dr. George Coker, Chief of NSA's Laboratory for Advanced Cybersecurity Research, remarked that the papers being recognized demonstrated scientific rigor and directly addressed the cutting edge of cybersecurity.

Dr. Deborah Frincke, NSA Research Director, emphasized the importance of Science of Security and the Paper Competition to NSA, adding that leaders of NSA's new Cybersecurity Directorate would be meeting with the papers' authors after the recognition ceremony to discuss their work. Dr. Frincke, who annually reads the top ten papers, stressed that these ten were at least as good as the papers that had won the previous six Paper Competitions, calling them "stellar papers". She pointed out several other aspects of this competition including that this was only the fourth time that Honorable Mentions had been awarded, the fifth time that an international paper had been recognized, the



Dr. Hicks presented the research for the winning paper

Details on the Paper Competition and the winning papers can be found at https://cps-vo.org/group/sos/papercompetition



Honorable Mention for "Continuous Formal Verification of Amazon s2n": Dr. Eric Mullen, Dr. Stephen Magill, Dr. Andrey Chudnov, Dr. Deborah Frincke, NSA Research Director

first time there was a repeat winner (Dr. Michael Hicks of the University of Maryland), and the first time that one of the authors was not yet in college when they co-authored a winning paper (George Klees, now a student at the University of Maryland). Dr. Frincke described the winning paper as an example of "steely-eyed determination" for its focus on establishing a scientific basis for the findings in the literature under review.

Following a presentation of the paper by Dr. Hicks, Dr. Adam Tagert, Technical Lead, Science of Security and Privacy, moderated a panel discussion between the papers' authors in attendance and the audience of NSA cybersecurity experts. In response to Dr. Tagert's question, "what would you suggest to practitioners about what tool to use and how to begin?" the panelists emphatically responded "Don't give up!" They recommended looking at blogs and trying multiple tools. Dr. Coker's question about differences in tools led to a spirited discussion that compared and contrasted the available fuzzing tools and how they may be applied. Researchers might consider using a variety of analytic tools to verify bugs and suggest seeds for testing. Neal Ziring, Technical Director of Cybersecurity, raised the issue of metrics, asking whether the authors noted other types of potentially useful metrics in the literature they reviewed. The authors noted that some heuristics haven't been tested and therefore may not provide good metrics; other possible metrics might require after-the-fact triage. A question on whether we can envision a time when a fuzzer will actually create a patch led the panelists to point out that some researchers are, in fact, doing that, and while it's doable, the approach merits further discussion. Another question about how to make best practices more broadly available prompted the panelists to recommend the Empirical Evaluation Checklist from SIGPLAN, and they encouraged reviewers to apply the checklist in evaluating papers thereby promoting the publication of sound results.



Honorable Mention for "Meltdown: Reading Kernel Memory from User Space.": Prof. Daniel Genkin, Mr. Werner Haas, Mr. Michael Schwarz, Dr. Deborah Frincke, NSA Research Director, Mr. Moritz Lipp, Mr. Anders Fogh



Dr. Adam Tagert, Technical Lead, Science of Security and Privacy, moderated the panel discussion

Intel International Science and Engineering Fair



NSA ISEF Award Winners at the ISEF Special Award Organizations Ceremony. (L-R): Dr. Adam Tagert, AnaMaria Perez, Suha Hussain, John Piwinski, Gabrielle Liu, Daniel Santiago, Divya Amirtharaj, Kaitlyn Zuravel, Alice Guo, Milidu Jayaweera, Kevin Meng, and NSA Executive Director, Harry Coker

Led by the Science of Security and Privacy (SoS) initiative, the NSA Research Directorate again sponsored awards at the Intel International Science and Engineering Fair (ISEF) held in Phoenix, Arizona from May 12-17, 2019. ISEF is the largest pre-college science and engineering fair in the world and there were approximately 1,800 students from 80 countries. To qualify for ISEF, a student must have won a position from an affiliated regional fair, which may be at the county, state, or country level. Participants are judged on various criteria to include skill set, creative ability, presentation clarity and scientific thought. The finalists this year competed for nearly $5 million in awards and scholarships.

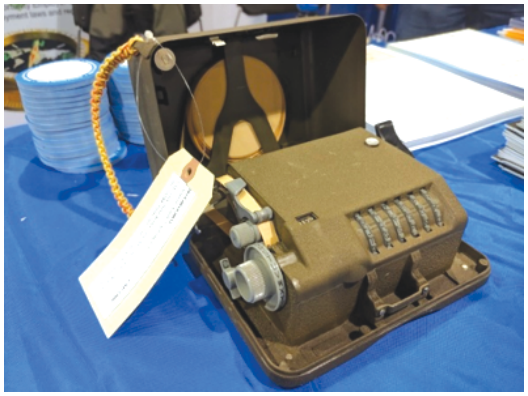
This is the fifth time the SoS initiative has organized NSA participation at ISEF. The SoS initiative is dedicated to developing a scientific discipline focused on cyber security, and the Research Directorate uses ISEF to spark student interest in research. These awards were created to encourage high school students to pursue scientific research in cybersecurity and related fields. This year, NSA was the only Intelligence Community participant at ISEF and the only organization to award cybersecurity prizes. A total of $12,000 was awarded to 17 ISEF participants in the categories of Science of Security, Mathematics, Cyber Pioneer, and Material Science. The NSA Executive Director, Harry Coker, presented the NSA awards to the winning students.

The team of NSA judges from the Laboratory for Advanced Cybersecurity Research and other research groups reviewed over 1850 student submissions presentations (papers, posters, and supplemental information) to narrow down the field. To select the top projects, the NSA judges conducted student interviews and reexamined all the student materials before awarding the 17 prizes. Notably, four of the 2019 winners had received Research Directorate ISEF awards in 2018. The judges also recognized approximately 60 ISEF finalists with an NSA Lapel Pin.

In addition to judging the entries, the NSA team organized an Agency booth that promoted Science, Technology, Engineering, and Math (STEM) outreach to over 8,000 students, parents, and teachers in attendance. Booth attractions included two World War II cipher technologies: an Enigma machine and a CSP-1500 device. These machines provided opportunities to teach cryptologic history and link NSA's commitment to STEM for national and international security. Other booth activities included a mathematics research festival where team members leveraged women in mathematics society materials to promote hands-on math activities.



Enigma Machine


CSP1500


The SoS Outreach page in Science News

To introduce interested students in careers at NSA, four NSA attendees served on a Working at NSA panel moderated by Adam Tagert. In addition to answering questions from the moderator, panelists answered questions from students and teachers. Topics included: How were you hired? What do you do? Do you travel? Do you carry a gun? What's the difference between CIA and NSA? and How do I work at the NSA?

For the second year, published a STEM outreach page in Science News magazine. Science News is published by Society for Science and Public (SSP), the organizers of ISEF. Science News is given to each attendee at ISEF as well as distributed to schools and subscribers and has over 200,000 readers. While the 2018 outreach page highlighted methods to do cybersecurity science fair projects, the 2019 page provided a coded message for subscribers to decode. The page contained only 3 elements: 1) A 15x15 grid of letters; 2) a QR code; and 3) the NSA and SoS logos. Like real secret messages, the method of decoding was not specified. The page did contain a hidden hint, the QR code, which contained an embedded URL to a SoS VO page containing information about NSA and Caesar Ciphers. This page taught the concepts so that a student could solve the message. The message was encoded using a Caesar Cipher where the key was 11, 4, 19, 5, 2 or the NSA's founding date of November 4, 1952. This number was chosen to reflect the human behavior Hard Problem of cybersecurity. Specifically, it reflected that people choose poor passwords such as birth dates. The message was:

HASHTAG NAILED IT WE ARE CODE MAKERS AND CODE BREAKERS OUR TWO FOLD MISSION IS TO COLLECT FOREIGN SIGNALS INTELLIGENCE AND TO DEFEND NATIONAL SECURITY SYSTEMS TO JOIN US IN THE FIGHT AND LEARN MORE ABOUT OUR MISSION VISIT CPSDASHVODOTORGSLASHTHERESNOTHINGTOSEEHERE

The embedded url, (cps-vo.org/nothingtoseehere) was a webpage on the SoS-VO that congratulated the person on solving the puzzle and provided the SoS email alias if they wanted to let us know how they solved the puzzle. About a dozen people emailed saying they solved the puzzle.

NSA Research Directorate Awards

Science of Security

First Place:
Suha Hussain, Ozone Park, NY (repeat winner)
Detecting Privacy Violations in Children's Apps Using HPCs

Second Place:
Divya Amirtharaj, Beaverton, OR (repeat winner)
A Secure Implementation of Mendelian Randomization via Multi-Party Computation

Alice Guo, Morgantown, WV
Protection of Deep Neural Networks against Adversarial Attacks with Application to Facial Recognition

Gabrielle Liu, Nashville, TN
Weight Friction: A Simple Method to Overcome Catastrophic Forgetting and Enable Continual Learning in Neural Networks

Honorable Mention:
Michael Brockman, Bartlesville, OK
The Next Artificial Intelligence Revolution: AI Making Decisions without Human Models or Knowledge of Rules to Create Completely Independent Solutions

Adam Kelly, Dublin, Ireland
Improved Gate Level Simulation of Quantum Circuits

Advay Koranne, West Linn, OR
Non-Periodic Pseudo-Random Number Generator Using Sinai Billiards

Cyber Pioneer

First Place:
Kevin Meng, Plano, TX (repeat winner)
Looking through Walls with Artificial Intelligence: An Innovative Solution for Real-Time Retrieval of the Human Figure behind Visual Obstruction

Second Place:
Milidu Jayaweera, Albuquerque, NM
A Game of Jamming: A Multi-Agent Game Theoretic Learning BasedCognitive Anti-Jamming Communication System to Combat an AI Jammer

Honorable Mention:
Ryan McDowell, Annapolis, MD
Should I Trust What's in My Computer? Using Current Draw Analysis to Identify Malicious Firmware in Solid State Drives

Material Science

First Place:
Kaitlyn Zuravel, Fayetteville, NC
Glue Busters II: The Effects of Accelerated Cure Time on the Ultimate Shear Strength and Efficiency of CA and PVA Glue

Honorable Mention:
Jacqueline Prawira, Mountain House, CA
OceanBioplas: The Plasticity of Marine Exoskeleton-Inspired Materials and Their Degradability in the Environment (Soil and Seawater/Saltwater)

Mathematics

First Place:
Daniel Santiago, Anasco, Puerto Rico (repeat winner)
Loop Spaces, P-Curvature, and Homotopy

Second Place:
AnaMaria Perez, Albuquerque, NM
Classifying Magic Squares and Their Associated Symmetries Using a --Chord Diagram Approach

John Piwinski, San Antonio, TX
On the Application of Heat Diffusion across a Manifold for Dimensionality Reduction

Honorable Mention:
Aayush Karan, Muskego, WI
Generating Set for Nonzero Determinant Links under Skein Relation

Brendon Matusch, Ontario, Canada
Improving Particle Classification in WIMP Dark Matter Detection Experiments Using Neural Networks

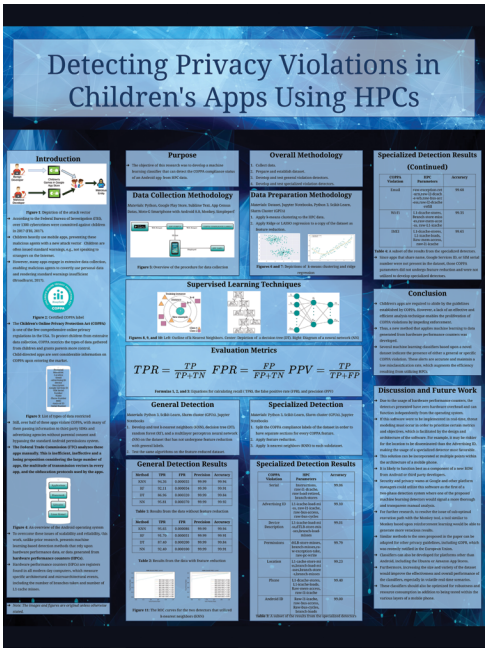
Ten of the seventeen ISEF awardees, along with their chaperones, visited NSA July 30-31, 2019. In addition to a poster session with NSA researchers, the awardees presented their research to NSA Director GEN Nakasone and NSA's Board of Directors.



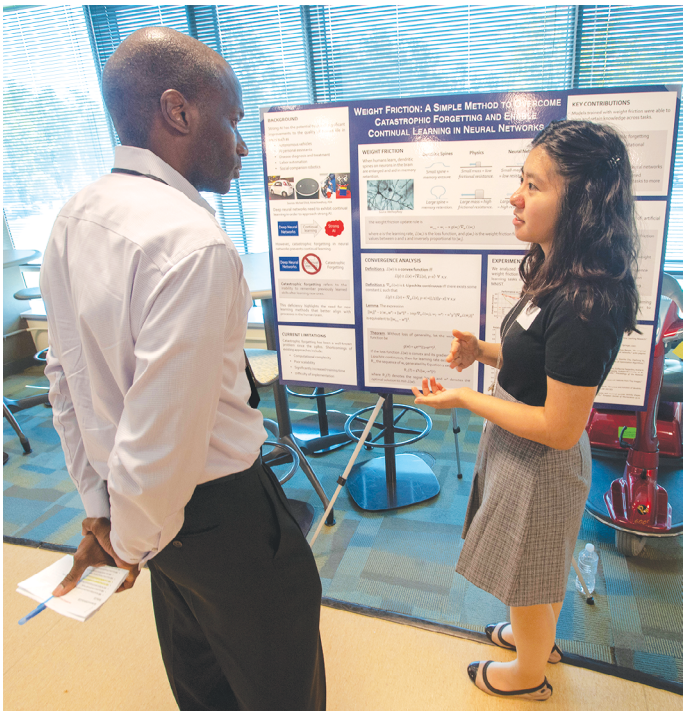
ISEF winners are pictured with the Director of NSA, GEN Paul Nakasone (far right), and the Deputy Director of NSA, George Barnes (far left). Students (L-R): Divya Amirtharaj, Alice Guo, Suha Hussain, Adam Kelly, Advay Koranne, Gabrielle Liu, Ryan McDowell, Kevin Meng, Daniel Santiago, Kaitlyn Zuravel

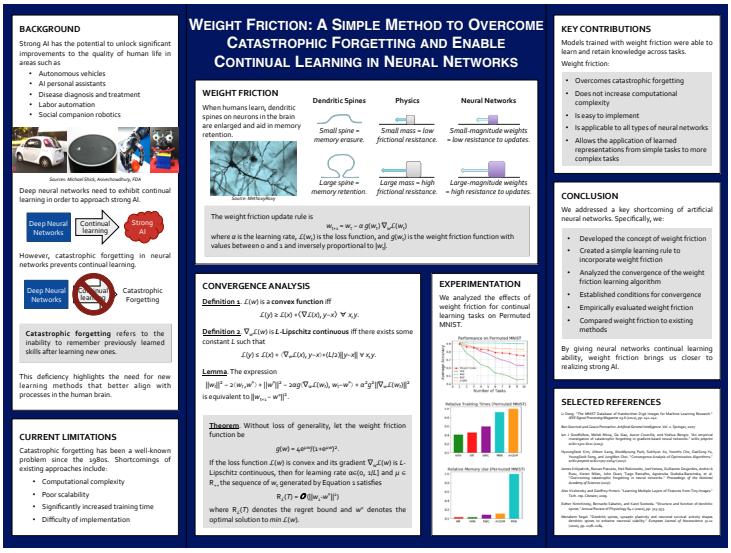
Poster Presentations





Suha Hussain, first place winner in the SoS category, presents her research



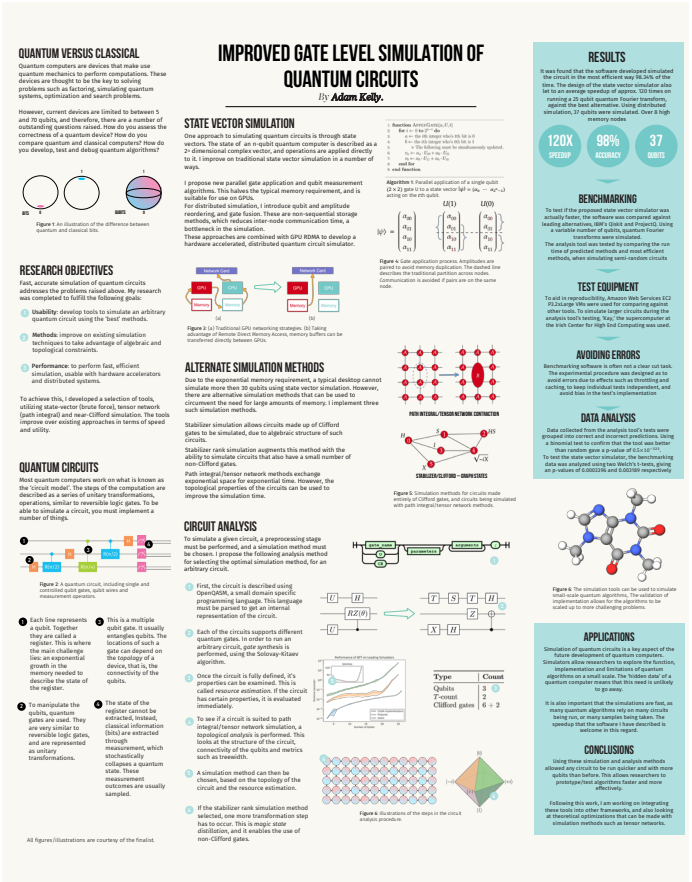




Suha Hussain, first place winner in the SoS category, presents her research to Dr. Deborah Fricke, NSA Research Director, and other NSA researchers



Adam Kelly presents his research to Dr. Deborah Frincke, NSA Research Director, and NSA researchers



IMPROVED GATE LEVEL SIMULATION OF QUANTUM CIRCUITS

By **Adam Kelly.**

QUANTUM VERSUS CLASSICAL

Quantum computers are devices that make use quantum mechanics to perform computations. These devices are thought to be the key to solving problems such as factoring, simulating quantum systems, optimization and search problems.

However, current devices are limited to between 5 and 70 qubits, and therefore, there are a number of outstanding questions raised. How do you assess the correctness of a quantum device? How do you compare quantum and classical computers? How do you develop, test and debug quantum algorithms?

RESEARCH OBJECTIVES

Fast, accurate simulation of quantum circuits addresses the problems raised above. My research was completed to fulfill the following goals:

1. **Usability:** develop tools to simulate an arbitrary quantum circuit using the 'best' methods.
2. **Methods:** improve on existing simulation techniques to take advantage of algebraic and topological constraints.
3. **Performance:** to perform fast, efficient simulation, usable with hardware accelerators and distributed systems.

To achieve this, I developed a selection of tools, utilizing state-vector (brute force), tensor network (path integral) and near-Clifford simulation. The tools improve over existing approaches in terms of speed and utility.

QUANTUM CIRCUITS

Most quantum computers work on what is known as the 'circuit model'. The steps of the computation are described as a series of unitary transformations, operations, similar to reversible logic gates. To be able to simulate a circuit, you must implement a number of things.

1. Each line represents a qubit. Together they are called a register. This is where the main challenge lies: an exponential growth in the memory needed to describe the state of the register.
2. To manipulate the qubits, quantum gates are used. They are very similar to reversible logic gates, and are represented as unitary transformations.
3. This is a multiple qubit gate. It usually entangles qubits. The locations of such a gate can depend on the topology of a device, that is, the connectivity of the qubits.
4. The state of the register cannot be extracted. Instead, classical information (bits) are extracted through measurement, which stochastically collapses a quantum state. These measurement outcomes are usually sampled.

All figures/illustrations are courtesy of the finalist.

STATE VECTOR SIMULATION

One approach to simulating quantum circuits is through state vectors. The state of an n-qubit quantum computer is described as a 2^n dimensional complex vector, and operations are applied directly to it. I improve on traditional state vector simulation in a number of ways.

I propose new parallel gate application and qubit measurement algorithms. This halves the typical memory requirement, and is suitable for use on GPUs.

For distributed simulation, I introduce qubit and amplitude reordering, and gate fusion. These are non-sequential storage methods, which reduces inter-node communication time, a bottleneck in the simulation.

These approaches are combined with GPU RDMA to develop a hardware accelerated, distributed quantum circuit simulator.

ALTERNATE SIMULATION METHODS

Due to the exponential memory requirement, a typical desktop cannot simulate more than 30 qubits using state vector simulation. However, there are alternative simulation methods that can be used to circumvent the need for large amounts of memory. I implement three such simulation methods.

Stabilizer simulation allows circuits made up of Clifford gates to be simulated, due to algebraic structure of such circuits.

Stabilizer rank simulation augments this method with the ability to simulate circuits that also have a small number of non-Clifford gates.

Path integral/tensor network methods exchange exponential space for exponential time. However, the topological properties of the circuits can be used to improve the simulation time.

CIRCUIT ANALYSIS

To simulate a given circuit, a preprocessing stage must be performed, and a simulation method must be chosen. I propose the following analysis method for selecting the optimal simulation method, for an arbitrary circuit.

1. First, the circuit is described using OpenQASM, a small domain specific programming language. This language must be parsed to get an internal representation of the circuit.
2. Each of the circuits supports different quantum gates. In order to run an arbitrary circuit, gate synthesis is performed, using the Solovay-Kitaev algorithm.
3. Once the circuit is fully defined, it's properties can be examined. This is called *resource estimation*. If the circuit has certain properties, it is evaluated immediately.
4. To see if a circuit is suited to path integral/tensor network simulation, a topological analysis is performed. This looks at the structure of the circuit, connectivity of the qubits and metrics such as treewidth.
5. A simulation method can then be chosen, based on the topology of the circuit and the resource estimation.
6. If the stabilizer rank simulation method selected, one more transformation step has to occur. This is magic state distillation, and it enables the use of non-Clifford gates.

Type	Count
Qubits	3
T-count	2
Clifford gates	6 + 2

RESULTS

It was found that the software developed simulated the circuit in the most efficient way 98.34% of the time. The design of the state vector simulator also let to an average speedup of approx. 120 times on running a 25 qubit quantum Fourier transform, against the best alternative. Using distributed simulation, 37 qubits were simulated. Over 8 high memory nodes

120X SPEEDUP | **98%** ACCURACY | **37** QUBITS

BENCHMARKING

To test if the proposed state vector simulator was actually faster, the software was compared against leading alternatives, IBM's Qiskit and ProjectQ. Using a variable number of qubits, quantum Fourier transforms were simulated.

The analysis tool was tested by comparing the run time of predicted methods and most efficient methods, when simulating semi-random circuits

TEST EQUIPMENT

To aid in reproducibility, Amazon Web Services EC2 P2.24xLarge VMs were used for comparing against other tools. To simulate larger circuits during the analysis tool's testing, 'Kay,' the supercomputer at the Irish Center for High End Computing was used.

AVOIDING ERRORS

Benchmarking software is often not a clear cut task. The experimental procedure was designed as to avoid errors due to effects such as throttling and caching, to keep individual tests independent, and avoid bias in the test's implementation

DATA ANALYSIS

Data collected from the analysis tool's tests were grouped into correct and incorrect predictions. Using a binomial test to confirm that the tool was better than random gave a p-value of 6.5×10^{-19}.

To test the state vector simulator, the benchmarking data was analyzed using two Welch's t-tests, giving an p-values of 0.0003196 and 0.003169 respectively

APPLICATIONS

Simulation of quantum circuits is a key aspect of the future development of quantum computers. Simulators allow researchers to explore the function, implementation and limitations of quantum algorithms on a small scale. The 'hidden data' of a quantum computer means that this need is unlikely to go away.

It is also important that the simulations are fast, as many quantum algorithms rely on many circuits being run, or many samples being taken. The speedup that the software I have described is welcome in this regard.

CONCLUSIONS

Using these simulation and analysis methods allowed any circuit to be run quicker and with more qubits than before. This allows researchers to prototype/test algorithms faster and more effectively.

Following this work, I am working on integrating these tools into other frameworks, and also looking at theoretical optimizations that can be made with simulation methods such as tensor networks.

Growing the Science of Security

The Science of Security and Privacy (SoS) initiative sees the need for a scientific basis for cybersecurity as a large-scale problem that one entity alone cannot solve--a community is needed. In 2019 the SoS continued its efforts to grow the Science of Security community of interest.

In 2019 the SoS team began engagements with other research universities to expand the current partnership network and tap into existing national pipelines of diverse cybersecurity talent. The team recruited nine universities to attend Lablet quarterly meetings to collaborate with the current six Lablets, twelve Sub-Lablets, and industry partners. This new tier of cybersecurity research partner universities are identified as Associate Lablets. Associate Lablets are universities that are active members of the SoS community. Not directly funded by the Science of Security and Privacy initiative, they participate in SoS community events (Quarterly Meetings, HotSoS, etc.) and collaborate with Lablet researchers at Lablet and/or Sub-Lablet schools to enhance cybersecurity research. Collaborations are not directed by NSA. Associate Lablets consist of the following universities:

- Alabama A&M University
- Columbia University
- Drexel University
- Harvard University
- Jackson State University
- Morgan State University
- Naval Postgraduate School
- North Carolina A&T University
- Simmons University

The above Associate Lablets include four Historically Black Colleges/Universities (HBCUs), one women's university, two Ivy League schools, and one military graduate school. The team leveraged existing NSA academic partnerships with the Associate Lablets to assess their cybersecurity research capabilities and identified them as candidate institutions to invite to Lablet quarterly meetings.

For the sixth year, SoS sponsored the Hot Topics in the Science of Security Symposium (HotSoS) which was held at Vanderbilt University on 2-3 April 2019. HotSoS is held in cooperation with the Association for Computer Machinery (ACM), and over 100 participants from government, industry, and academia attended this year's meeting.

The SoS-Virtual Organization (SoS-VO), a longstanding initiative designed to grow the Science of Security community, grew to over 1700 members in 2019. It continues to provide a centralized location for cybersecurity research, events, and news.

SoS also supports a variety of focused outreach efforts. The monthly Science of Security Reviews and Outreach (R&O) newsletter which links to the SoS-VO now has over 1700 subscribers. This year also saw a growing SoS presence on social media, with the SoS Facebook group nearing 100 members. SoS again co-sponsored with Vanderbilt University the Cyber-Physical Systems Security Summer Camps which expanded to include teachers in 2019. The SoS initiative is routinely raised by SoS team and Lablet researchers at international conferences. As a result of all of these efforts, the SoS community continued to expand. Details on HotSoS, and other outreach initiatives are found in the following pages.



Hot SoS



The seventh Hot Topics in the Science of Security (HotSoS) Symposium was hosted by Vanderbilt University (VU) on April 2-3, 2019 in Nashville, Tennessee. This community meeting is sponsored by the NSA Science of Security and Privacy (SoS) initiative. It assembles participants from academia, government, and industry to interact and learn about the latest rigorous scientific approaches to prevent, detect, and mitigate cyber threats. The symposium was a mix of refereed papers, invited talks, panel discussion, tutorial, and paper discussion workshop as well as a session on Works-in-Progress and Works-already-Published (WiP/WaP).

Xenofon Koustoukos, VU Principal Investigator (PI) served as General Chair of HotSoS 2019 along with Program Co-Chairs Alvaro Cardenas of the University of Texas at Dallas and Ehab Al-Shaer of the University of North Carolina Charlotte. The HoTSoS 2019 program included seven papers, six WiP/WaP papers as well as three special workshop papers. The papers included the work of 43 authors representing 19 universities. Along with the papers there were 21 posters submitted by 66 authors from 14 universities. Both the best paper and the best poster were selected on use of strong research methodology, innovation and relevance to SoS. There was a tutorial on Game Theory for Cyber Detection prior to the formal opening of HotSoS.

HotSoS 2019 proceedings have been published by ACM and are available online at the ACM digital library: https://dl.acm.org/doi/proceedings/10.1145/3314058

Keynote Presentations

1. Accountability in Computing
Rebecca Wright, Barnard College and Rutgers University

Prof. Wright's presentation explored using accountability as a useful paradigm to shift away from prevention-based approaches. She stated that "accountability" is used to describe computer-security mechanisms that complement preventive security, but the word lacks a precise, agreed-upon definition. Current definitions focus on such concepts as detection, evidence for third parties, judgment/blame, punishment, associating actions and identifications, and answerability. She proposed using punishment as a deterrent, not as retribution and proposed a definition of accountability as follows: an entity is accountable with respect to policy whenever the entity violates it, then the entity is or could be punished in a targeted fashion. She said the research goal is to explicate the relationship between accountability and related properties: identification, authorization, detection, closed systems, punishment, and compensation. She concluded her presentation by noting that today's technology creates opportunities for convenience and connection, but also for security breaches, invasion of privacy, and weaponization of information. She said that accountability is an important goal, both for individuals and for providers of technical applications and infrastructure. Finally, insights gained will lead to better system design and more informed choices by decision makers and users.

2. Toward a Science of Cyberdeception
Kevin Hamlen, University of Texas at Dallas

Prof. Hamlen said that cybersecurity is the most asymmetric form of warfare with attackers having a broad choice of targets, plenty of time to probe defenses, and the need to find only one critical vulnerability. Defenders, on the other hand, have only seconds to react to attacks and must close all vulnerabilities. He proposes deception to level the asymmetry, quoting Sun Tzu that all warfare is based on deception. The goal of a cyber deceptive defense is to give the attacker false certainty (rather than confusion) through refining defenses, collecting reconnaissance and anticipating the attacker's future moves. To succeed, a more rigorous science of software cyber deception is needed. This approach requires an interdisciplinary, data driven science that includes risk management, machine learning, computer networking, cognitive sciences, game theory, and the economics of security. He cited cyber deception science challenges as automation, assimilation, adaptation, and assurance. He suggested that using the approach he discussed can enable the transition of cyber deception from an art to a science.

3. The Science of Attack Surfaces and its Applications
Trent Jaeger, The Pennsylvania State University

Prof. Jaeger began by making several points about attack surfaces: new, unexpected attack surfaces may be created by administrators and others; programmers are not proactive about identifying and defending their programs' attack surfaces, and; adversaries seem to exploit nearly every undefended attack surface. He gave several reasons why he believes the study of attack surfaces relates to the science of security including because they can be computed systematically and tracked

at runtime to aid detection and can be leveraged to drive improvement in software security. One conclusion he presented was that vulnerabilities are caused by a combination of factors, but a key factor is that the adversary must have access to run code to exploit flaw. He also said that the lack of knowledge of attack surfaces leaves programs open to a variety of attacks.

Research Papers

Seven refereed papers were selected for presentation. For the second year, HotSoS had a Best Paper Award which recognized the paper that exhibits outstanding achievement in science. The selection committee was made up of the HotSoS 2019 Program Co-Chairs and two members of the NSA Research Directorate. The winning paper receives an automatic nomination into the 8th Annual Best Scientific Paper Competition. The paper "Integrated Data Space Randomization and Control Reconfiguration for Securing Cyber-Physical Systems" by Bradley Potteiger, Zhenkai Zhang and Xenofon Koutsoukos, Vanderbilt University, received the 2019 HotSoS Best Paper Award. This paper was written by the researchers from the Vanderbilt University Lablet.

1. Integrated Data Space Randomization and Control Reconfiguration for Securing Cyber-Physical Systems
Bradley Potteiger, Zhenkai Zhang, and Xenofon Koutsoukos, Vanderbilt University

Non-control data attacks have become widely popular for circumventing authentication mechanisms in websites, servers, and personal computers. Moreover, in the context of Cyber-Physical Systems (CPS), attacks can be executed against not only authentication but also safety. With the tightly coupled nature between the cyber components and physical dynamics, any unauthorized change to safety-critical variables may cause damage or even catastrophic consequences. Moving Target Defense (MTD) techniques such as Data Space Randomization (DSR) can be effective for protecting against various types of memory corruption attacks including non-control data attacks. However, in terms of CPS it is also critical to ensure timely Cyber-Physical interactions after attacks thwarted by MTD. This paper addresses the problem of maintaining system stability and security properties of a CPS in the face of non-control data attacks by developing a DSR approach for randomizing binaries at runtime, creating a variable redundancy-based detection algorithm for identifying variable integrity violations, and integrating a control reconfiguration architecture for maintaining safe and reliable operation. The researchers' security framework is demonstrated by utilizing an autonomous vehicle case study. This work was supported by the NSA SoS Initiative

2. A Copland Attestation Manager
Adam Petz and Perry Alexander, University of Kansas

Copland is a domain specific language designed for describing, analyzing and executing attestation protocols. The Copland framework is a joint effort with NSA and John Hopkins University Applied Physics Lab. Its formal semantics define evaluation, sequencing, and dispatch of measurements resulting in evidence describing a system's state. That evidence is in turn appraised to determine if and how an external system will interact with it. The contribution of this work is a description of the first Copland interpreter and the attestation manager built around it. This work was supported by the NSA SoS Initiative.

3. BEEER: Distributed Record and Replay for Medical Devices in Hospital Operating Rooms
Avesta Hojjati, Yunhui Long, Soteris Demetriou, and Carl A. Gunter University of Illinois Urbana-Champaign and Imperial College London

Medical devices in hospital operating rooms are getting increasingly interconnected which enables them to download instructions and report results with less risk of error compared to traditional manual techniques. However, many of these devices are safety-critical and, thus, any risks from cyberattacks can be extremely high. This paper describes BEEER, a distributed record and replay framework suitable for environments where more than one safety critical device is in simultaneous use. A prominent example is a hospital operating room where a number of networked devices work together. In such scenarios, a key step to forensically analyze an incident is understanding the causality of events produced by devices. BEEER orders events during recording by leveraging the fact that it takes significantly more time for a drug's effects to come to prominence on a patient compared to device-to-device communication on a local network. During replay, BEEER uses a newly developed token mechanism to coordinate execution of the events. The researchers implemented and evaluated a prototype of BEEER and found that synchronization for short medical operations can be achieved using Network Time Protocol (NTP). For longer operations they designed and developed a new event ordering protocol (projection protocol) based on vector timestamps. BEEER's replay mechanism is efficient and therefore suitable for forensics analyses in practice. This work was supported by the National Science Foundation.

4. Limitations on Observability of Effects in Cyber-Physical Systems
Suresh K. Damodaran and Paul Rowe, MITRE Corporation

Increased interconnectivity of Cyber-Physical Systems, by design or otherwise, increases the cyberattack surface and attack vectors. Observing the effects of these attacks is helpful in detecting them. This paper shows that many attacks on such systems result in a control loop effect, Process Model Inconsistency (PMI). The formal approach elucidates the relationships among incompleteness, incorrectness, safety, and inconsistency of process models. Incomplete process models are demonstrated to lead to inconsistency. Surprisingly, inconsistency may arise even in complete and correct models. The approach is illustrated through an Automated Teller Machine (ATM) example, and characterize the practical implications of the theoretical results.

5. Game Theory for Adaptive Defensive Cyber Deception
Kimberly Ferguson-Walter, U.S. Department of Defense
Sunny Fugate, Justin Mauger, Maxine Major, Space and Naval Warfare Systems Center Pacific

In prior game theory work, cyber deception games have generally been modeled as non-cooperative, sequential games, where the quality and authenticity of the signal is being manipulated. This work relies on a signal, which can be deceptive, being actively pushed to the attacker. The researchers have formulated a new framework for cyber deception games in which each player has their own perception of the game being played and the moves being taken. A particular player may manipulate other players' perceived payoffs to induce them to take sub-optimal actions. This model of deception seeks to address defender interactions

with an attacker following the initial deployment of cyber defenses addressed by previous work. The primary contribution is a model of defensive cyber deception that incorporates defender control of attacker perception of the cyber environment. The ultimate goal is to inform future cyber defense systems, enabling more sophisticated responses to attacker behaviors and improving defensive posture. This work was supported by the NSA Laboratory for Advanced Cybersecurity Research.

6. Situated Information Flow Theory
Sebastian Benthall, New York University

A key component of recent privacy rules is restriction on the flows of personal information or data based on information categories. This tendency conflicts with the fact that data's meaning is not stable but depends on how it was formed and with what other information it is combined. These properties of information challenge naive intuitions that information "flows" like a fluid, such as water or oil. Rather, the authors built on Dretske, Pearl, and Nissenbaum to develop Situated Information Flow Theory (SIFT), a view of information flows as causal flows with nomic associations due to a larger context of causal relations. The semantics of situated information flow are precise within the statistical framework of Bayesian networks. This new understanding of information flow has three policy implications: (1) Restrictions on data transfers are more precise and enforceable than restrictions on information flow; (2) Information "categories" or meanings must be defined relative to a particular class of observers and take into account their reasonable background information; and (3) The semantics of data are ambiguous when there is uncertainty about causal structure, and this structure is learned from data aggregation. Hence, the information asymmetry between data aggregators and individual data subjects are one reason why data processors are "opaque" and difficult to regulate.

7. Browser Fingerprinting using Combinatorial Sequence Testing
Bernhard Garn, Dimitris E. Simos, Stefan Zauner, Rick Kuhn and Raghu Kacker SBA Research, University of Applied Sciences Campus Vienna, and National Institute of Standards and Technology

This paper reports on the applicability of combinatorial sequence testing methods to the problem of fingerprinting browsers based on their behavior during a TLS handshake. Creating an appropriate abstract model of the TLS handshake protocol enabled mapping browser behavior to a feature vector. From this vector, the researchers derive a distinguisher. Using combinatorial methods, they created test sets consisting of TLS server-side messages as sequences that are sent to the client as server responses during the TLS handshake. This approach was validated with a case study showing that combinatorial properties have an impact on browsers' behavior.

Special Research Presentation

Traditionally, the winner of the previous Best Cybersecurity Paper Competition is invited to HoTSoS to present the winning work. This year, Tiffany Bao presented the work.

How Shall we Play a Game? A Game-theoretical Model for Cyber-warfare Games
Tiffany Bao, Yan Shoshitaishvili, Ruoyu Wang, Christopher Kruegel, Giovanni Vigna, and David Brumley
Carnegie Mellon University and University of California, Santa Barbara
Winner of the 6th Annual Best Scientific Cybersecurity Paper Competition

The researchers endeavored to "identify the best strategy for the use of an identified zero-day vulnerability in a cyber-warfare scenario where any action may reveal information to adversaries." They developed a game-theoretic model and the ability to quickly find optimal solutions to it. These strategies aid humans and computers in making decisions when dealing with previously unknown vulnerabilities in computer systems. The model accounts for both attack and defensive actions and imperfect information about the current status. Actions that can be taken include attacking by using this vulnerability, patching one's own systems, stockpiling for later, or taking no action. The model also develops steps for one to follow over time, such as patching one's own computers for a period and then later attacking.

Works-in Progress/Works-already Published Session
The Works in Progress/Works already Published (WiP/WaP) Session at HotSoS was introduced in 2019 and aims to provide researchers with an opportunity to receive early feedback on emerging research. During this session authors discuss their ideas, present ongoing work, gather feedback from experts, and/or introduce work published in journals or elsewhere that is relevant to the science of security community. Papers were solicited that dealt with the five Hard Problems as well as privacy, Cyber-Physical Systems, and the Internet of Things. Submissions with an emphasis on scientific work which examines the foundations of trustworthy systems for security and privacy were especially encouraged. WiP submissions focused on ongoing, unpublished work discussing evolving and early-stage ideas, new research directions, and current trends in the area. The intention of WaP presentations was to raise awareness of work published elsewhere that is relevant to the science of security community.

The following papers were presented in this session:

- *Attacks on Electricity Markets* Carlos Barreto and Xenofon Koutsoukos
- *Ethics in Norm-Aware Agents* Nirav Ajmeri, Hui Guo, Pradeep Kumar Murukannaiah and Munindar Singh
- *A Graph-Based Analysis of Industrial Control Systems Network Traffic* Imani Palmer, Shane McFly and Edmond Rogers
- *A Multidimensional Multilevel Model for Smart Environment* Amir Modarresi, John Symons and James Sterbenz
- *Search Prevention using Captchas against Web Search Engine: A Proof of Concept* Luke Sample and Donghoon Kim

- *The Tularosa Study: An Experimental Design and Implementation to Quantify the Effectiveness of Cyber Deception* Kimberly Ferguson-Walter, Temmie Shade, Andrew Rogers, Elizabeth Niedbala, Michael Trumbo, Kevin Nauer, Aaron Jones, Angela Combs and Robert Abbott

Special Workshop Session

The goal of the Special Workshop Session at HotSoS was to provide authors and the security community with early, in-depth feedback on new research. Three subjects were presented. Each manuscript had 45 minutes, in which a non-author discussant presented the manuscript for up to 10-15 minutes. The discussant identified the strengths and weaknesses along multiple dimensions:

- Is the problem framing clear, is this the right problem given the approach? Is the threat model described, the right model?
- Is the chosen analysis appropriate for the claims being made? Is there an alternative analysis better suited to the research questions or hypotheses? Do threats to internal validity exist, if so, how can those be appropriately acknowledged and/or mitigated?
- Are the results being presented clearly, what alternative presentations exist?
- Is the scope of impact consistent with the results, are connections to a broader impact missing or available for further research?

Following the discussant questions and author reactions, the remaining time was used by the audience to raise their own questions and to engage in audience discussion. Audience members prepared their own questions in advance after reading the manuscript prior to attending the workshop. The ultimate goal for each session was to provide authors with detailed actionable feedback, which they will then use to improve their manuscripts prior to submission for publication at a different venue. Since the manuscripts were in the early stages, no information on the specific material was published.

Panel Discussion

Are there fundamental truths in cybersecurity?
Moderator: Adam Tagert, NSA Panelists: Tiffany Bao, Arizona State University Dan Geer, In-Q-Tel Trent Jaeger, The Pennsylvania State University Jonathan Spring, Carnegie Mellon University

After noting that the panel discussion topic was prompted by the Best Scientific Cybersecurity Paper Competition, Dr. Tagert introduced the panelists and led off with Jonathan Spring who said that part of science is how to ask questions well. Dr. Spring discussed his paper in which he wanted to tease out why he thinks what he is doing is science, and why it's producing unsatisfactory results.

Dan Geer, one of the Distinguished Experts (DEs) who evaluated the papers in the paper competition, was asked what he thought fit into a paradigm for quality science. Dr. Geer said that the DEs evaluate papers on the basis of what they are doing to advance the not-yet-fully-formed idea of a science of security and whether the paper provides a new direction and a clarity of description for that purpose.

Dr. Tagert asked Trent Jaeger what he tried to instill in his students to become good scientists in the future. Dr. Jaeger said that if you look at a system you are likely to find some anomalies and he encourages his students to study how principles relate to that anomaly.

Dr. Tagert's question to Tiffany Bao dealt with her winning paper in the paper competition, and he asked what her team did to validate the model presented in the paper. She said that in the beginning they didn't have the validation element, but were told by reviewers that they needed to evaluate it; she noted that there needs to be collaboration in order to provide feedback to theorists.

Panelists then began to discuss various issues dealing with the science of security. In response to Jonathan Spring's assertion that considering science as a process can lead to useful questions about changes to the processes to improve it, Dan Geer noted that a lot of people also say security is a process. Trent Jaeger said that the things that are most predictive, boiled down, seem the most "science". Another comment was that many papers reflect the excitement of tool building and the possibility of using the new tool, but few papers give us the predictive power we are hoping for. A lively give-and-take among the panelists and with the audience ensued—some concluding observations were that the permanence of the result was important and that science is generalizable new knowledge.

Posters

Twenty-one posters were presented on topics including intrusion and anomaly detection, resilience, APT chains, reliability measurement, privacy leaks, moving target defenses, and others. The Best Poster was "A Synopsis of Static Analysis Alerts on Open Source Software" by Nasif Imtiaz and Laurie Williams, from NCSU. The Best Poster Award recognizes cybersecurity research with scientific rigor, clarity of presentation, and global impact. It is to encourage scientists across multiple disciplines to address the fundamental problems of security in a principled manner. The following are 19 of the posters presented along with their abstracts.

Best Poster
1. A Synopsis of Static Analysis Alerts on Open Source Software
Nasif Imtiaz and Laurie Williams North Carolina State University

Static Application Security Testing (SAST) tools detect potential

code defects (alerts) without having to execute the code. SASTs are now widely used in practice by both commercial and Open Source Software (OSS). Prior work found that half of the state-of-the-art OSS projects have already employed automated static analysis. However, little public information is available regarding the actionability (important to developers to act upon) of SAST alerts.

2. A Bird's Eye View of Knowledge Needs Related to Penetration Testing
Akond Rahman and Laurie Williams North Carolina State University

According to the National Institute of Standards and Technology (NIST), penetration testing is an assessment conducted on software systems to identify vulnerabilities that could be exploited by adversaries. Despite the importance of penetration testing in software security, practitioners search for strategies and guidance on how to get started in the domain of penetration testing. We hypothesize that practitioners have knowledge needs related to penetration testing, which can be synthesized using penetration testing-related questions posted on Questions and Answer (Q&A) websites. A systematic investigation can identify the knowledge needs of practitioners related to penetration testing, helping the cybersecurity community in advancing the field of cybersecurity education. The goal of this paper is to help cybersecurity researchers in advancing the field of cybersecurity education by analyzing penetration testing-related questions posted by practitioners.



3. A Formal Verification of Mutation-Based Moving Target Defense Muhammad Abdul Basit Ur Rahim, Qi Duan, and Ehab Al-Shaer University of North Carolina at Charlotte

Static system configuration provides the advantage for the attackers to discover the vulnerabilities of the system and launch attacks. Moving Target Defense (MTD) can break this asymmetry for the defenders by mutating certain configuration parameters proactively and at the same time maintaining the run-time correctness and operational integrity. MTD is essentially distributed by nature since the processes or actions in MTD are executed in an interleaved manner. For any distributed system, it is important to verify the correctness and integrity, since they may be jeopardized by design errors or run time inconsistencies. In this work we present a framework for formal verification of MTD techniques. We describe MTD techniques with formal ontology and model the system behaviors with timed automata, and verify the correctness, liveness, fairness and deadlock-free properties of the system. We use Random Host Mutation (RHM) as the case study or MTD formal verification. Our experimentation validates the feasibility and scalability of the formal verification framework.

4. A New Metric to Compare Anomaly Detection Algorithms in Cyber-Physical Systems
Jairo Giraldo, University of Texas at Dallas, and Alvaro A. Cardenas, University of California Santa Cruz
The performance of different anomaly detection algorithms is typically compared using metrics that depend on the True Positive Rate (TPR) and the False Positive Rate (FPR). However, to obtain the TPR it is necessary to generate attacks that will be detected, which is useless to evaluate detection strategies against more realistic adversaries that can adapt their attacks to remain undetected. On the other hand, the FPR can be misleading and hard to interpret in practical applications since the amount of time a process is observed is not fixed. In this poster, we present a novel metric that is based on the maximum impact an adversary can cause while remaining stealthy, and on the expected time between false alarms. Our metric is useful for the evaluation and comparison of anomaly detection strategies in CPS.

5. CLIPS/ActiveSDN for Automated and Safe Cybersecurity Course-of-Actions Orchestration
Md Mazharul Islam, Ehab Al-Shaer, Ashutosh Dutta, and Mohammed Noraden Alsaleh University of North Carolina at Charlotte

Continuous attack reports such as data breach, malware, phishing and spamming attacks published daily indicate that cyberattack is inevitable in our daily life. Sometimes it takes days, even months, to detect and mitigate such stealthy attacks. This requires making network systems resilient against attacks with a high assurance of defense mechanisms that can go beyond attack detection with safe mitigation. That's why we developed a flexible yet expressive policy specification language called CLIPS for Active Cyber Defense, and provably-correct policy refinement engine, and ActiveSDN to enable a safe, efficient construction and execution of Course-of-Action workflow composed of investigating for analysis and mitigating for reconfiguration actions to support cyber resilience automation.

6. Cyber Defense Matrix—A New Model for Optimal Composition of Cybersecurity Controls for Constructing Resilient Risk Mitigation Ashutosh Dutta and Ehab Al-Shaer University of North Carolina at Charlotte

Most enterprises aiming to deploy cost-effective security configuration follow information security standards and guidelines to adopt cybersecurity controls such as CIS Critical Security Controls (CIS CSC). With the increased dependency over cyber, the landscape of cyberattacks is escalating quickly, and, as a consequence, hundreds of cybersecurity controls have to be delineated to implement NIST Cybersecurity Framework (i.e., identify, detect, protect, respond and recover). The security configuration comprised of the appropriate set of security controls requires that it not only be optimized regarding Return on Investment (RoI) but also to be resilient in order to tackle the failures against diversified cyberattacks. However, the composition of such optimal and resilient cybersecurity portfolio (security configuration) is a highly complex and error-prone task as there are exponential numbers of ways to construct the portfolio due to a large number of security controls, network vulnerabilities, and budget constraints. The objective of this research is to develop a novel and an automated approach to compose the optimal and resilient risk mitigation planning by selecting the most Critical Security Controls (CSC) considering affordable residual risk (risk appetite), budget, resiliency, and enterprise-oriented usability constraints. We developed a model named Cyber Defense Matrix (CDM) that resembles the deployed cyber defense strategy. The structure of CDM incorporates three dimensions: Security Function

(what), Enforcement Level (where), and Kill Chain Phase (why) that enables the composition of multi-layer and multi-stage resilient defense configurations. Our approach leverages CDM to determine "what" security controls are needed and for which security function (Identify, Protect, Detect, Respond, and Recover), "where" each security control should be enforced in the cyber systems (Network, Device, People, Application, and Data), and "why" they are effective (i.e., against what attack and wherein the kill chain). We formulate the approach to compute the correct-by-construction cybersecurity planning as CDM using SMT constraints.

7. Design and Simulation Platform for Evaluation of Grid Distribution System and Transactive Energy Himanshu Neema, Harsh Vardhan, Carlos Barreto, and Xenofon Koutsoukos Vanderbilt University

With the advent of remarkable development of solar power panel and inverter technology as well as long-term interest in getting rid of polluting energy, countries are trying to migrate from conventional energy sources (carbon-based energy sources) to nonconventional energy sources (e.g., solar, wind, and geothermal). That gave rise of a new paradigm called Transactive Energy (TE) that utilizes economic and control techniques to effectively manage Distributed Energy Resources (DERs). Another goal of TE is to improve grid reliability and efficiency. However, to evaluate various TE approaches, a comprehensive simulation tool is needed that is easy to use and capable of simulating the power-grid along with various grid operational scenarios that occur in the Transactive Energy paradigm. In this research, we present a web-based design and simulation platform (called a design studio) targeted toward evaluation of power-grid distribution system and Transactive Energy approaches. This design studio is a step towards creating a platform to edit and visualize existing power-grid models graphically, create new power-grid network models, simulate those networks, and inject various scenario-specific perturbations to evaluate specific configurations of transactive energy simulations. The design studio provides (1) a novel Domain-Specific Modeling Language (DSML) using the Web-based Generic Modeling Environment (WebGME) for the graphical modeling of power-grid, cyber-physical attacks, and TE scenarios: and (2) a reusable cloud-hosted simulation backend using the Gridlab-D power-grid distribution system simulation too.

8. Design for an Educational Cyber Range
Ryan Brunner, Sang Keun Oh, Jesse Ramirez, Paul Houck, and Nathaniel Stickney,
United States Military Academy, and Raymond Blaine, Cyber Research Center
The "Cyber Range" concept has arisen as a method for organizations to provide realistic training in penetration testing and incident response. These ranges typically utilize virtualization to reduce cost and allow for rapid reconfiguration. However, most management software is expensive or difficult to install. This paper presents a management application using open-source software with the intent of allowing educators to rapidly create virtual ranges at minimal cost.

9. Designing a Raspberry Pi Sensor Network for Remote Observation of Wildlife
Brian Curtin, Rachelle H. David, Emmet D. Dunham, Cullen D. Johnson, Nikhil Shyamkumar, Thomas A. Babbitt, and Suzanne J. Matthews United State Military Academy

Scientists and the military need unobtrusive methods of observing wildlife. In this poster, we assess the feasibility of a Raspberry Pi sensor network for wildlife detection and monitoring. While technology for wildlife observation such as camera traps exist, they are expensive and/or require human intervention to collect the data. An inexpensive sensor network that can take pictures of wildlife with negligible human intervention will enable scientists and military personnel to discretely detect the presence of endangered wildlife. Each Raspberry Pi sensor node collects data for transmission to a remote user via an android app interface. In order to improve battery efficiency, the system features an adaptable network with four distinct sleep modes. We also explore increasing the range and durability of the network by integrating a mesh network.

10. Exploring RNNs for Analyzing ZEEK HTTP Data
Daniel Andrews, Jennifer Behn, Danielle Jaksha, Jinwon Seo, Madeleine Schneider, James Yoon, Suzanne J. Matthews, and Alexander S. Mentis, United States Military Academy, and Rajeev Agrawal, United States Army Engineer Research and Development Center
Cyber vulnerabilities pose a threat across systems in the Department of Defense. Finding ways to analyze network traffic and detect malicious behavior on a network will help keep these systems safe. This paper looks at the data collection techniques, model creation, and results of building a recurrent neural network to classify incoming traffic as normal or malicious. Additionally, it considers how the information will be best portrayed on a GUI to network administrators. The model's initial accuracy is 83.45% when trained on 500,017 connections. With increased accuracy, this tool may be used by the Department of Defense to help defend its networks.

11. Learning APT Chains from Cyber Threat Intelligence
Ghaith Husari, Ehab Al-Shaer, Ruhani Faiheem Rahman, and Bill Chu University of North Carolina at Charlotte

With rapidly evolving cyberattacks, cybersecurity specialists are actively using cyber threat intelligence to identify and respond to cyberattacks in a timely manner. However, this information will be highly useful for attack detection and mitigation if we can construct structured cyber threat intelligence and accurately generate TTP Chains to understand the steps of cyberattacks. In this poster, we present our preliminary Natural Language Processing (NLP) analysis to characterize the temporal relationship of attack actions of an APT attack to extract and construct the reported TTP chains using the popular standard, MITRE ATT&CK, and the Structured Sharing Language, STIX 2, a machine readable language that will help automated the process of understanding and responding to the cyberattacks shared in unstructured text via blogs, emails, and social media.

12. Load Forecasting with Adversarial Attacks in Power Systems using DeepForge
Xingyu Zhou, Yi Li, and Carlos Barreto, Vanderbilt University
Peter Volgyesi and Xenofon Koutsoukos, Vanderbilt University, Institute for Software Integrated Systems
For electricity markets, the profit for a retailer is based on a dynamic balance between supply and demand. As a result, an accurate demand prediction system is essential for the pricing strategy to maximize the profit. The retailers may try to reduce the uncertainties by forecasting the demand of their customers. In particular, the accuracy of the forecasts improves aggregating the demand of large groups of customers. With recent

developments of machine learning techniques especially deep learning, neural network prediction models can be constructed to gain good results. From distributed meter readings of different areas, we can build a time series network model to predict the future demand after a certain amount of time. However, the complexity of current black-box inference systems leads to vulnerabilities when exposed to adversarial attacks. Since the discovery of adversarial examples in CNNs from 2013, security and robustness has become a hot topic and valid concern. As a result, even though we can build a power load prediction model easily with current deep learning techniques, we still need to take a more cautious view on security issues beneath current machine learning applications. One difficulty for viewing and testing these kinds of security risks lies in the domain-specific knowledge and state-of-art machine learning techniques. To address this issue in the power system CPS domain, we followed a model-based design approach and developed our forecasting method and security testing framework using DeepForge.

13. Measures and Metrics for the Enforcement of Critical Security Controls: A Case Study of Boundary Defense Mohiuddin Ahmed and Ehab Al-Shaer University of North Carolina at Charlotte

The Center for Internet Security publishes a set of defense actions which forms a set of defense-in-depth best practices known as Critical Security Controls (CSC) to detect, prevent, respond, and mitigate the cyberattacks against cyber systems and networks. However, no well-defined automated measures and metrics are developed to validate the enforcement of these CSC. Additionally, it is infeasible to directly analyze implementation of security products to verify and validate the enforcement of CSC in those security products. In this poster, we are going to present our preliminary analysis to develop automated measures and metrices for CSC based on threat model each CSC is targeting, measurable features, cyber artifacts which can be used to develop metrics for key enforcement indicators. We also present a case study to develop measures and metrics for a CIS CSC called boundary defense.

14. Simulation Based Evaluation of Security and Resilience in Railway Infrastruture
Bradley Potteiger, Hamzah Abdel-Aziz, Himanshu Neema, Xenofon Koutsoukos Vanderbilt University

Autonomous and connected vehicles technologies are rapidly emerging in the modern railway domain. Traditional standalone systems are now becoming more sophisticated, resulting in more distributed interfaces and a larger attack surface. As such, there now exists a tightly coupled nature between the cyber and physical components that now makes a cyberattack not only capable of exfiltrating sensitive data, but also manipulating safety-critical operations of the system. This manipulation can cause the train to experience unsafe conditions, and can potentially lead to devastating consequences such as derailments, explosions, and collisions. These facts make it as important as ever to ensure the security, and resilience of railway infrastructure at the earliest stages of design, making deployment infrastructure more safe, reliable, and predictable in a compromising environment. In this poster we demonstrate how to analyze the security and resilience of railway infrastructure by utilizing a domain specific simulation-based framework. Our framework includes a simulation cluster with both transportation and network modeling capabilities, a graphical experiment manager to rapidly develop attack scenarios against railway networks, and a results dashboard that includes real time visualization of various metrics from within the simulation. Additionally, we not only support software-in-the-loop, but have integrated a hardware-in-the-loop testbed to test realistic attacks that occur within the deployment environment. We illustrate the capabilities of our developed framework with a connected railway case study of the Washington D.C Metro rail network.

15. Taming Web Views in the Detection of Android Privacy Leaks
Xue Qin, Robert Neuhaus, Diego Gonzales, Xiaoyin Wang, and Jianwei Niu, The University of Texas at San Antonio, and Travis Breaux, Carnegie Mellon University

Billions of smartphone users force both technical and non-technical facilities to publish their applications in market. One of the easiest ways to create an application for non-techies is by transferring their existing website using WebView. Current approaches for privacy policy analysis are mainly focused on APIs and predefined user interfaces and cannot cover these types of information. In this paper, we proposed a novel solution to trace the input data that is collected from WebView.

16. THREATZOOM: Neural Network for Automated Vulnerability Mitigation
Ehsan Aghaei and Ehab Al-Shaer University of North Carolina at Charlotte

The increasing variety and quantity of cyber threats is making it clear that traditional human-in-loop approaches are no longer sufficient to keep systems safe. To address this, forward-thinking pioneers propose a new cyber security strategy using automation to build a more efficient and cheaper defense. Associating a large number of unpatchable CVEs (vulnerability descriptions) generated every day to appropriate CWE (weakness) and CAPEC (attack pattern) can be used to automatically infer the expected impact and corresponding mitigation course of action for that new CVE. Routinely, an adversary exploits a vulnerability to trigger a cyberattack where this vulnerability results from a product or system weakness. Hence, finding a common system weakness associated with a vulnerability within a particular product can help to identify the software, system, or architecture flaw and the potential attack impacts. This identification enables us to prevent, detect, and mitigate those flaws. On the other hand, after recognizing the cause and the effect of a vulnerability, discovering the procedural-oriented description of the attack to create behavioral observables for detection and mitigation is necessary and can be derived from CAPEC and ATTCK. Mapping the CWE to CAPEC and ATTCK provides pre-TTP and post-TTP respectively where TTP stands for Tactics, Techniques, and Procedures. Having all CWE, CAPEC, and ATTCK in one hand enables us to find corresponding mitigations for each one. On the other hand, extracting threat actions provided by each of these concepts leads to finding another type of mitigation coming from Critical Security Controls (CSC). In this proposal, the target is to do mapping all the way from CVE to CAPEC and ATTCK automatically using machine learning, deep learning, and natural language processing and find the appropriate mitigation for each one and then find a proper patch as course of action defense. So far, we have introduced a neural network model which successfully classifies CVE to CWE automatically and as working on a deep learning model to classify CWEs to CAPEC

17. Toward a Reliability Measurement Framework Automated Using Deep Learning
John Heaps, Xiaoyin Wang, Jianwei Niu, The University of Texas at San Antonio, and Travis Breaux, Carnegie Mellon University

Many software bug detection measurement models and frameworks exist, however most of these approaches and tools are process-based and suffer from many limitations. We propose a framework to detect software bugs based on code pattern detection. Our framework will mine and generate bug patterns, detect those patterns in code, and calculate a reliability measure of software. Based on our knowledge, we are the first pattern-based model for the detection and measurement of bugs in software. Advancement in machine learning techniques in recent years have led researchers to consider applying deep learning to source code. While there is little research available on the subject, the work that has been done shows great potential. We believe that this work can be leveraged to obtain new insight into helping automate parts of our framework for better practicality in the real world.

18. Towards an Understanding of Resilience with Complex Networks
Francisco Pipa and John Symons University of Kansas

Resilience and robustness are key concepts to understanding the behavior of complex systems. However, they are often confused. Whereas robustness refers to the tendency of systems to remain unchanged under perturbations, resilience emphasizes the tendency of the systems to adapt, thus involving some sort of change. The formalization of both concepts using complex networks can provide common methodological tools for different sciences. This study clarifies the distinction between resilience and robustness in complex networks and provides a plausible solution to formalize resilience.

19. TRACTION: An Infrastructure for TRusted Alert Sharing and Collaborative MitigatION Keywhan Chung, Phuong Cao, Yuming Wu, Zbigniew T. Kalbarczyk, Ravishankar K. Iyer, University of Illinois Urbana-Champaign, and Alexander Withers, National Center for Supercomputing Applications

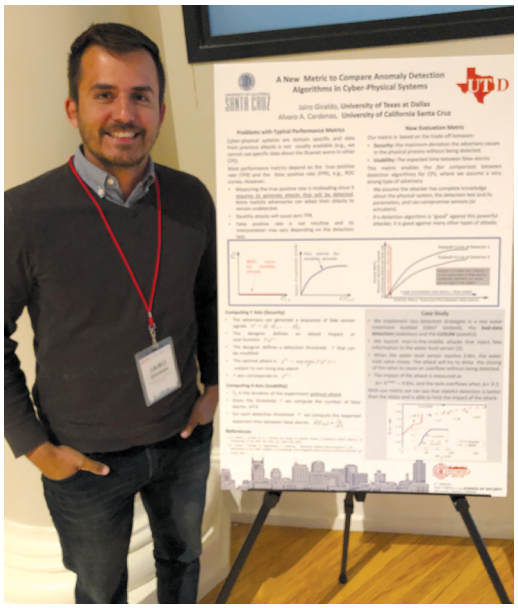
Advanced Persistent Threats (APTs) are among the most sophisticated attacks targeting networked systems. Instead of exploiting a single vulnerability, an APT uses multiple attack vectors to achieve objectives and may remain undetected for an extended period of time by staying under the radar of the defender's detection techniques. Such threats are not only growing in scale but coordinating to attack high-value sites, including both cyber and physical systems. As coordinated APTs are hard to detect with the limited data that can be collected from a single site, there is a need to enrich the observation of attacks by sharing information on monitored events with trusted sites. In this paper, we present our preliminary design of a new and unique shared infrastructure, TRACTION (Trusted Alert Sharing and Collaborative Mitigation), which at its core is a probabilistic graphical model, specifically, a distributed factor graph (DFG) anchored at each site by a local FG. The DFG provides an umbrella for automated and secure threat intelligence sharing. The overarching goal is to perform analysis and stop coordinated APTs in a manner previously not possible. Our initial design, at the scale of a single site, has been demonstrated in a production network at the National Center for Supercomputing Applications (NCSA) at the University of Illinois.

For members of the Science of Security-Virtual Organization the agenda and presentations are available at https://cps-vo.org/group/hotsos/agenda

The eighth annual HotSoS event hosted by Kansas University will be held virtually September 22-24, 2020. The symposium will again feature a mix of invited talks, a Works-in-Progress session, and refereed papers and posters to be published by ACM.

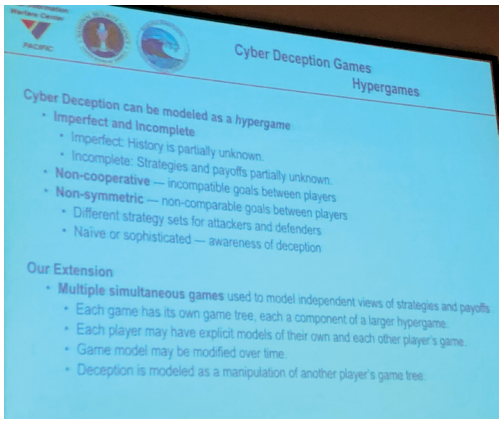


HotSoS *in Pictures*

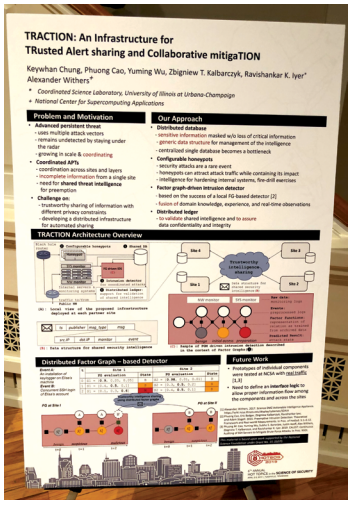


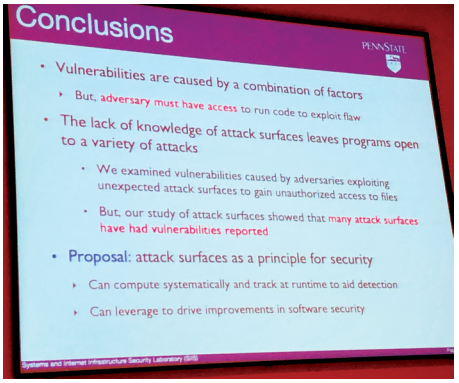








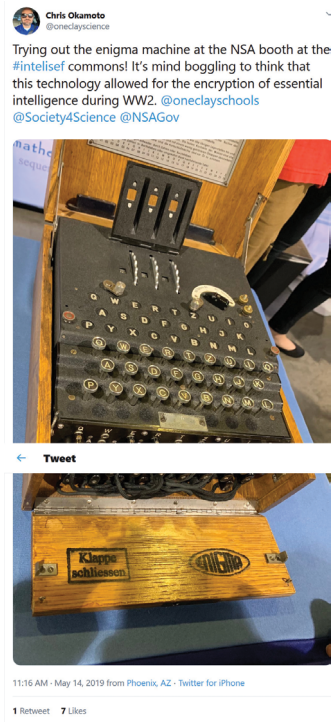




SoS *in the* News











Hicks Wins NSA's Annual Best Scientific Cybersecurity Paper Competition



Wed Oct 02, 2019

Michael Hicks, a professor of computer science, helped lead a team of researchers to victory in the National Security Agency's (NSA) 7th Annual Best Scientific Cybersecurity Paper Competition.

The research team's winning paper, "Evaluating Fuzz Testing," examines prior research on *fuzz testing*—a quality assurance technique used to discover coding errors and security loopholes in software—and evaluates whether that research produces trustworthy results. It was presented at ACM SIGSAC Conference on Computer and Communications Security (CCS '18) in Toronto.

The research was supervised by Hicks and conducted by George Klees, a sophomore majoring in computer science; Benji Cooper, a computer science alum now working as a software engineer at Audible; Andrew Ruef, who received his doctorate in computer science in 2018 and is now a research scientist at the Institute for Defense Analyses' Center for Computing Sciences; and Shiyi Wei, an assistant professor of computer science at the University of Texas at Dallas and a former postdoctoral researcher at MC2.

According to the awards website, the paper was selected because it embodies the attributes of outstanding science and the criteria of the competition: rigorous research, generalizable results and clarity of presentation.

"This paper is a step forward in bringing scientific understanding to the security community," the awards site states. "It is grounded to current understanding by its methodological survey of evaluation practices, then advances the science through quantitative analysis and proposes conclusions that apply broadly in the fuzzing community. This paper is already having tremendous impact on fuzzing research, setting the standard for how evaluations should be done."

The purpose of the competition is to recognize security papers that "best reflect the conduct of good science" in the field of cybersecurity.

Hicks, in addition to his appointment in computer science, is part of the University of Maryland Institute for Advanced Computer Studies and the Maryland Cybersecurity Center.

Eileen Buckholtz ▸ **Science of Security**

Dec 21, 2019 · · Following up on the posting about schools being attacked by ransomware, new k-12 Cybersecurity Act legislation has been introduced in congress. The act proposes stronger cybersecurity review and measures for k-12 educational institutions in the USA. #ScienceofSecurity #Cybersecurity #Schools #Ransomware https://www.securitymagazine.com/articles/91447-new-k-12-cybersecurity-act-mandates-dhs-review-schools-cybersecurity-policies



3 Seen by 34

NSA/CSS @NSAGov · Feb 7

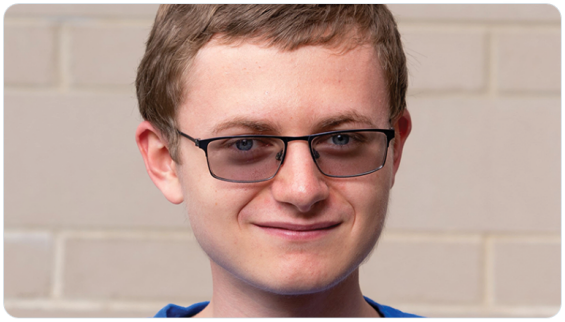
#NSA Research Director, Deborah A. Frincke delivered the opening keynote at @RiceUniversity's 1st Symposium on #DataPrivacy. The event attracted over 200 attendees from the community and raised awareness of NSA advances in this area. View the talk here: youtu.be/XoO2QS5eD14



19 35 85

ACES_umd @ACES_UMD · Oct 18, 2019

It's rare that an undergraduate researcher is listed as first author on an award winning paper. This was a reality for George, an ACES Sophomore, whose his research paper won first place for @NSAGov's Best Cybersecurity Paper Competition. Read more at aces.umd.edu/news!



5 30

ICSI Retweeted

Serge Egelman @v0max · Aug 12, 2019

Alisa Frik (@ICSIatBerkeley) presenting our work on privacy and security threats for older adults at #soups2019!

Paper here: blues.cs.berkeley.edu/blog/2019/07/0...



10 33

Gregory Yarcia Rigby

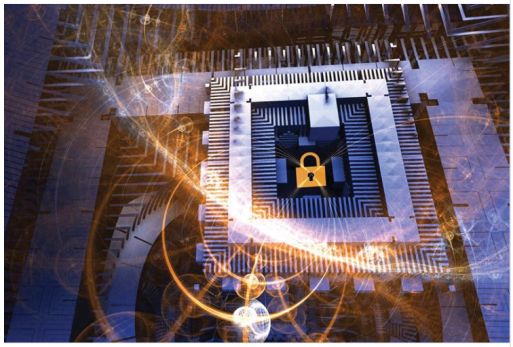
Admin · December 18, 2019

There needs to be an increase in mathematical research aimed at developing quantum-proof cryptography as the era of fully developed quantum computers approaches. Such research is needed because quantum computers are expected to be capable of breaking current encryption algorithms, RSA and ECC. Quantum-mechanical properties would allow the technology to perform calculations at a significantly faster

Linda Home Fin

https://cps-vo.org/node/64295

#ScienceOfSecurity #Research #Security #Privacy #Cryptography #QuantumComputing



Eileen Buckholtz and 3 others Seen by 43

Like Comment

Posted Jun 11, 2019 0

Adam Ekwall

Admin · January 6 at 4:40 PM

The number of automotive cybersecurity incidents has increased dramatically. Since 2016, the number of annual cyber events against automobiles has risen by 605%, with the number of incidents more than doubling in the last year alone. If you would like to read more, then please click on the link below.

#ScienceofSecurity #Cybersecurity #AutomotiveCybersecurity

https://cps-vo.org/node/64522



Linda Home Fin

Like Comment

Toni Scott Yup. I've been saying this....

How is Subaru have a car out and marketing that it monitors eye movement.... ... See More

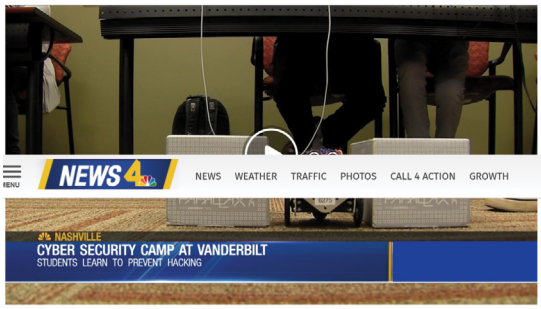
Like · Reply · 2w 2

Adam Ekwall Good point Toni, Subaru is also trying to add facial recognition. The new feature is called DriverFocus, Subaru describes it as a "driver monitoring system that uses facial recognition software to identify signs of driver fatigue or driver distraction.... See More

Like · Reply · 2w · Edited 1

Toni Scott replied · 2 Replies

Write a comment...



As featured on



Teens learn how to prevent hacking at Vandy camp

Some teenagers are learning how to keep hackers from getting your personal information.

Outreach

Researchers working at the Lablets and Sub-Lablets and collaborators from elsewhere in academia, government and industry continued to serve as Science of Security ambassadors in 2019. Whether within their own organizations or at local and international symposia and conferences, SoS champions ensured that the science of security and privacy was included in discussions and presentations.

One of the primary means of outreach remains the Science of Security Virtual Organization (SoS-VO) which was established to provide a focal point for the SoS initiative's research results, activities, and artifacts. It emphasizes community development, information sharing, and interaction among researchers in the field. SoS-VO membership grew to over 1700 members in 2019, extending the SoS presence to universities, research centers, private companies, and government agencies worldwide. The SoS-VO provides a forum to discover resources, connect to others, and share and survey cybersecurity research. The goal of the SoS-VO is to help establish and support true collaboration in advancing cybersecurity science.

The SoS-VO provides information on SoS activities, to include Lablet research, Lablet quarterly meetings, HotSoS, the annual Best Scientific Cybersecurity Paper Competition, the Intel International Science and Engineering Fair (Intel ISEF), and other SoS activities. The SoS-VO enables its members to post research findings and publications done elsewhere, advertise community events, host chats, blog, create and participate in forums. The SoS-VO provides information on upcoming events, position openings, calls for papers for conferences, and general cybersecurity news. It also serves to highlight other cybersecurity activities such as the NSA Codebreaker challenge. In 2019 there were over 5900 unique items posted on the SoS-VO.



New members are encouraged and can join by signing up via the SoS VO website at www.sos-vo.org

The SoS Reviews and Outreach (SoS R&O) newsletter published 11 editions in 2019 to over 1700 subscribers. The purpose of the R&O is to highlight research, news, and events that impact the SoS technical community. All materials included in the R&O are available on or through the SoS-VO, and are organized as follows:

- **Pub Crawl**: A summary, organized by Hard Problem, of publications that have been peer reviewed and presented at SoS-related conferences or referenced in current work. The topics are chosen for their usefulness for current researchers. There were approximately 2000 Pub Crawl items published in 2019 covering over 280 topics and representing the curated work of over 7500 authors.
- **In the News**: A consolidated list of selected articles from recent SoS-VO postings that are focused on SoS-related research, advancements, and discoveries, and are published daily on the SoS-VO. In 2019 approximately 750 news items were included in the R&O.
- **Upcoming Events**: Information on SoS-related conferences, symposia, and workshops.
- **Cyber Scene**: Material that provides an informative, timely backdrop of events, thinking, and developments that contribute to the technological advancement of SoS Cybersecurity collaboration and extend its outreach. This section explores other dimensions of cyber research beyond the academic, and also addresses US and international policy issues, proposed regulations here and abroad, congressional inquiries and testimony, and in-depth articles from non-technical publications.
- **Musings:** Brief articles on areas of concern or interest in areas of Science of Security

In addition to HotSoS, the Intel ISEF awards, and the Annual Best Scientific Cybersecurity Paper Competition, SoS participated in a number of conferences and workshops where SoS personnel shared information on the science of security and SoS initiatives. Dr. Eric Clemons, a member of the SoS team, attended the Black Engineer of the Year (BEYA) STEM conference in Washington, DC in February and discussed the science of security and privacy. Dr. Adam Tagert,

Technical Lead, SoS, presented a webinar on NSA's SoS initiative which was hosted by the Cyber Security and Information Systems information Analysis Center.

Ms. Heather Lucas, SoS team lead, presented at the NIST NICE K12 Cybersecurity Education Conference (National Institute for Standards and Technology National Initiative for Cybersecurity Education), December 9-10 in Garden Grove, CA. Ms. Lucas spoke of innovative ways to make cyber security principles visual through the use of robots in K-12 classrooms, highlighting the work done at the CPS Security Summer Camps.

Lablet summer programs geared towards K-12 students and SoS' engagement at the Intel ISEF provided opportunities to engage the next generation of researchers. The incorporation of fundamental research findings in Lablet and Sub-Lablet undergraduate coursework increases emphasis on Science of Security principles in a wide range of related disciplines. In February Dr. Adam Tagert, Technical Lead, SoS, presented a webinar on NSA's SoS initiative which was hosted by the Cyber Security and Information Systems Information Analysis Center.

CPS Summer Camp

Vanderbilt University CPS Summer Camps

Vanderbilt's Institute for Software Integrated Systems (ISIS) again offered Cyber-Physical Systems (CPS) Security Summer Camps in June 2019. Each session was a free, five-day experience using a hands-on curriculum that teaches cybersecurity, programming and robotics. Two camps were for students, and one camp, new this year, was for qualified middle and high school teachers.



Summer Camp participants put their robots to the test

Campers worked in teams using WiFi-enabled robots and NetsBlox, a visual block-based programming environment specifically designed to teach distributed programming basics and computer networking. The NetsBlox environment enables students to disrupt or hijack each other's robots, allowing them to be both attacker and defender. These camps demonstrate to students the need for various detection and defense mechanisms and spark interest in further studies in cyber-physical systems security.

Professor of Computer Engineering Akos Ledeczi and his team created NetsBlox. The team has worked with students as young as middle schoolers and the June camps are another opportunity to work with young learners. "I have an appreciation for K-12 teachers and their bright and challenging students," Ledeczi said, and smiled as the noise level rose in the classroom.

"If you consider the popularity of Netflix, Facebook, YouTube, Twitter, Google Maps, Siri, Amazon Echo, all distributed programs, distributed programming is rapidly becoming part of basic computer literacy, so NetsBlox presents a unique opportunity because students already use this technology every day and their natural curiosity will motivate them to learn more about it," Ledeczi said. "We believe that NetsBlox will provide increased motivation to students to become creators and not just consumers of technology."

In all, there were 19 students in the first camp, 25 in the second camp, and 11 teachers in the camp created for educators. While students' programming experience levels varied, they were all engrossed in the camp activities. "Several students had no previous programming experience, yet they left the camp feeling they had the skills to experiment with programming using Netsblox and they will continue improving their understanding of how to write code. The hackers to watch out for were the quiet programmers. You could see the confidence and satisfaction in their faces as they quietly hacked competing team's robots," said Mary Metelko, a system architect at ISIS and a camp resource.

Students enjoyed their week at Vanderbilt, with one noting "the programming is tiring but the outcome is fun." Another student said "I like computer science in general, but programming and robots in particular," he said. "It was super interesting. I learned a lot."

This is the third summer for Vanderbilt CPS Security camps and camp partners, the National Science Foundation and National Security Agency. NSA provided three employees as teachers. Four engineering graduate students and two undergraduates also helped with the camps.

Anyone can leverage the free netsbox and robot platforms at https://netsblox.org/.

The technology and the curriculum used at the camps were transitioned to NSA for its cyber mission outreach program, and used at NSA sponsored activities. A highlight of this adoption was the development of the "Research Robot Raceway" for NSA's Family Day. During Family Day, families of NSA employees are invited to NSA facilities to learn about its missions. The SoS team along with the Research STEM advocate developed an interactive raceway. Visitors were provided a hands-on experience with the robots while being introduced to cybersecurity best practices.

Racers selected a robot from a set of five profiles. Each robot, named after a different cyber-pioneer, had a distinct set of cyber defenses that could affect the robot's performance. For example, a strong password or use of encryption would result in a decrease in speed. Using Xbox controllers, racers were to complete a lap around the track while traversing an attack zone. In this zone, each robot would be subjected to a randomly-selected cyber attacks. If the robot was vulnerable, the attack

would cause the robot to misbehave. During the four-hour event, roughly a hundred races were run.

The five robots were:

- **Art** – Art Levinson, a WWII and NSA cryptographer. Features: Pure Speed. Art has no time for security updates, he just wants to win and go fast.
- **Grace** – Admiral Grace Hopper, computer science pioneer. Features: Latest gadgets installed. Grace doesn't use a password but has mitigations to respond to denial of service attacks.
- **Whit** – Whitfield Diffie, cryptographer. Features: Fast with secure software. Whit has the latest tech but often forgets his password so he writes it down and posts it near his computer.
- **Ada** - Ada Lovelace, published the first machine algorithm. Features: Vintage racer. Ada uses old technology but has a difficult password. *Ada had the most wins of any other robot.*
- **Hedy** - Hedy Lamarr, inventor of frequency hopping spread spectrum technology. Features: Up-to-date security hardware and software. Hedy uses NSA's open source SPECK crypto cipher to encode her messages with a difficult password and sequence numbers. *While Hedy was the slowest robot, none of the attacks in the race compromised Hedy.*

The four attacks:

- Message injection: robots were sent commands that did not originate from the racers causing the robot to spin.
- Denial of Service: a flood of messages were sent which prevented the robot from responding to commands quickly.
- Replay Attack: previously sent messages were repeated.
- Password Cracking: spoofed messages were sent by determining the password by brute forcing the key. The command put the robot into a high-speed reverse.



Students learning how to program the robots



Teachers became students in their summer camp and learned about the Internet of Things

Lessons learned from NSA Family Day were incorporated in follow-on outreach events such as STEAM Night at the Museum, a public event at the National Cryptologic Museum. In addition to the NSA Family Day, the NSA Cryptologic Centers employed the technology for their family outreach events. The technology continues to be developed for NSA outreach programs.







produced by cyber pack ventures, inc.



sos-vo.org