



Science of Security
and Privacy



Annual Report

April 2021

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 2021



Table of Contents

Executive Summary	1
Section 1: Engaging the Academic Community for Foundational Research	3
Hard Problems	4
Science of Security Lablets	
Carnegie Mellon University	7
International Computer Science Institute	17
North Carolina State University	24
University of Illinois Urbana-Champaign	35
University of Kansas	46
Vanderbilt University	58
Science of Security Quarterly Meetings	67
SoS Virtual Seminar Series	69
Section 2: Promoting Rigorous Scientific Principles	70
Annual Best Cybersecurity Paper Competition	71
Section 3: Growing the Science of Security	74
HotSoS 2019	75
Outreach	82

Executive Summary

Despite the pandemic, the Science of Security and Privacy initiative (SoS) 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 impactful foundational research
- Promote rigorous scientific principles
- Grow the SoS community

Under the sponsorship of the National Security Agency (NSA) Research Directorate (RD) 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 is advancing the goal of safeguarding interactions in cyberspace.

Despite the pandemic, 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 in 2020.

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 2011 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 2020, 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 initiating a SoS Virtual Seminar Series to increase NSA exposure to Lablet projects. Lablet 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 CPS 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 2020 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 8th Annual Best Scientific Cybersecurity Paper Competition and an additional Best Paper Award at the Hot Topics in the Science of Security Symposium (HotSoS 2020). There were 52 papers submitted in the 8th Annual Paper Competition bringing the total submissions to over 300 during the competition's seven years. This year's winning paper, "Spectre Attacks: Exploiting Speculative Execution" was originally published at the 2019 IEEE Security & Privacy Symposium. The winning paper, along with "Meltdown: Reading Kernel Memory from User Space" a 7th Annual Honorable Mention winner by the same researchers, launched a global effort to mitigate critical vulnerabilities in processors. While SoS had sponsored prizes at the Intel International Science and Engineering Fair (ISEF) for five years, ISEF had to be cancelled because of the pandemic.

Details on SoS activities to promote rigorous scientific principles in 2020 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. At the outset of the pandemic when lockdowns were instituted and in-person activities were curtailed, additional effort was taken to ensure that the SoS community remained connected via postings on the SOS-VO and Facebook. The SoS-VO (www.sos-vo.org) grew to over 1800 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. Although delayed by the COVID pandemic and held virtually, HotSoS 2020 nonetheless brought together over 400 researchers and practitioners from academia, government, and industry for thought-provoking presentations on the cybersecurity science. This was the largest attendance in HotSoS history.

Details on what the SoS initiative did in 2020 to grow the Science of Security community can be found in Section 3.

In 2021, 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 through the continuation of the SoS Virtual Seminar Series. The Best Scientific Cybersecurity Paper Competition, the 9th Annual, will again recognize papers that exemplify the development of scientific rigor in cybersecurity research. HotSoS 2021 will, for the first time, be hosted by NSA and held virtually from April 13-15, 2021. HotSoS 2021 will encourage interaction among presenters and attendees and focus on Works-in-Progress (WiPs), already published work, posters, and student presentations. 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 cyberdefense.

Section 1

Engaging ^{the} Academic Community ^{for} Foundational Research



In 2020 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 at Urbana-Champaign (UIUC), the University of Kansas (KU), 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.

Despite the pandemic which forced the Lablets to either temporarily curtail or cease operations, the Lablets and their Sub-Lablets performed substantive research on 23 projects in 2020 and published 67 peer-reviewed articles or papers bringing to well over 700 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. This year, most of the presentations were virtual. The foundational research embodied by the papers has contributed significantly to enhancing the scientific rigor of research into cybersecurity.


In 2019 the SoS team recognized nine Associate Lablets (Alabama A&M University, Columbia University, Drexel University, Harvard University, Jackson State University, Morgan State University, Naval Postgraduate School, North Carolina

A&T University, and Simmons University) to expand the academic partnership network and tap into existing national pipelines of diverse cybersecurity talent. As was the case with the six Lablets, research activities were limited by the pandemic and interactions were virtual.

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, establish a common language and a way to assess 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.

Lablet researchers were limited in their community outreach and education initiatives this year, but were able to participate virtually in international conferences and workshops. The Lablets continued to provide quarterly reports on their activities, and were able to meet in person for the 2020 Winter Quarterly prior to the start of COVID-19 restrictions. The SoS initiative also launched the Virtual Seminar Series to allow NSA personnel, including those teleworking, to hear about Lablet projects. The Series also enabled Lablet student researchers to dry-run their presentations before a conference or thesis defense presentation, for example.

The Hard Problems, Lablet projects and activities, the Lablet Winter Quarterly meeting, and the Virtual Seminar Series are described 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 \(November 2012\)](http://cps-vo.org/node/6394) 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 \(August 2015\)](http://cps-vo.org/node/21590) 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. 9* Characterizing user behavior and anticipating its effects on computer security with a Security Behavior Observatory
- p. 11* Model-Based Explanation For Human-in-the-Loop Security
- p. 13* Obsidian Language for Blockchain
- p. 15* Securing Safety-Critical Machine Learning Algorithms

Lead:

Perry Alexander

Sub-Lablet:

University of Tennessee

Hard Problems:



Metrics



Scalability and Composability



Resilient Architectures

Projects:

- p. 48* Cloud-Assisted IoT Systems Privacy
- p. 50* Formal Approaches to the Ontology & Epistemology of Resilience
- p. 52* Scalable Trust Semantics & Infrastructure
- p. 54* Secure Native Binary Execution
- p. 56* Side-Channel Attack Resistance

Lead:

Serge Egleman

Sub-Lablets:

Cornell Tech
UC Berkeley

Hard Problems:



Human Behavior



Metrics



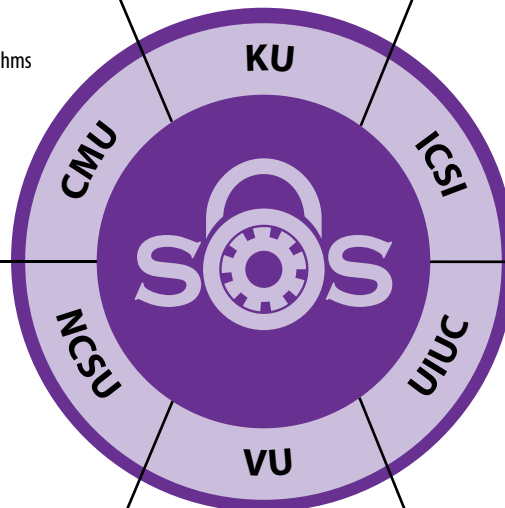
Scalability and Composability



Secure Collaboration

Projects:

- p. 18* Contextual Integrity for Computer Systems
- p. 19* Designing for Privacy
- p. 20* Governance for Big Data
- p. 21* Operationalizing Contextual Integrity
- p. 22* 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. 25* Coordinated Machine Learning-Based Vulnerability & Security Patching for Resilient Virtual Computing Infrastructure
- p. 27* Development of Methodology Guidelines for Security Research
- p. 29* Predicting the Difficulty of Compromise through How Attackers Discover Vulnerabilities
- p. 31* Principles of Secure BootStrapping for IoT
- p. 33* 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. 59* Analytics for Cyber-Physical System Cybersecurity
- p. 61* Foundations of a CPS Resilience
- p. 63* Mixed Initiative and Collaborative Learning in Adversarial Environments
- p. 65* Multi-model Test Bed for the Simulation-based Evaluation of Resilience

Leads:

Sayan Mitra
David M. Nicol

Sub-Lablets:

Illinois Institute of Technology
UT Austin

Hard Problems:



Human Behavior



Metrics



Resilient Architectures



Scalability and Composability




Secure Collaboration

Projects:

- p. 36* A Human-Agent-Focused Approach to Security Modeling
- p. 38* Automated Synthesis Framework For Network Security and Resilience
- p. 40* Monitoring, Fusion, and Response for Cyber Resilience
- p. 42* Resilient Control of Cyber-Physical Systems with Distributed Learning
- p. 44* Uncertainty in Security Analysis

Carnegie Mellon University




The Carnegie Mellon University (CMU) Science of Security Lablet has made advances in composability and usability within the science of security. The Lablet is focusing on four technical thrusts: (1) contrastive explanation in automated planning to assist designers in understanding how computer agents make decisions in multi-objective planning; (2) evaluating leakage of training data from statistically-learned models when explainability is used to interpret model decisions; (3) designing a typestate language for blockchain programs that combines strong technical guarantees with an explicit focus on programmer usability; and (4) understanding how computer users practice “security hygiene” on their personal computers and mobile devices.

Principal Investigator (PI) Travis Breaux and CO-PI Jonathan Aldrich lead a team of faculty, postdoctoral and PhD student researchers from CMU and four Sub-Lablets, George Mason University, University of North Carolina at Chapel Hill, University of California, Berkeley, and Indiana University. Many of the CMU Lablet faculty and graduate students are affiliated with the CMU CyLab, the coordinating entity for cybersecurity research.



Travis Breaux



Jonathan Aldrich

In 2020 the National Security Agency’s Research STEM outreach funded five Security Science Research Experience for Undergraduates (REUs) to work with Lablet researchers over the summer. Students were trained in security science and introduced to academic research practices. Most of these students came to CMU from universities that did not have active cybersecurity research programs, which effectively broadened participation and diversity in research. The Security Science students plus students from CMU’s sister program, the Research Experience for Undergraduate in Software Engineering (REUSE), met with Adam Tagert, NSA’s SoS Technical Lead, to discuss research and careers at NSA.

CyLab
6,043 Tweets Follow

options, making it easier for working professionals to receive #privacy engineering training and education while continuing to work. Read all about it: cylab.cmu.edu/news/2020/11/2...

Carnegie Mellon ECE Retweeted

CyLab @CyLab · May 18, 2020
Having a paper accepted to @IEEEESSP is a big deal. Having a paper receive the symposium's Test-of-Time award is a *huge* deal. Five @CarnegieMellon

Carnegie Mellon ECE
2,705 Tweets Follow

Five CMU security and privacy papers awarded IEEE's Test of Time award
Five CMU security and privacy papers awarded IEEE's Test of Time award -
CyLab Security and Privacy Institute
cylab.cmu.edu

The five participating students and their research projects were:

- Emma Hogan, Siena College, “picoCTF Cybersecurity & Education”

picoCTF is a free online Capture the Flag-style game implemented by and maintained at CMU to promote and facilitate cybersecurity education. To improve the design of the picoCTF platform in order to reach a larger number of students, especially in under-resourced communities, new empirical data and analyses are needed to inform new design enhancements and platform scalability. This project included a research plan to obtain the needed data through user studies, focus groups, and usability and scalability tests that examine picoCTF in a classroom setting.

- Alexander Billups, Pennsylvania State University at Johnstown, “Security Requirements”

This project focused on designing and implementing crowdsourcing and Natural Language Processing (NLP)-based tools to advance requirements extraction from text. Example NLP techniques included phrase structure grammars and typed dependencies, and feature engineering and deep learning to perform named entity recognition and semantic role labeling.

- Sang Heon Choi, Rose–Hulman Institute of Technology, “Safe and Robust Human-Machine Interfaces”

Unintuitive, badly designed Human-Machine Interfaces (HMI) are not merely an annoyance, but can pose significant risks to users in systems where safety is a key concern. The challenge is that humans are far from perfect and inadvertently make mistakes from time to time, but many interfaces are not designed to deal with such human errors; as a result, when a safety failure occurs, it is often the users who get blamed, even when a better designed interface could have prevented such a failure. To tackle this challenge, the project was a development of an interface that is explicitly designed to recognize and handle potential human errors and prevent them from causing safety failures.

- Timothy Mou, Swarthmore College, “Nominal Wyvern: Adapting Dependent Object Types for Decidable Subtyping”

Dependent Object Types (DOT) in languages such as Scala bring together aspects of both functional and object-oriented languages; however, subtyping in DOT is undecidable. Nominal Wyvern proposes a DOT-like system that recovers decidability by enforcing a “material-shape” semantic separation, distinguishing between concrete types and types used for constraints. Timothy and his team developed an implementation of Nominal Wyvern in a selfhosting compiler.

- Reed Oei, University of Illinois at Urbana Champagne, “Psamathe: A DSL for Safe Blockchain Assets”

The project was to develop a new abstraction, called Flows, representing an atomic transfer operation. Flows allow encoding semantic information about the flow of assets into the code. Flows leverages the Labet-developed Obsidian language’s existing capability to mark types with modifiers, such as asset, making some classes of bugs impossible. In addition, Flows allow programmers to define contracts more concisely. A paper describing this work was published in the ACM SPLASH 2020 Student Research Competition.



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 home computer users' computer use and online behavior choices that impact security and privacy. This work can be used to develop models and technologies to be targeted to realistic situations.

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. This team is conducting ongoing work on a number of research questions, including investigating how people browse the web, and a comparison between desktop/laptop use and mobile device use. This is leading to a model to detect that website that are likely to expose users to malware before traditional detection methods have identified it.

ACCOMPLISHMENTS

In December 2019, we conducted an analysis of the optional exit survey that SBO participants completed over the summer at the end of data collection. This offers useful insights into what participants liked and disliked about the experience of participating in this unusual longitudinal study, and some participants offered suggestions for improving that experience for future studies of this type. Since this type of study is rare, we may combine this with researchers' anecdotal experiences in running this study and submit it to a magazine or workshop to share what we have learned from conducting this study.

The original SBO collected data from home Windows computers (desktop or laptop). This year, some of the SBO team worked with researchers in Japan to develop a Mobile SBO that will

allow for similar data collection from mobile devices, allowing for ecologically valid data collection from these devices. Mobile usage has eclipsed desktop usage worldwide since the creation of the SBO, so the ability to instrument mobile devices will be essential in the future to obtain a representative snapshot of computing behavior and continue working towards a better understanding of users' everyday security and privacy challenges. A poster on this work (Mobile Security Behavior Observatory: Long-term Monitoring of Mobile User Behavior. Akira Yamada, Shoma Tanaka, Yukiko Sawaya, Ayumu Kubota, So Matsuda, Reo Matsumura, Shun Umamoto, Jun Nakajima, Kyle Crichton, Jin-Dong Dong, and Nicolas Christin) was presented remotely at SOUPS 2020 in August.

One of our ongoing projects has involved studying the extent to which security- and privacy-related information is presented to users through their social media or “friends”. By analyzing the actual social media logs of participants, we study the extent to which this information is exposed to users and how different aspects of this information impacts people’s security behaviors measured across the type of browsing they do, password use habits, and other system-related behaviors. A paper describing the results of this research is currently under review.

We are also reviewing findings from an ongoing project studying

home users’ web browsing patterns. As a result of analysis from 257 participants we report a substantial increase in tabbed browsing that highlights the need to include tab information in order to obtain accurate web measurements. We observe that web browsing is highly centralized, with about half of internet use spent on 1% of websites, but we also note that users spend a large portion of time on websites with overall low visit counts, where riskier content is likely to be found. We discuss how users may be getting to these sites and note implications for future security research. A paper on this research is under review.

IMPACT ON HARD PROBLEM

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.

PUBLICATIONS

- Sruti Bhagavatula, Lujo Bauer, and Apu Kapadia, “(How) Do people change their passwords after a breach?” Appeared at *Workshop on Technology and Consumer Protection (ConPro 2020)*, Virtual Conference, May 21, 2020.

Data breach victims aren't changing their passwords

By Anthony Spadafora June 03, 2020

Yet another reminder to always change your password after a company reports a data breach





(Image credit: Shutterstock)

A new study by academics from Carnegie Mellon University's Security and Privacy Institute (CyLab) has revealed that only a third of users actually change their passwords after a [data breach](#) announcement.

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

Human end-users of planning software agents require trust that the software agents make decisions in ways that conform to what the human users want. In many real-world applications of planning, multiple optimization objectives, including security, are often involved. Thus, planning agents' decisions can involve complex tradeoffs among competing objectives. It can be difficult for a human to understand why an agent decides on a particular planning solution on the basis of its objective values, particularly when there is a potential misalignment between the agent and the user's preference for the different planning objectives. As a result, the user may not know whether the agent's decision is the best option with respect to their own values and preferences. In this work, we contribute an explainable planning approach, based on contrastive explanation, that enables the agent to communicate its preference for the different planning

objectives and help the user better understand whether the agent's decision is optimal with respect to their own preference, despite a potential value misalignment. We conducted a human-subject experiment to evaluate the effectiveness of our explanation approach in the mobile robot navigation domain. The results show that our approach significantly improves the users' ability and reliable confidence in determining whether the agent's decisions are in line with their preferences.

Explanation can be helpful to allow the human to understand why a system is making certain decisions, as demonstrated in our human-subject experiments. However, explanations come with costs in terms of, e.g., delayed actions, or the possibility that a human may make a bad judgement. Hence, it is not always obvious whether explanations will improve the satisfaction of

system goals and, if so, when to provide them to a human. We defined a formal framework for reasoning about explanations of adaptive system behaviors and the conditions under which they are warranted. Specifically, we characterized explanations in terms of their impact on a human operator’s ability to engage in adaptive actions. We leverage a probabilistic reasoning tool to determine when an explanation should be used as a tactic in an adaptation strategy in order to improve overall system utility. The approach is illustrated in a representative scenario of an adaptive news website in the context of potential denial-of-service attacks, and can be used to decide when an explanation should be provided, based on knowledge about a human operator’s capability and the cost associated with generating an explanation.

A key element of our approach to automating decisions about both *when* and *how* a system should offer explanations to humans in order to improve overall human-system collaboration has been the use of formal planning models that capture the multi-dimensional nature of the domain, tradeoffs among various qualities, and explicit modeling of uncertainty. Such models have advantages over other AI approaches to decision making where the rationale for autonomous actions is less explicit (such as with machine learning), and they form the basis for a rich theory of human-system interaction that can be leveraged to improve system performance, especially in critical decision-making contexts such as in support of machine-assisted security.

IMPACT ON HARD PROBLEM

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 at appropriate times for automated mitigation plans generated by self-protecting systems that use formal models of the software, network, attack, and collaborating humans.

PUBLICATIONS


- Roykrong Sukkerd, Reid Simmons and David Garlan, “Tradeoff-Focused Contrastive Explanation for MDP Planning,” In Proceedings of the *29th IEEE International Conference on Robot & Human Interactive Communication*, Virtual, 31 August - 4 September 2020.
- Nianyu Li, Javier Camara, David Garlan and Bradley Schmerl, “Reasoning about When to Provide Explanation for Human-in-the-loop Self-Adaptive Systems” In Proceedings of the *2020 IEEE Conference on Autonomic Computing and Self-organizing Systems (ACSOS)*, Washington, D.C., 19-23 August 2020.
- Nianyu Li, Sridhar Adepu, Eunsuk Kang, and David Garlan, “Explanations for Human-in-the-loop: A Probabilistic Model Checking Approach,” In Proceedings of the *15th International Symposium on Software Engineering for Adaptive and Self-managing Systems (SEAMS 2020)*.

The screenshot shows the SEAMS 2020 website interface. At the top, it says 'VIRTUAL SEAMS 2020' and '15th INTERNATIONAL SYMPOSIUM ON SOFTWARE ENGINEERING FOR ADAPTIVE AND SELF-MANAGING SYSTEMS'. The dates 'JUN 29 - JUL 3' are also visible. Below the header, there are navigation menus for 'Attending', 'Info', 'Program', 'Track/Call', 'Organization', 'Search', and 'Series'. A breadcrumb trail indicates the current page is 'ICSE 2020 (series) / SEAMS 2020 (series) / 15th International Symposium on Software Engineering for Adaptive and Self-Managing Systems /'. The main title of the paper is 'Explanation for Human-on-the-loop: a probabilistic model checking approach'. The authors listed are 'NIANYU LI, Sridhar Adepu, Eunsuk Kang, David Garlan'. The track is 'SEAMS 2020'. The abstract states: 'Many self-adaptive systems benefit from human involvement and oversight, where a human operator can provide expertise not available to the system and can detect problems that the system is unaware of. One way of achieving this is by placing the human operator on the loop – i.e., providing supervisory oversight and intervening in the case of questionable adaptation decisions. To make such interaction effective, explanation is sometimes helpful to allow the human to understand why the system is making certain decisions and calibrate confidence from the human perspective. However, explanations come with costs in terms of delayed actions and the possibility that a human may make a bad judgement. Hence, it is not always obvious whether explanations will improve overall utility and, if so, what kinds of explanation to provide to the operator. In this work, we define a formal framework for reasoning about explanations of adaptive system behaviors and the conditions under which they are warranted. Specifically, we characterize explanations in terms of explanation content, effect, and cost. We then present a dynamic adaptation approach that leverages a probabilistic reasoning technique to determine when the explanation should be used in order to improve overall system utility.'

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







Jonathan Aldrich



Brad Myers

PI:	Jonathan Aldrich	CO-PI	Brad Myers
SUB-LABELT	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

We performed an empirical study with 20 participants comparing Obsidian to Solidity, which is the language most commonly used for writing smart contracts today. We observed that most of the Obsidian participants were able to successfully complete most of the programming tasks we gave them. We also found that asset-related bugs, which Obsidian detects at compile time, were commonly accidentally inserted by the Solidity participants. We identified potential opportunities to improve the usability of typestate as well as to apply the

usability benefits of Obsidian's ownership system to other languages.

Most distributed payment schemes (cryptocurrencies) suffer from lack of privacy and anonymity for their users due to the public nature of their ledger. We developed MiniLedger, a permissioned, distributed payment system which not only guarantees the privacy of transactions but also offers built-

in functionalities for various types of audits by any external authority (i.e., audits on transaction values or of total assets of participants). As a starting point we use the recently proposed zkLedger architecture [DBLP:conf/nsdi/NarulaVV18]. We define system functionalities and security properties and we address vulnerabilities and shortcomings identified in zkLedger. Most importantly, MiniLedger is the first private and accountable payment system with nearly constant storage costs. To achieve such a storage improvement, we implement pruning functionalities for the transaction history without hurting security or auditing, while we extend MiniLedger to perform fine-grained audits in a client level. Our evaluation results show that MiniLedger is not only practical in terms of storage, but also has minimal computational overhead compared to zkLedger, as the pruning-related functionalities can be tailored for maximum efficiency on each deployment use-case.

We developed a new approach for detecting injection vulnerabilities in applications by harnessing the combined power of both human developers' test suites and automated dynamic analysis. Our new approach, RIVULET, monitors the execution of developer-written functional tests in order to detect information flows that may be vulnerable to attack. Then, RIVULET uses a white-box test generation technique to repurpose those functional tests to check if any vulnerable flow could be exploited. When applied to the version of Apache Struts exploited in the 2017 Equifax attack, RIVULET quickly identifies the vulnerability, leveraging only the tests that existed

in Struts at that time. We compared RIVULET to the state-of-the-art static vulnerability detector Julia on benchmarks, finding that RIVULET outperformed Julia in both false positives and false negatives. We also use RIVULET to detect previously unknown vulnerabilities.

We adapted HCI methods to make them more suitable for programming language design and integrated these methods into a new process, PLIERS, for designing programming languages in a user-centered way. We evaluated PLIERS by using it to design two new programming languages. Glacier extends Java to enable programmers to express immutability properties effectively and easily. Obsidian is a language for blockchains that includes verification of critical safety properties. Summative usability studies showed that programmers were able to program effectively in both languages after short training periods.

We describe two case studies that evaluate Obsidian's applicability to the domains of parametric insurance and supply chain management, finding that Obsidian's type system facilitates reasoning about high-level states and ownership of resources. We compared our Obsidian implementation to a Solidity implementation, observing that the Solidity implementation requires much boilerplate checking and tracking of state, whereas Obsidian does this work statically.

We worked with one of the students in the Security Science Research Experience for Undergraduates program to develop a new abstraction, Flows, which is described in the CMU introduction.

IMPACT ON HARD PROBLEM

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.

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

Prior to the onset of COVID-19, we engaged the blockchain development and research community with Obsidian as follows:

- We gave invited talks about Obsidian to three companies:

Ripple, RogueWave, and Perforce.

- We gave invited talks at Harvard, MIT, UC Berkeley, University of New Orleans, and Dagstuhl.

PUBLICATIONS

- Michael J.Coblentz, "User-Centered Design of Principled Programming Languages," Ph.D. Dissertation, 2020, Carnegie Mellon University, Pittsburgh, PA. <http://reports-archive.adm.cs.cmu.edu/anon/2020/CMU-CS-20-127.pdf>.

- Reed Oei, "Psmathe: A DSL for Safe Blockchain Assets," In Proceedings of the ACM SPLASH 2020 Student Research Competition.



Securing Safety-Critical Machine Learning Algorithms




Lujo Bauer



Matt Fredrikson

PI:	Lujo Bauer	CO-PI Matt Fredrikson
PARTICIPATING SUB-LABELT	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 continued our study of network pruning techniques to enhance robustness. Our approach is based on attribution measurements of internal neurons, and aims to identify features that are pivotal for adversarial examples but not necessary for correct classification of normal inputs. Our experiments to date suggest that it is possible to identify and remove such non-robust features for norm-bounded attacks, but suggest that physical attacks may rely on different sets of features that cannot be pruned without significant impact on model performance.

Our team's research focused on revising and extending previous results on n-ML (which provides robustness to evasion attacks via ensembles of topologically diversified classifiers) and attacks on malware detection. We've discovered that the utility of and best approaches for tuning n-ML differ depending on the dataset and its complexity, e.g., tunings of n-ML that lead to particularly good performance for MNIST lead to sub-ideal performance for GTSRB and vice versa. In the process, we're discovering tunings of n-ML that further improve its performance compared to other approaches for making classifiers more robust. We have improved our algorithm for attacking malware classifiers

to better gauge the impact of small changes to the binary (e.g., swapping a pair of instructions) on the correctness of classification of the binary as benign or malware. We have also identified engineering errors in libraries that our code builds on; fixing these should result in significantly lower resource usage, enabling more comprehensive experiments.

In addition to continuing our research on n-ML and on evasion attacks against malware classifiers, we are also working on improving experimental infrastructure and methodology, which will enable more automated, much quicker experiments with malware evasion and a more comprehensive examination of the effects of hyperparameter tuning on both attacks on defenses.

We are also continuing our research on investigating the leakage of training data from models. We are currently examining the increased risk that techniques from explainability pose to this leakage, as well as the role that robustness plays in this risk. We aim to determine the feasibility of leakage attacks in black-box settings, where explainability methods are most likely to be used.

IMPACT ON HARD PROBLEM

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

PI Lujo Bauer presented work that was part of this project as follows:


- At a seminar in CMU-Africa in Kigali, Rwanda
- As part of a keynote for the German American Chambers of Commerce East Coast Industry Forum's meeting in Pittsburgh
- At a seminar at ETH Zurich
- At the CPS Verification & Validation Workshop at CMU

PUBLICATIONS

- Christopher Bender, Yang Li, Yifeng Shi, Michael Reiter, and Junier Oliva, "Defense through Diverse Directions," In Proceedings of the 37th International Conference on Machine Learning, July 2020.

Machine learning (ML) algorithms are becoming ubiquitous; they're used in applications from playing chess and predicting the weather to cancer diagnosis and self-driving cars. In this project we first try to understand how robust ML algorithms are in the face of an adversary. Specifically, we study whether an adversary can fool ML classifiers in *practical settings* without arousing the suspicion of a human. For instance, we showed that it is possible to 3d print a pair of eyeglasses that, when worn by an adversary, can cause a state-of-the-art face-recognition algorithm to identify the adversary as (a specific) someone else. We leverage what we learn of ML algorithms' weaknesses to design ML algorithms that are more resistant to attack.

Video demonstrating targeted impersonation: Mahmood impersonates Ariel against VGG10. The video shows that the face recognizer isn't confused by non-adversarial eyeglasses, including large, bright ones, but adversarial eyeglasses generated specifically to fool the recognizer into classifying Mahmood as Ariel are overwhelmingly successful at doing so. Targeted impersonation is achieved via the method described in "Adversarial generative nets: neural network attacks on state-of-the-art face recognition" (see below).





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.



Serge Egelman

On March 17, 2020, the city of Berkeley issued a public health order instructing all residents to remain at home and mandating ICSI's closure. Consequently, ICSI was forced to stop work. Some research on selected projects was accomplished prior to the shutdown. That research is described in the relevant project writeup.

The five research projects are identified below:

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

ICSI Retweeted

Irwin Reyes @irwinreyescom · Jul 16, 2020

What a fun project that was, and it was exciting to see its impact in the security and privacy community *and* beyond. Really thrilled to share this award with my coauthors @primalw, Joel, @r_pannah, @AmitElazari, @narseo, and @v0max

PETS @PET_Symposium · Jul 16, 2020

"Won't Somebody Think of the Children?" Examining COPPA Compliance at Scale by @@irwinreyescom, @primalw, Joel Reardon, @AmitElazari, @r_pannah, @narseo, and @v0max (PoPETs 2018.3) #pets20

Show this thread

ICSI Retweeted

AppCensus @AppCensusInc · Jul 16, 2020

Congrats to Serge Egelman (@v0max), Narseo Vallina-Rodriguez (@narseo), Joel Reardon and co-authors for the Caspar Bowden Award for Outstanding Research in Privacy Enhancing Technologies announced today!

Link to paper: petsymposium.org/2018/files/pap...

#privacy #pets20 #coppa

5 retweets, 15 likes

ICSI @ICSIatBerkeley · Jul 16, 2020

Congratulations to all of you!

ICSI Retweeted

Center for Long-Term Cybersecurity @CLTCBerke... · Feb 21, 2020

Privacy policies + permission requests don't disclose all the ways that consumers are being tracked: "We need watchdog groups and regulators. They should be doing this so that consumers don't need to." says @ICSIatBerkeley director @v0max

ICSI 1,436 Tweets Following

Your phone talks about you behind your back. These researchers are list... They're keeping track by tracking the trackers that track you.

cnet.com

Contextual Integrity for Computer Systems





Michael Tschantz



Helen Nissenbaum

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



No substantive research was performed during the past year.

GOAL

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.

ABSTRACT

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.

IMPACT ON HARD PROBLEM

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.

Designing for Privacy



Deirdre Mulligan

PI:	Deirdre Mulligan
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 the 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

Researchers started performing narrative review of the factors/barriers to Privacy-by-Design (PbD) and other privacy frameworks, which will inform a meta-analysis.

IMPACT ON HARD PROBLEM

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.

Governance for Big Data



Deirdre Mulligan

PI:	Deirdre Mulligan
SUB-LABELT	University of California, 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

Researchers made progress in planning content for workshops, as well as identifying potential participants. They also designed a study to interview “privacy champions” within

large organizations to better understand their data governance practices.

IMPACT ON HARD PROBLEM

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

- Met with FTC officials to discuss data governance issues.

Operationalizing Contextual Integrity

PI:	Serge Egelman	CO-PI: Helen Nissenbaum
SUB-LABEL	Cornell Tech	
URL:	https://cps-vo.org/ICSI_OperationalizingContextualIntegrity	

HARD PROBLEM: Scalability and Composability, Policy-Governed Secure Collaboration, Security Metrics and Models, Human Behavior



Serge Egelman



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 to determine what parameters are actually important to users are: 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

Paper accepted on examining privacy expectations vs. reality surrounding disaster preparedness/response apps.

Designed multiple studies all focusing on users' expectations surrounding in-home data collection and the types of controls they would like to have.

IMPACT ON HARD PROBLEM

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

- Attended Privacy Papers for Policymakers award reception

to discuss this research (which received the Best Student Paper Award)

PUBLICATIONS

- Madelyn R. Sanfilippo, Yan Shvartzshnaider, Irwin Reyes, Helen Nissenbaum, and Serge Egelman, "Disaster

Privacy/Privacy Disaster," Journal of the Association for Information Science and Technology (JASIST), 2020; 1- 13.

Scalable Privacy Analysis

PI: Serge Egelman CO-PI: Narseo Vallina-Rodriguez

URL: https://cps-vo.org/ICSI_ScalablePrivacyAnalysis

HARD PROBLEM: Scalability and Composability, Policy-Governed Secure Collaboration, Security Metrics and Models



Serge Egelman



Narseo Vallina-Rodriguez

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

We began building a taxonomy of the ways in which apps attempt to detect whether or not they are being monitored

(specifically, whether they're running on jailbroken/rooted devices).

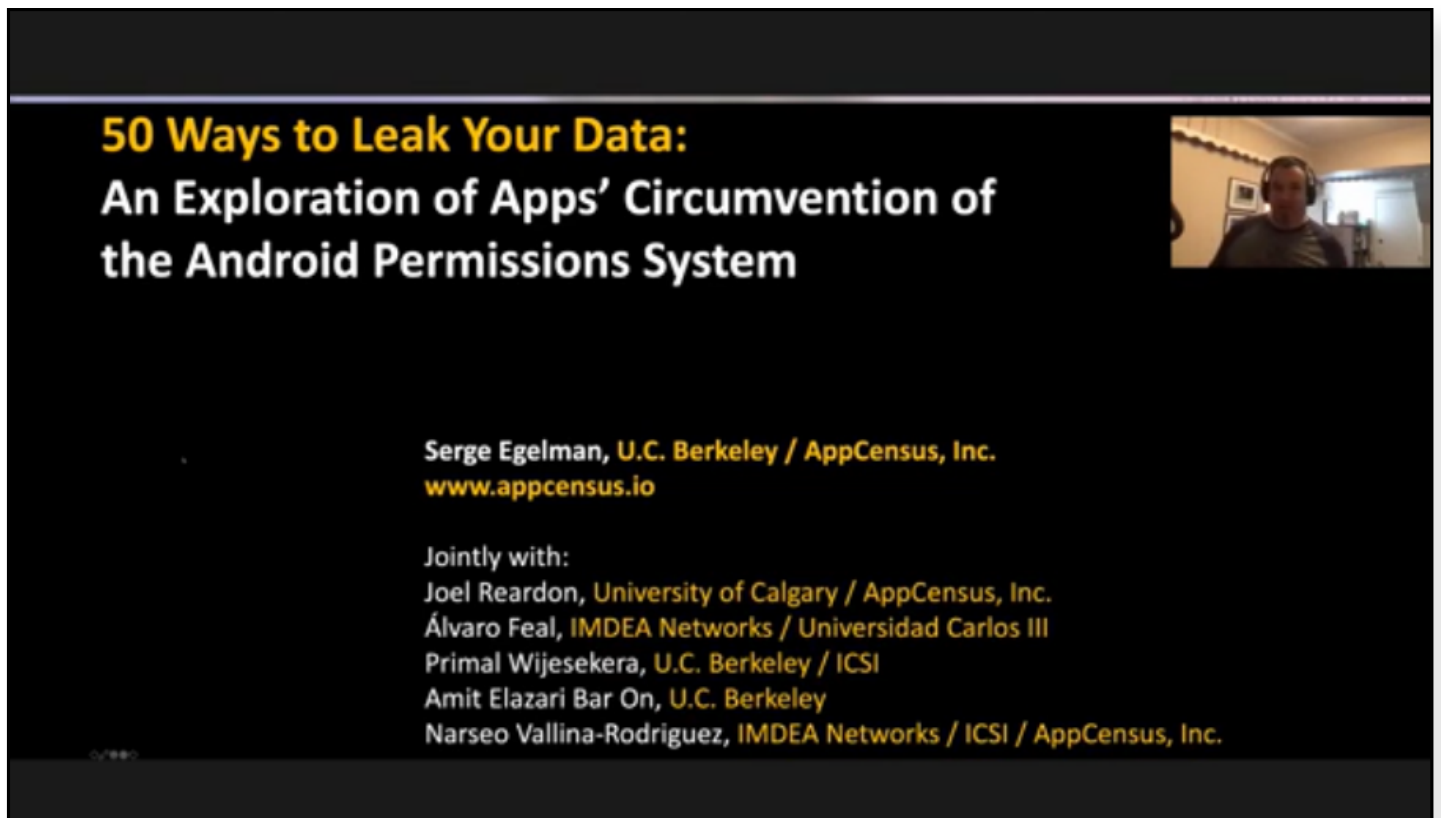
IMPACT ON HARD PROBLEM

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


- Serge Egelman gave numerous talks and media interviews about this project, specifically how apps are tracking users. In 2020 he was interviewed by publications including the Washington Post, Consumer Reports and Cnet.
- Serge Egelman gave a guest lecture in a law school class about this research.



**50 Ways to Leak Your Data:
An Exploration of Apps' Circumvention of
the Android Permissions System**

Serge Egelman, U.C. Berkeley / AppCensus, Inc.
www.appcensus.io

Jointly with:
Joel Reardon, University of Calgary / AppCensus, Inc.
Álvaro Feal, IMDEA Networks / Universidad Carlos III
Primal Wijesekera, U.C. Berkeley / ICSI
Amit Elazari Bar On, U.C. Berkeley
Narseo Vallina-Rodriguez, IMDEA Networks / ICSI / AppCensus, Inc.



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

Project research and SoS outreach continued despite the pandemic, and we continued to produce science of security outcomes.

We have developed models on how vulnerabilities are detected by analyzing the data we have collected from the National Collegiate Penetration Testing Competition (CPTC) from years 2018 and 2019.

We held discussions with a local startup, Airgap Inc., regarding secure manufacturing.

We gave a conference tutorial that involved concepts of sociotechnical systems, norms, and privacy at the International Conference on Autonomous Agents and MultiAgent Systems (AAMAS) and the Conference on Programming Language Design and Implementation (PLDI).

Several Lablet members had a discussion session with Cisco regarding upcoming challenges in cybersecurity from the standpoint of authentication that Cisco is considering addressing in new products.



Laurie Williams



Munindar Singh

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

Come postdoc with us!

This position will be part of the NSA-sponsored Science of Security lablet which North Carolina State has been a part of since 2011. The position will also be part of the [Secure Computing Institute](#).

Reporting to:

- Laurie Williams
- Munindar Singh

The postdoctoral research associate will conduct research into the science of security (including writing and publishing the results of that research in peer-reviewed venues); mentor students in the science of security lablet; assist in preparing quarterly and annual reports from the science of security lablet; and assist in organizing meetings for the science of security lablet.

A successful candidate would have completed a PhD in computer science or in a closely related field, and would have completed the main requirements, including successful defense of a PhD dissertation.

Duties & Responsibilities:

- Conduct research into the science of security.
- Mentor students in the science of security lablet.
- Assist in preparing quarterly and annual reports from the science of security lablet.
- Assist in organizing meetings for the science of security lablet.

The screenshot shows the ACM/SIGAI website with a navigation bar (ABOUT, ACTIVITIES, AWARDS, ORGANIZATION). The main content area features a 'Membership' section with details for students (\$15) and others (\$25), and a list of 'Upcoming conferences' including FDG, ASE, EAAME, MSR, ICK, ICC, RSOEWS, KSPWAM, EE, and EAAME. A prominent section titled 'The ACM/SIGAI Autonomous Agents Research Award' includes a photo of Professor Munindar Singh and text stating he is the 2020 award winner. The text describes the award's history and the selection committee's decision to honor Professor Singh for his contributions to social interaction and autonomy in AI.

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

We designed and implemented runtime targeted patching techniques by extracting exploit signatures for different vulnerabilities and performing targeted patching for the detected vulnerabilities. Compared to existing techniques, our approach aims at achieving a more intelligent and efficient security patching scheme for container systems. We evaluated our scheme over 32 real world security vulnerabilities in 23 commonly used server applications. Results show that we can timely detect and classify 78% of the attacks before they succeed in exploiting the tested vulnerabilities. Compared to traditional patching approaches (i.e., whole software upgrade), our targeted patching scheme can reduce memory footprint by over 50% and disk consumption by 23%. The traditional upgrade approach can only fix 4 out of 32 tested vulnerabilities.

We further refined our runtime targeted patching system implementation and design and have started to refine our detection scheme to improve on those mis-detections by adding system call arguments into our analysis. We continued

our work on an aggregated learning framework to further improve anomaly detection accuracy for microservices system consisting of many ephemeral containers.

We completed the design and initial prototype implementation of Self-Patch, a new self-triggering patching framework for applications running inside containers. We further refined the design and implementation of CDL, a classified distributed learning framework to achieve efficient security attack detection for containerized applications.

We presented Self-Path, a self-triggering patching framework for containerized applications at ACSOS 2020. We completed CDL, a classified distributed learning framework for security attack detection in container-based systems. We further started to investigate automatic security patch generation techniques and enhanced security attack detection solutions using reinforcement learning.

A poster based on this work by Olufogorehan Tunde-Onadele,

IMPACT ON HARD PROBLEM

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

PUBLICATIONS

- Yuhang Lin, Olufogorehan Tunde-Onadele, and Xiaohui Gu, "CDL: Classified Distributed Learning for Detecting Security Attacks in Containerized Applications," In Proceedings of the *Annual Computer Security Applications Conference (ACSAC)*, Austin, TX, December, 2020.
- Olufogorehan Tunde-Onadele, Yuhang Lin, Jingzhu He, and Xiaohui Gu, "Self-Patch: Beyond Patch Tuesday for Containerized Applications," In Proceedings of the *IEEE International Conference on Autonomic Computing and Self-Organizing Systems (ACSOS)*, 2020, pages 1--7.

NC STATE UNIVERSITY

Toward Just-in-Time Patching for Containerized Applications

Olufogorehan Tunde-Onadele, Yuhang Lin, Jingzhu He, Xiaohui Gu
North Carolina State University

Motivation

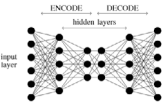
Containerized applications pose a set of new security challenges to distributed computing environments

- An alarming degree of vulnerability exposures exist in official image repositories (Shu et al. 2017)
- Significant resource increase in resource-limited containers can result after patching
- Traditional patching schemes that follow a scheduled whole upgrade approach (e.g., every Tuesday), do not work well for short-lived containers

Anomaly Detection

OPatch applies the unsupervised autoencoder neural network to detect abnormal system call frequency changes

- Does not require labelled training data which makes it robust to unknown attacks
- Achieves good accuracy with relatively few neurons and low training cost

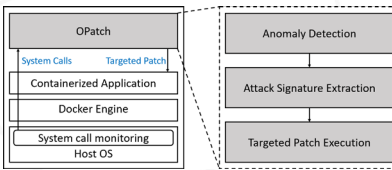


Contribution

OPatch is a new on-demand targeted patching framework for container environments

- We perform lightweight vulnerability exploit detection
- We achieve practical and effective security protection using a signature extraction scheme for identifying vulnerabilities
- We evaluate OPatch on 31 real world security vulnerability exploits in 23 commonly used server applications

Overview



OPatch is composed of three modules

- **Anomaly Detection** detects vulnerability exploits
- **Attack Signature Extraction** creates a signature to map exploits to the culprit vulnerability identifier
- **Targeted Patch Execution** triggers the proper software library update

Attack Signature Extraction

Attacks are characterized by a Secure Hash Algorithm (SHA) signature of the top frequent system calls

Application	Top System Calls				
Apache	write	fcntl	geteuid	getegid	switch
ActiveMQ	read	lseek	open	switch	futex
ImageMagick	read	lseek	open	switch	futex
Nginx	switch	poll	stat	writew	read

Results

- We evaluated OPatch over **31** real-world vulnerabilities discovered in **23** common containerized applications

Targeted Patch Execution

Targeted patching only installs the packages required by the application

```

# install files
of ghoscript-*.xx
./configure
make install

# remove files
of ghoscript-*.xx
apt-get purge -y wget gcc make
apt-get autoclean -y
cd ..
rm -f ghoscript-*.xx.tar.gz ghoscript-*.xx
                    
```

Whole Upgrade only updates applications handled by the package manager

```

> apt-get update
> apt-get upgrade
                    
```

Detection Rate


False Positive Rate (FPR)


Patching Cost

Conclusion

Our initial experimental results of OPatch are promising

- We can increase detection rate to over 80% and reduce false alarm rate to 0.7%
- Compared to the whole upgrade approach, OPatch can reduce the memory overhead by up to 84% and disk overhead by up to 40%





7TH ANNUAL
HOT TOPICS in the SCIENCE OF SECURITY
APRIL 7-8, 2020 | LAWRENCE, KANSAS

Development of Methodology Guidelines for Security Research



Jeffrey Carver

PI: Jeffrey Carver

SUB_ LABELLET: University of Alabama

URL: https://cps-vo.org/UIUC_AutomatedSynthesisFramework

HARD PROBLEM: Security Metrics and Models



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 guidelines that provide insight into the scientific rigor of the information included in a research report (i.e., a journal or conference paper). By providing guidelines to help 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 these guidelines, we are interviewing experts from the cybersecurity community, both from within the Lablets and from

outside. The goal of these interactions is to determine what type of information is most important for judging scientific rigor in different portions of the cyber security community. To ensure that the guidelines are 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 also 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 security research community to have input into the contents of the guidelines. By gathering input from a wide swath of researchers representing different perspectives, our guidelines should be more widely accepted in the larger cyber security research community.

ACCOMPLISHMENTS

Our interviews with experts have yielded valuable insights and knowledge into what makes cybersecurity research scientific. In addition to interviewing experts from inside the Lablets, we broadened the scope of expertise by including experts who are associate researchers to the Lablets and focus on in industry-led research, the philosophy of ethics in cybersecurity, reverse-engineering and hardware assurance, security in distributed web systems, and research ethics. The diverse representation of topics within and around cybersecurity research allowed us to include new perspectives and valuable data points in our analysis that we would have otherwise missed from the traditional cybersecurity settings. To this end, we continue towards a better and more versatile rubric or set of guidelines as we continue to interview additional experts. The interviews

reveal the scope of information the Science of Security and the Paper Review guidelines will have to contain to address different types of cybersecurity research papers. Our key finding from the interviews is that guidelines that address a wide variety of cybersecurity topics will be complex and potentially large. Our engagement with the community has been frequent and largely positive in regards to the project goals despite the challenge of the pandemic.

We have made significant progress on our "Good Examples" paper that presents examples of good practices in scientific reporting from papers published in IEEE S&P and ACM CCS. The knowledge gained from analyzing these publications will be helpful not only for providing validity to our interview findings

but also for increasing the acceptance rate of the conclusions drawn in the larger community. We plan to develop an initial draft of the guidelines based on the data from the interviews and "Good Examples" paper. We would like to get feedback on this initial draft either through a workshop or through some other types of interactions. We hope to get feedback on how well the guidelines are organized, how well they can be used,

and what is missing or needs modification. In parallel to this first draft of the guidelines, we will begin identifying cyber security experts outside the Lablets who can help ensure the validity of our conclusions and provide any additional perspectives that we may have missed. We have begun a new study that will supplement our interview data by analyzing the comments left by reviewers on submissions to the HotSoS symposium.

IMPACT ON HARD PROBLEM

This project addresses the challenges of metrics in regard to research methods and community. The guidelines we construct will provide a basis upon which to judge the scientific rigor in the reporting of cyber security research. These guidelines

will have community input to ensure they capture all relevant perspectives and is applicable to different types of research methodologies and problems.

Development of Methodology Guidelines for Security Research

Jeffrey Carver & Matthew Armstrong
carver@cs.ua.edu
maarmstrong3@crimson.ua.edu
University of Alabama



WHERE LEGENDS ARE MADE®



College of Engineering

Predicting the Difficulty of Compromise through How Attackers Discover Vulnerabilities



Andy Meneely



Laurie Williams

PI: Andy Meneely CO-PI Laurie Williams

SUB-LABEL: Rochester Institute of Technology

URL: https://cps-vo.org/NCSU_PredictingDifficultyofCompromise

HARD PROBLEM: Security Metrics and Models



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

National collegiate penetration testing competitions in 2018 and 2019 allowed us to collect a massive data set. 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.

We collected 9 terabytes and over a billion events from CPTC 2019 Regionals and CPTC 2019 Nationals. Having been more involved with the instrumentation this time, we have an even better look at what attackers are doing in a controlled, competition environment. We adapted our CPTC 2018 observations to queries so that we can extract timelines even faster for the 2019 data set.

We have conducted a literature search on existing stochastic

models of attacker behavior and believe we know the existing literature on the topic. None of the approaches have such a fine-grained data set as ours, and many of the models were theoretical to begin with. Using what we have learned from these theoretical models, we developing our own stochastic models for simulating attacker behavior. If we can develop these models and use CPTC data to train it, we can accurately predict which vulnerabilities are more discoverable based on various conditions.

We collated the vulnerability reports from CPTC 2019 and found a total of 67 vulnerabilities reported from the teams, and we constructed timelines from those vulnerabilities using the techniques we developed on studying CPTC 2018 data.

We began work on a model to assist in helping tag the timeline according to the MITRE ATT&CK framework. Our initial model trained on the CPTC 2019 data had an F1 measure of 59%, meaning that it is likely that we will be able to construct a robust model to assist in mapping timeline events to the MITRE ATT&CK framework, improving curation efforts.

After we finished collecting our timeline data from the CPTC 2019 competition data logs, we developed improved queries for collecting these timelines to speed up future data collection. We continued to evaluate (nine) state-of-the-art vulnerability detection tools and are interpreting our findings.

We advanced our analysis of the CPTC 2019 and CPTC 2018 competition data. Our ML classifier so far achieves precisions and recalls of over 70% on a sample dataset.

We assisted in the instrumentation for the CPTC 2020 competition. Our data collection is reduced from last year to focus on analyzing existing data.


IMPACT ON HARD PROBLEM

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.

NationalCPTC

- About
- Compete
- Regions
- Volunteers
- Research
- Contact Us



The Collegiate Penetration Testing Competition:

Training and Evaluating the next generation of cyber security professionals.

Principles of Secure Bootstrapping for IoT



Ninghui Li

PI: Ninghui Li

SUB-LABEL: 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.

ABSTRACT

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

ACCOMPLISHMENTS

We designed an enhanced bootstrapping protocol for Zigbee that prevents a wide range of attacks. Our investigation uncovered a number of critical security and privacy issues in the connection establishment (also known as the “joining”) procedure of Zigbee protocol. To mitigate these issues, we designed and implemented an enhanced connection establishment procedure. In this solution, we leverage the existing installation code mechanism to use it as public-key cryptography and combine it with the Elliptic-Curve Diffie-Hellman (ECDH) mechanism to ensure better security and privacy guarantees.

We evaluated our proposed enhancements to the Zigbee protocol, which we proposed to avoid vulnerabilities in Zigbee that we previously identified. We used ProVerif to verify the correctness of the proposed protocols. We implemented and deployed our enhanced protocol to evaluate and compare with the Zigbee standard implementations in terms of delay, memory usage, and message size. We found that the enhanced

protocol does not introduce extra messages and induces only 3.8% overhead on average for the entire join procedure.

We started working on two new topics, going beyond our previous work with Zigbee: one is Connected Vehicle Systems security, and the other is contact tracing. We identified important challenges in connected vehicles, specifically, their keyless (i.e., fob-based) entry systems and on-board diagnostic systems. We found that existing approaches expose a large threat surface that could be exploited to impersonate a vehicle owner, gain control of a vehicle, or steal private information. We also looked at security and privacy concerns in mobile contact tracing apps.

With a focus on studying contact tracing protocols, we developed a framework to analyze Proximity-based Contact Tracing (PCT) protocols. We have identified two main dimensions along with which different designs for PCT protocols can be made.

IMPACT ON HARD PROBLEM

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.

PUBLICATIONS

- Weicheng Wang, Fabrizio Cicala, Syed Rafiul Hussain, Elisa Bertino, and Ninghui Li, "Analyzing the Attack Landscape of Zigbee-enabled IoT Systems and Reinstating Users'

Privacy," *13th ACM Conference on Security and Privacy in Wireless and Mobile Networks (ACM WiSec 2020)*.


ACM
WiSec
2020


13th ACM Conference on
Security and Privacy in Wireless and Mobile Networks


Teaser


Analyzing the Attack Landscape of Zigbee-enabled IoT Systems and Reinstating Users' Privacy

Weicheng Wang, Fabrizio Cicala, Syed Rafiul Hussain, Elisa Bertino, and Ninghui Li

 Association for
Computing Machinery
Advancing Computing as a Science & Profession

 INSTITUTE
OF NETWORKS
AND SECURITY

 JKU
JOHANNES KEPLER
UNIVERSITY LINZ



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. Through manual collection and examination of app reviews that describe spying activities with apps, we have determined the necessity of considering app reviews for identifying apps that can aid spying, either explicitly or through misuses. Specifically, we are concerned about intimate partner surveillance spying. Based on this understanding, we are developing a computational framework for spotting such apps, in which we first identify apps that can potentially be misused for spying based on their metadata (e.g., their descriptions and permissions), collect their reviews, and determine their spying capability based on user-reported stories.

We started building a computational framework which by analyzing app reviews identifies if that app facilitates spying

activity. We conducted a preliminary investigation to identify app reviews that were relevant to spying. We observed that relevant app reviews differ greatly in terms of the severity of the problem leading us to investigate how we can automatically determine the severity of the app's spying capability described in an app review. We are designing an annotation scheme for crowdsourcing the annotation of reviews based on their severity.

We performed a systematic analysis of the network protocol exchanges used by Payment Service Providers (PSPs). Through formal modeling using the Tamarin Prover, we identified four vulnerabilities in these SDKs and demonstrated proof-of-concept exploits for four payment service providers. We have reported these vulnerabilities to these providers.

We continued our analysis of Payment Service Provider (PSP) application programming interfaces (APIs), developing models

for analyzing the security of code in Software Development Toolkits (SDKs). We completed a systematic literature review of research works on mining threat intelligence from unstructured textual data.

We extended our scope from spying to Unexpected Information Gathering (UIG) in mobile apps, and identified 124 UIG-enabling apps from our current dataset of apps. We identified an additional 131 UIG-enabling apps in a snowball fashion.

Healthcare professionals use mobile apps to store patient information and communicate with their patients, but not all such apps are HIPAA compliant. We began investigating HIPAA compliance of medical mobile apps on the Apple App Store. In a preliminary investigation, we identified 899 medical apps that were potentially relevant but did not mention HIPAA compliance in their descriptions. We are investigating these 899 medical apps further to determine their compliance with HIPAA.

IMPACT ON HARD PROBLEM

This project addresses the hard problem of policy-governed secure collaboration. Specifically, Cardpliance 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

- Munindar Singh gave a talk on Engineering Ethical Multiagent Systems at the University of Wollongong, Australia, in which he discussed concerns of policy and privacy.

PUBLICATIONS

- Md Rayhanur Rahman, Rezvan Mahdavi-Hezaveh, and Laurie Williams, "A Literature Review on Mining Cyberthreat Intelligence from Unstructured Texts," in Proceedings, 1st IEEE ICDM Workshop on Deep Learning for Cyber Threat Intelligence (DL-CTI), November, 2020, pages 1—10.
- Samin Yaseer Mahmud, Akhil Acharya, Benjamin Andow, William Enck, and Bradley Reaves, "Cardpliance: PCI DSS Compliance of Android Applications," in Proceedings of the USENIX Security Symposium, 2020.
- Hui Guo and Munindar Singh, "Caspar: Extracting and Synthesizing User Stories of Problems from App Reviews," in Proceedings of the 42nd International Conference on Software Engineering, 2020, pages 628--640. doi: 10.1145/3377811.3380924.

NC STATE UNIVERSITY

Application Study

Cardpliance Analysis

- We ran our 6 PCI checks on 358 applications
- Cardpliance reported 20 applications violated at least one PCI check
- Another 20 applications have bad **SocketFactory** classes


Manual Validation

- We manually validated 40 apps over a one month period with the JEB decompiler
- We confirmed **15 PCI violations across 6 applications**

App Name	Downloads	T1	T2	T3	T4
Credit Card Reader	500K+	✗			✗
Fast Toll Illinois	10K+	✗	✗		✗
Bens Soft Pretzels	10K+	✗	✗	✗	✗
The Toll Roads	100K+	✗	✗		✗
Connect Network by GTL	1M+			✗	
Peach Pass GO!	50K+	✗			✗



University of Illinois *at* 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-PI David Nicol. 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.



PI Sayan Mitra



CO-PI David Nicol

We continued our research projects during the pandemic publishing multiple papers, participating in conferences and engaging in multiple education and outreach initiatives virtually.

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 Security and Resilience
- Monitoring, Fusion, and Response for Cyber Resilience
- Resilient Control of Cyber-Physical Systems with Distributed Learning
- Uncertainty in Security Analysis



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

We extended our work focused on a metamodeling-based approach to sensitivity analysis and uncertainty quantification in complex security models. To review, many realistic security models run slowly and have input variables whose values are uncertain, which makes it difficult to conduct sensitivity analysis and uncertainty quantification. It is possible to create metamodels of the base security model that trade some accuracy for speed using machine learning techniques. Earlier, we had investigated this method by applying it to a previously-published work that models the growth of peer-to-peer botnets, and we applied it to two new models to test its general applicability.

We investigated two ways to solve an issue with applying our metamodeling approach to certain models that contained a mix of quantitative and qualitative input variables. The two approaches were one-hot encoding and splitting. We implemented the two approaches and evaluated them on an AMI ADVISE model, and found that at least in that one case that splitting substantially outperformed one-hot encoding. This work can help modelers apply the metamodeling approach we developed to a broader class of security models. The metamodeling approach helps modelers perform sensitivity analysis and uncertainty quantification on complex slow-running security models that contain uncertain input variables.

This project concluded in 2020.

IMPACT ON HARD PROBLEM

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.

PUBLICATIONS


- Michael Rausch and William Sanders, "Sensitivity Analysis and Uncertainty Quantification of State-Based Discrete-Event Simulation Models through a Stacked Ensemble Metamodel," In Proceedings of the *Quantitative Evaluation of SysTems (QEST)*, 2020. Best Paper Award.



An Automated Synthesis Framework for Network Security and Resilience







Matt Caesar



Dong (Kevin) Jin

PI:	Matt Caesar	CO-PI Dong (Kevin) Jin
SUB-LABEL:	Illinois Institute of Technology	
URL:	https://cps-vo.org/UIUC_AutomatedSynthesisFramework	
HARD PROBLEM:	Resilient Architectures (Primary), 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

We continued the transfer of our technology to industry through interactions with Veriflow and VMWare. Current collaborations target enhancement of our verification technology to operate on real-time traffic data. We have continued our collaboration with AT&T to customize and deploy our technology in their environments, and we have continued our collaborations with Boeing on constructing a resilient IoT platform for the battlefield. We have also made progress constructing a real drone implementation which we will use to evaluate our design, though this work has been somewhat delayed due to the lockdown. Most recently, we have been focusing on improving the attack resilience of our algorithms. We have developed new deep learning mechanisms that are resilient to data sets that are “constructed” by adversaries, and our early simulation results show some benefits to these approaches in practical settings.

We have begun a collaboration with a security company (Censys) to perform wide-area verification of networked services, and have begun discussions on approaches to incorporate their data into our synthesis frameworks.

We continued the exploration of self-healing network management to address the resilient architecture hard problem with applications in wide-area monitoring systems in power grids. We developed an optimization-based self-healing scheme that quickly recovers PMU network connectivity and restores power system observability against cyber-attacks. We developed a proof-of-concept system based on scenarios derived from the IEEE 30-bus and 118-bus systems. A paper describing this work received the best paper award in the 2020 SmartGridComm conference.

We continue to study the interdependence between the power system and the communication network with the goal of improving resilience in critical energy infrastructures. We developed an accurate model representing the interdependencies between the two systems so that one can construct realistic communication networks to meet specific power system requirements for planning and evaluation. To address the scalability problem, we are conducting extensive experiments with a large-scale power system consisting of thousands of buses.

IMPACT ON HARD PROBLEM

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

- Matthew Caesar was selected to serve as the General Chair for ACM SIGCOMM 2021. He will also serve on the program committee.
- Matthew Caesar created and operates a new Slack workspace for the SIGCOMM community. The platform, now with over 1400 members, enables participants to discuss security and networking topics and includes a channel to discuss topics related to the science of security.
- Matthew Caesar helped create and served as co-chair for an ACM SIGCOMM workshop on "Teaching and Learning Computer Networking During the Pandemic". The workshop was a great success, attracting over 200 participants across academia and industry.
- Matthew Caesar was selected to serve as the mentoring chair for ACM SIGCOMM 2020 and helped design the conference to be the first "virtual" SIGCOMM conference.
- Matthew Caesar was selected to serve on the program committee for ACM CCS 2021.
- Matthew Caesar was selected to serve on the program committee for ACM NSDI 2021.
- Matthew Caesar continues to serve as an Editor for IEEE/ACM Transactions on Networking, and also serves as the editor for the Education track of the ACM SIGCOMM Computer Communication Review.
- Matthew Caesar has continued an engagement with the University of Illinois Center for Digital Agriculture towards securing our nation's food supply. His work leverages machine learning to detect anomalies in supply-chain operations. He is in the process of conducting a prototype deployment of his work within the Imported Swine Research Laboratory (ISRL) farm on the UIUC campus.
- Kevin Jin was chosen to serve on the program committee for ACM SIGSIM-PADS 2021.
- Kevin Jin was chosen to serve on the program committee for IEEE SmartGridComm 2020.
- Kevin Jin organized a virtual Ph.D. colloquium as part of the ACM SIGSIM-PADS conference in June 2020. The Ph.D. colloquium included a keynote speech and multiple student presentations with 99 attendees.
- Kevin Jin served as the web chair for the 2020 ACM SIGCOMM Symposium on SDN Research (SOSR).

PUBLICATIONS



- Dong Jin, Yanfeng Qu, Xin Liu, Christopher Hannon, Jiaqi Yan, Alex Aved, and Philip Morrone, "Dynamic Data-Driven Approach for Cyber Resilient and Secure Critical Energy Systems," Handbook on Dynamic Data-Driven Application Systems (DDDAS) (Vol. II), Book Chapter. Accepted for publication.
- Yanfeng Qu, Gong Chen, Xin Liu, Jiaqi Yan, Bo Chen, and Dong Jin, "Cyber-Resilience Enhancement of PMU Networks Using Software-Defined Networking," 2020 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm), November 2020. Best Paper Award.
- Yanfeng Qu, Xin Liu, Jiaqi Yan, and Dong Jin, "Dynamic Data-Driven Self-Healing Application for Phasor Measurement Unit Networks," Third International Conference on InfoSymbiotics/DDDAS 2020. October 2020.
- Steve Uhlig, Ralph Holz, Olivier Bonaventure, Matt Caesar, et al, "Update on ACM SIGCOMM CCR Reviewing," ACM SIGCOMM Computer Communications Review, October 2020.
- Bingzhe Liu, Ali Kheradmand, Matthew Caesar, and Brighten Godfrey, "Towards Verified Self-Driving Infrastructure, ACM Workshop on Hot Topics in Networks (HotNets), November 2020.
- Santhosh Prabhu, Kuan Yen Chou, Ali Kheradmand, Brighten Godfrey, Matthew Caesar, "Plankton: Scalable Network Configuration Verification Through Model Checking," NSDI, February 2020.
- Xin Liu, Bo Chen, Chen Chen, and Dong Jin. "Electric Power Grid Resilience with Interdependencies between Power and Communication Networks — A Review," IET Smart Grid.

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 goal of this project is to improve the resilience of a system's intrusion and incident detection infrastructure against monitor compromise.

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 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

Our Response and Recovery Engine (RRE) work incorporates modules to monitor current state of a system, detect intrusions, and respond to achieve resilience-specific goals. Intrusion detection in large-scale distributed systems, which is a necessary precondition for intrusion tolerance and resilience, is highly susceptible to malicious manipulation of system data used for detection (e.g., using rootkits and log tampering), which we term "monitor compromise". Existing literature attempts to counteract the problem using reputation systems, which weight the trustworthiness of monitor data based on past trustworthiness of the data, but such systems are themselves subject to "betrayal attacks" and "sleeper attacks". We instead propose the use of data-driven methods for detecting potential monitor compromise. We leverage the insight that systems usually contain multiple monitors that provide redundant

information about system activity so we can use discrepancies between observations of system activity across different monitors to identify potential monitor compromise.

Our work on using metamodels to indirectly perform sensitivity analysis and uncertainty quantification on complex and long-running cyber security models won a Best Paper Award at QEST 2020. Using our work, sensitivity analysis and uncertainty quantification can be accomplished thousands of times faster than using traditional methods, and with more accuracy than competing metamodeling approaches. The work we did should make it easier to validate the performance of the cyber security models and allow modelers to gain confidence in the model results.

IMPACT ON HARD PROBLEM

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 through 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).

PUBLICATIONS

- Benjamin Ujcich, Adam Bates, and William Sanders, “Provenance for Intent-Based Networking”, *IEEE Conference on Network Softwarization (NetSoft '20)*.
- Benjamin Ujcich, Samuel Jero, Richard Skowyra, Steven Gomez, Adam Bates, William Sanders, and Hamed Okhravi, “Automated Discovery of Cross-Plane Event-Based Vulnerabilities in Software-Defined Networking”, *2020 Internet Society’s Network and Distributed System Security Symposium (NDSS '20)*.
- Michael Rausch and William Sanders, “Stacked Metamodels for Sensitivity Analysis and Uncertainty Quantification of AMI Models,” In *Proceedings of the IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm), 2020*.

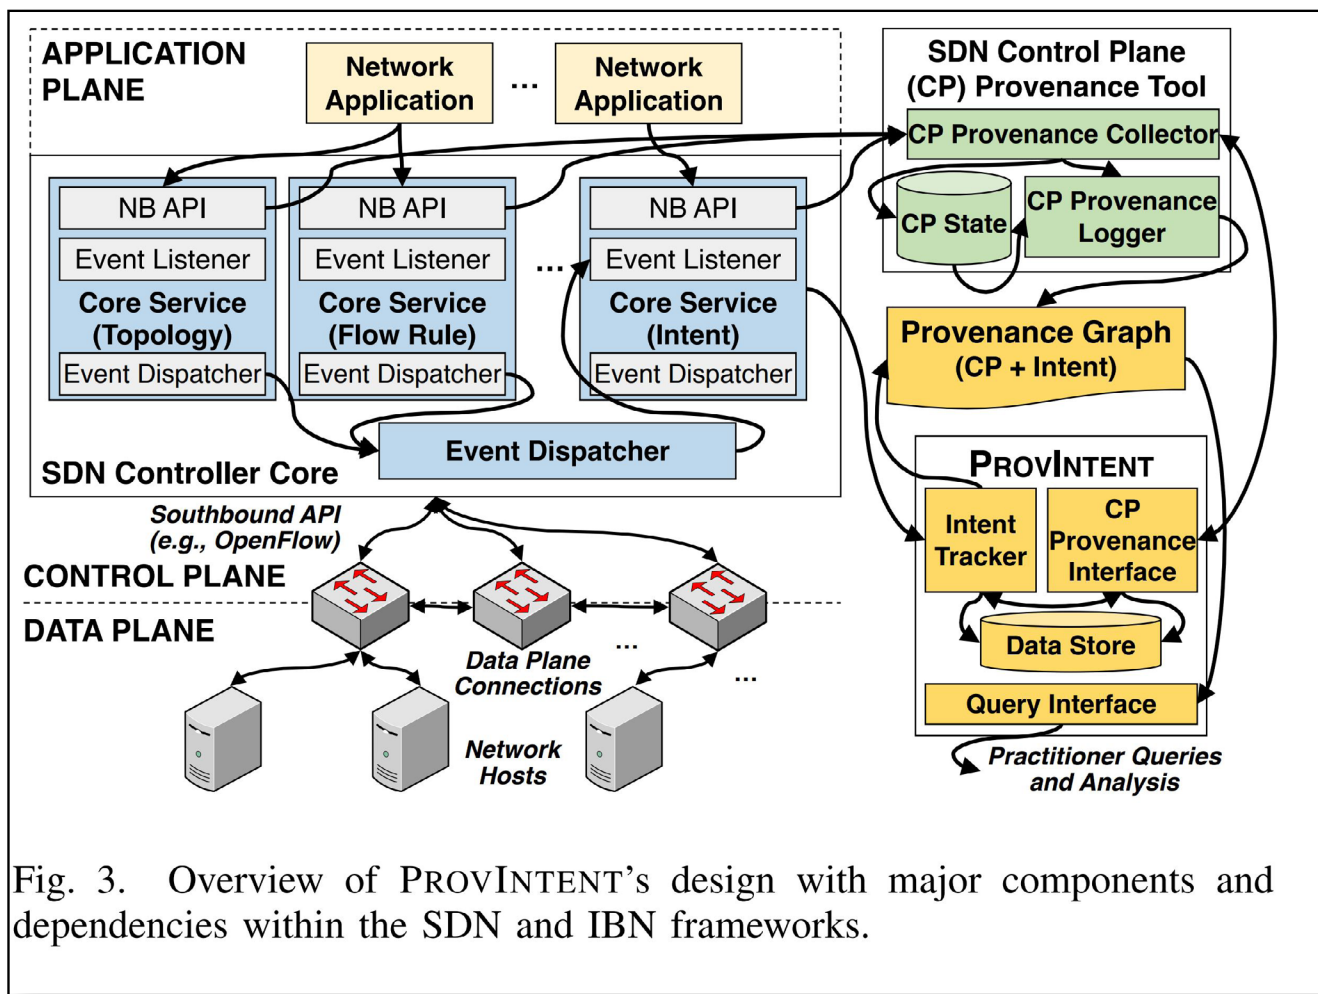


Fig. 3. Overview of PROVINTENT’s design with major components and dependencies within the SDN and IBN frameworks.

Resilient Control of Cyber-Physical Systems with Distributed Learning



Sayan Mitra



Geir Dullerud



Sanjay Shakkotai

PI: Sayan Mitra CO-PIs: Geir Dullerud, Sanjay Shakkotai

SUB-LABEL: University of Texas at Austin (UT Austin)

URL: https://cps-vo.org/UIUC_ResilientControlofCPSDistributedLearning

HARD PROBLEM: Security Metrics and Models, Resilient Architectures



GOAL

The goal of the project is to bring together techniques from Machine Learning (ML) and formal verification to improve resiliency and risk-reduction in autonomous systems and Cyber-Physical Systems (CPS) (e.g., autonomous vehicles, drones, and manufacturing systems). Verification solves the problem of catching design bugs and vulnerabilities and it can provide high-assurance guarantees. We are exploring how existing model-based verification methods like statistical model checking and

reachability analysis can be improved using ML techniques like multi-armed bandits algorithms. The research aims to provide theoretical bounds of how testing and validation budget can be best allocated based on metrics like sample complexity (i.e., how many test samples do we need to draw to answer a risk / resiliency question). 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 exposes 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, we exploit models

when it makes sense but otherwise get (gracefully degrading) guarantees with black-box executables. We also focus on the marginal benefits of models/executable fidelity in security and risk reduction. The developed approaches are being evaluated on case studies drawn from autonomous vehicles and aircraft systems.

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 existing model checking approaches such as Storm and Prism. The prototype tool has been made available online. In a related thread, we have also developed an algorithm for probabilistic reachability analysis that can work on black-box dynamical systems. A paper based on this research is under review.

IMPACT ON HARD PROBLEM

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

- The third edition of PI Sayan Mitra's new course *Principles of Safe Autonomy* at the University of Illinois came to a successful conclusion. The course takes a deep dive into the seminal topics in object recognition, localization, decision making, path planning, and safety verification. With support from the Illinois Center for Autonomy, we have set up a laboratory with 7 workstations with GPUs for performing simulation-based experiments. In addition, this is a unique class in which student get to deploy and test their code on an actual autonomous vehicle platform.
- Sanjay Shakkottai, Lecture, "Hyper-parameter Tuning for ML Models: A Monte-Carlo Tree Search (MCTS) Approach," EE Seminar, Indian Institute of Science, January 2, 2020.
- Geir Dullerud, Invited Speaker, Formal Methods in Mathematics Workshop, Mathematics, Carnegie Mellon University, "Learning and Statistical Validation of Complex Cyber-Physical Systems," January 7th, 2020 on logic in engineering systems.
- Sanjay Shakkottai, Lecture, "On the Throughput vs Accuracy Trade-Off for Streaming Unsupervised Classification," Workshop on Learning Theory 2, Tata Institute of Fundamental Research, January 3, 2020.
- Sanjay Shakkottai, Invited Speaker, "On the Throughput vs Accuracy Trade-Off for Streaming Unsupervised Classification", Workshop on Learning Theory 2, Tata Institute of Fundamental Research, January 3, 2020.
- Sanjay Shakkottai, Invited Speaker, "Learning and Resource Allocation in Networks: Finite Time Bounds and Insights," at the Symposium on Advances in Communication Networks, Indian Institute of Science, Bangalore, July 10, 2020.
- Sanjay Shakkottai, Invited Panelist, "Pursuing high-impact activities: Advise from distinguished scientists in the NeTS community," at the NSF NeTS Community CAREER Webinar, June 17, 2020.
- Sanjay Shakkottai, Invited Panelist and Panel Moderator, "Warm starting adaptive interventions with side information from confounded logs," The First NSF NeTS Community Workshop Call to Arms Workshop on modeling, analysis and mitigation of COVID-19, April 13, 2020; invited talk University of Arizona at Tuscon Applied Math COVID-19 Working Group, May 12, 2020.
- Geir Dullerud, Plenary Lecture, "Statistical Validation and Principle-Based Simulation of Complex Cyber-Controlled Systems," Forum on Robotics & Control (FoRCE), September 4th, 2020.
- Sayan Mitra, Lecture, "Formal Verification of Cyber-physical Systems: Opportunities and Challenges," at a mini-course on at NIT Jalandhar, India, September 12, 2020.
- Sayan Mitra's text book "Verifying Cyber-Physical Systems" will be published by MIT Press in Spring of 2021.

PUBLICATIONS

- Ronshee Chawla, Abishek Sankararaman, Ayalvadi 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.
- Negin Musavi, Dawei Sun, Sayan Mitra, Sanjay Shakkottai, and Keir Dullerud, "Optimistic Optimization for Statistical Model Checking with Regret Bounds" Available online from <https://arxiv.org/abs/1911.01537> April 2020; *Workshop on Symbolic and Numerical methods for Reasoning about Cyber-Physical Systems*, Jul 2020.
- Nihal Sharma, Soumya Basu, Karthikeyan Shanmugam and Sanjay Shakkottai, "Warm Starting Bandits with Side Information from Confounded Data," arXiv 2002.08405, 2020. Available at: <https://arxiv.org/abs/2002.08405>
- Yu Wang, Nima Roohi, Matthew West, Mahesh Viswanathan, and Geir Dullerud, "Verifying PCTL Specifications on Markov Decision Processes via Reinforcement Learning," *IEEE Conference on Decision and Control*, December 2020.
- Joao Porto, Bin Hu, and Geir Dullerud, "Convergence Guarantees of Policy Optimization Methods for Markovian Jump Linear Systems," in *Proceedings of the American Control Conference*, 2020.
- Joao Porto, Bin Hu, and Geir Dullerud, "Policy Learning of MDPs with Mixed Continuous/Discrete Variables: A Case Study on Model-Free Control of Markovian Jump Systems," in *Proceedings of the Learning for Dynamics and Control Workshop*, 2020.
- Sayan Mitra, "Verifying Cyber-Physical Systems: A Path to Safe Autonomy," to be published by MIT Press, February 16, 2021.

Uncertainty in Security Analysis



David Nicol

PI: David Nicol

URL: https://cps-vo.org/UIUC_UncertaintyinSecurityAnalysis

HARD PROBLEM: Policy-Governed Secure Collaboration, 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 inflict 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 critical components in the CPS and to inflict significant damage* 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 *uncertainty* 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 this study, we introduce the *attack loss*, a function that quantifies the loss to the system given the event of an adversary reaching a specific set of hosts. As the ability of an adversary to reach a set of hosts is uncertain and as we model uncertainty using probability, the overall loss caused by an attack is represented by an attack loss distribution. While previous analysis focused on computing *reachability*, i.e., the probability that a pathway exists between a specifically chosen source and destination host, current analysis focuses on techniques for quantifying the attack loss distribution. In particular, the right tail of the distribution contains the "worst-case scenarios" where the attack inflicts the largest amount of losses. Understanding high-impact security events and their probabilities allows businesses to make risk-informed defense decisions whether to reduce the risk, e.g., by investing in network hardening solutions, or to transfer the residual risk, usually by purchasing some form of cyber-insurance

ACCOMPLISHMENTS

Determining the exact probability of high-impact security events is a computationally hard problem, which leaves us with Monte Carlo methods as the primary tool for evaluation. However, the problem also falls in the category of *rare event estimation* as in a well-designed system, the possibility of an attacker being able to inflict large losses is small. It is well known that a naive Monte Carlo approach is inefficient in such cases. In this study, we proposed a simulation technique based on *importance sampling* to quantify the *loss tail probability*, the probability that the attack loss is greater than a selected threshold. The technique provides an unbiased estimate of the loss tail probability with the assurance of bounded relative error, meaning the relative error of the estimate remains bounded by

a constant as the edge probabilities tend to zero. Performing importance sampling requires solving the mindset-maxprob problem, problem, which is unfortunately NP-hard; however, the problems is easier to solve for the cases of very large losses.

We also came up with another simulation technique based on importance sampling for estimating the loss tail probability. The main idea is to strike the right balance between achieving a high "hitting rate", i.e., by maintaining a sufficient number of "rare event" samples to use in the computation, while at the same time minimizing the effect of *weight degeneracy*, i.e., by minimizing the maximum weight of rare event samples. Under the assumption that the attack loss is *monotone* with respect to

the edge probabilities, we formulate the “minimization of the maximum weight” problem as a stochastic program, translate it into a convex optimization problem, and use commercially available tools to solve and obtain the optimal parameters for importance sampling. Initial results show that this technique is faster and more accurate than the cross-entropy method and

our previous method based on solving the mindset-maxprob problem.

A poster based on this work by Hoang Hai Nguyen, “An uncertain graph-based approach for cyber-security risk assessment,” was presented at Hot Topics in the Science of Security (HotSoS 2020), September 22-24, 2020.

IMPACT ON HARD PROBLEM

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.

PUBLICATIONS

- Hoang Hai Nguyen and David M. Nicol, “Estimating Loss Due to Cyber-attack in the Presence of Uncertainty,” *IEEE TrustCom 2020*.



ILLINOIS

Information Trust Institute

GRAINGER COLLEGE OF ENGINEERING

Illinois Science of Security (SoS) Lablet

HOME
PROJECTS – IN PROGRESS
DIRECTORY
EVENTS
SUMMER INT


Uncertainty in Security Analysis



Investigator: David Nicol

Researcher: Hoang Nguyen

Models proliferate in security analysis. Models of systems typically include identification of devices and their interconnection. The models may include identification of software services running on some or all of those devices, device configurations, and descriptions of attackers and defenders. The problem is that in practice the information one has about the system is incomplete. There is considerable uncertainty about the predictions of a model as a result of uncertainty about the components of the model. This project aims to provide some rigor (both formal and experimental) to uncertainty analysis of security models. There is a large literature on uncertainty quantification in physical systems; there are in fact at least two journals dedicated to it. While there is much to learn from the ideas developed in that context, a significant difference exists. Physical models have some “ground truth” to which analysis can appeal and quantify variation in results relative to ground truth. In discrete models, and for security analysis in particular, such ground truth as exists is very specialized. For example, it may or may not be possible for an attacker to make lateral movement through a network from host h_A to h_B . Whether it is possible or not is a binary outcome. A security analysis focused on the possibility cannot with certainty say that it is not possible if the model analyzed fails to include crucial specifics. Early work on capturing uncertainty in connectivity is our own paper “An Approach to Incorporating Uncertainty in Network Security Analysis,” which first brought an uncertainty model to network reachability analysis. This paper analyzed uncertainty in the existence of paths allowing an attacker lateral movement. The project will extend the approach to include uncertainty in attacker and defender behavior, and use a more sophisticated model of belief function for traversing a potential connection.



University of Kansas



Perry Alexander

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.

This year the Lablet continued work on five projects on resiliency, IoT and cloud privacy, preventing side channel communication, and developing semantics and infrastructure for trust, and secure native binary execution. Specifically, we are: (i) developing a method to enable cloud-assisted, privacy-preserving machine learning classification over encrypted data for IoT devices; (ii) reducing micro-architectural side-channels by introducing new OS abstractions while minimally modifying micro-architecture and OS; (iii) developing an epistemology and ontology for framing resilience; (iv) formalizing the remote attestation and defining sufficiency and soundness; and (v) developing a framework for client-side security assessment and enforcement for COTS software.

The KU Lablet is supported by The Information and Telecommunications Technology Center (ITTC), an interdisciplinary research center focused on all aspects of information and its applications. ITTC is a designated KU Research Center and receives core support for research activities, including the Lablet. ITTC researchers working on Lablet projects are drawn from departments 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, and Prasad



Kulkarni. The KU Sub-Lablet at the University of Tennessee is led by Michael Jantz. Approximately 12 graduate students and 6 faculty are supported by the program assisted by several research staff and undergraduate students. Our 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.

The KU Lablet hosted the Hot Topics in The Science of Security (HoTSoS 2020) conference online September 22-24. Perry Alexander (KU) served as General Chair with Drew Davidson (KU) and Baek-Young Choi (UMKC) as Program Co-Chairs. Over 400 individuals registered for the conference. HoTSoS hosted four Keynote presentations 12 traditional papers, 20 posters and 6 Works in Progress papers. The Works-in-Progress (WiP) papers are a distinct feature of HoTSoS that allows researchers to present preliminary results prior to formal publication. We saw this collection of papers grow from 2 to 6 this year and anticipate WiP papers becoming a tradition at future HoTSoS conferences. Details on HoTSoS 2020 can be found in Section 3 of this report.

With our partners Syracuse University, University of Minnesota, Case Western Reserve University, and Indiana University we continued executing our NSF Industry University Cooperative Research Center (I/UCRC) planning grant awarded in fall 2019. Over 150 companies were interviewed to gauge interest in participation including all members of the KU SoS Industry Advisory Board. The topic of the Center will be High-Assurance and Secure Systems with numerous researchers and advisory board participants from the KU Lablet. This topic is synergistic with a number of our Lablet research efforts and the I/UCRC presents an excellent technology transfer opportunity. This

successful proposal is a direct result of leveraging our Science of Security effort.

The KU EECS department is beginning our first full academic year supporting a new undergraduate Cyber Security certificate program. Even with Covid uncertainties, all seats are taken and the waiting list is at capacity. We anticipated high demand for the certificate among our Computer Science and Interdisciplinary Computing students. The certificate is not limited to computer science students, but will be difficult to obtain without computer science background. The program integrates hands-on experience through cyber-competitions with traditional classroom learning.

KU is sharing technology with Missouri University of Science and Technology in support of the Kansas City National Security Campus (KCNSC), a major DoE research and development facility. We are integrating our Lablet's remote attestation capabilities with their IoT blockchain to facilitate strong identity and secure information exchange in a model-based manufacturing environment. The resulting system leverages Lablet and will serve as a prototype for new digital protection mechanisms at KCNSC as well as an experimental environment for attestation research.


Dr. Alexander met with Kansas Department of Commerce and Enterprise KC representatives to discuss the KU Lablet and its contribution to the local economy and workforce development. Enterprise KC envisions the greater Kansas City area becoming a commercial hub for cybersecurity research and commercialization. The meetings explored how our Lablet specifically, and KU generally, can contribute this initiative. The result is a strategic plan to be submitted to Kansas Department of Commerce for their consideration.



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, Resilient Architectures		

GOAL

This project investigates security and privacy threats in IoT systems, including IoT devices, the companion mobile applications, the cloud endpoints, and the communication channels. The goal of the project is to model privacy

requirements and threats and develop protection mechanisms to secure cloud-assisted IoT systems. The work aims to address the hard problems of resilient architectures, security metrics as well as scalability and composability

ABSTRACT

The rapid adoption of Internet of Things (IoT) devices and applications also raises many security and privacy concerns. In this year, we have worked on four projects on the security and privacy of IoT devices and applications. First, IoT devices and apps especially the consumer IoT devices collect various types of privacy-sensitive data about the end users in both explicit and implicit forms. It is of great importance to identify the potentially privacy-sensitive data collected by the IoT apps and provide usable app descriptions that correctly describe the behavior of IoT apps and their privacy practices. We studied 319 SmartApps on the Samsung SmartThings cloud platform, which adopts a trigger-action programming paradigm and a capability-based permission model. We conducted a program analysis and a natural language processing (NLP) analysis to extract capabilities and trigger-action dependencies from an app's code and description. By comparing the results, we find that most IoT apps have overly brief functionality descriptions which are too general and non-informative, or inaccurate functionality descriptions with key capabilities or trigger-action dependencies missing.

The second project focuses on protecting privacy in cloud-assisted IoT data analysis, more specifically the federated learning (FL) applications that analyze locally generated IoT data in a decentralized way to avoid uploading potentially privacy-sensitive data to a centralized server. The federated learning approach protects the raw IoT data, however, two security and privacy concerns remain. First, the local model update computed by each IoT device in each round may disclose raw IoT data and therefore needs to be protected using appropriate privacy-preserving techniques that are affordable to the resource-constrained IoT devices. Secondly, the adversary may compromise IoT devices to launch the model poisoning attack, which tampers with the federated learning process and pollutes the global model. To tackle these problems, we developed a

novel blockchain-based privacy-preserving federated transfer learning framework and implemented a proof-of-concept prototype. Our framework adopts a novel homomorphic encryption scheme to encrypt the model updates and store them in a permissioned blockchain, which allows not only the FL server to compute the aggregate of model updates without knowing the individual values but also stores the encrypted model updates in the immutable blockchain for model pollution detection.

Compromised IoT devices can collect sensitive data about the users or inject polluted data to tamper with data analysis. Therefore, in the third project, we studied the IoT device identification problem in which the adversaries connect unauthorized devices to public wireless networks by hiding or forging the device identities. We extracted the unique features in network packets introduced by the subtle differences in the implementations of network modules on different mobile and IoT devices, constructed six views, and developed a novel multi-view classifier to detect abnormal devices, since forged devices often behave inconsistently in different views.

Finally, we worked on IoT device firmware security analysis. Finding bugs in microcontrollers (MCUs) is challenging. Many existing sophisticated software testing tools on x86 cannot be directly used, because the MCU runs different instruction sets than x86 and exposes a very different development environment. To maintain a unified developing and testing environment, a straightforward way is rehosting that re-compiles the source code into the native executable for a commodity machine, however, ad-hoc re-hosting is a daunting and tedious task and subject to many issues such as library dependence, kernel dependence and hardware dependence. To tackle this challenge, we systematically explored the portability problem of MCU software, and proposed a novel para-rehosting approach to ease the porting process.

ACCOMPLISHMENTS

First, we developed an IoT privacy checker, consisting of a program analyzer and an NLP analyzer. The program analyzer utilized the Abstract Syntax Tree (AST) transformation of Groovy to identify the entry point, sink nodes, subscribe statements and functions of the app, from which capabilities and trigger-action dependencies can be extracted. The NLP analyzer utilized the spaCy POS tagger to process the app descriptions and extracted trigger-action phrases and the corresponding capabilities. Finally, the privacy checker can pinpoint the inconsistent capabilities and the corresponding trigger-action paths.

In the second project, we developed a proof-of-concept prototype for the proposed PPFTL framework and ran experiments on a small network with two FL servers and 13 IoT clients, and an Ethereum blockchain. We tested our framework in several federated transfer learning tasks and showed it

achieved a desirable performance with a high prediction accuracy and acceptable delays. The preliminary results of this project have been accepted by the 2020 HotSoS Work-in-Progress session, and the final manuscript is under preparation.

We also implemented the multi-view classifier for device identification and conducted experiments on a data set with over 30,000 device samples. Our approach outperformed existing schemes in accuracy and coverage. The results have been accepted by the 2020 USENIX Security Symposium.

Finally, we developed a para-rehosting prototype on a x86-based PC running the Ubuntu 16.04 OS. We evaluated our approach against nine MCU OSs and successfully compiled and executed over 63% of the libraries. We conducted fuzz testing using our tool to analyze several popular libraries used in MCUs and identified 28 previously-unknown bugs. The results of this project have been accepted by the 2021 NDSS conference.

IMPACT ON HARD PROBLEM

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

- Fengjun Li served on the “Cybersecurity Research and Education” panel in the 2nd International Workshop on Cyber Security and Data Privacy, July 18, 2020.
- Fengjun Li gave an invited talk on “Achieving Accountable Single Sign-on with Ticket Transparency” in the Secure Multiparty Computation Symposium, May 30, 2020.
- Bo Luo gave an invited talk on “Mobile and IoT Device Identification” in the Secure Multiparty Computation Symposium, May 30, 2020.
- Bo Luo gave a guest lecture on “Privacy, a Computer Science perspective” to JOUR790 Social Media Research & Analysis, the University of South Carolina, April 15, 2020.
- Fengjun Li gave a Professional Skill-Building Webinar on “Social Network Security & Privacy: Learning the Truth While Protecting the Sensitive” in the Advisory Boards Meeting of the CEBC Center, University of Kansas, April 6, 2020.
- Fengjun Li gave an invited talk on “Privacy-Preserving Collaborative Learning” in the Frontiers Informatics Meetup: Healthcare Data Analytics and Security, Kansas City, KS, USA, March 5, 2020.
- Bo Luo was invited to visit the Center for Trustworthy IoT Infrastructure at Japan Advanced Institute of Science and Technology and gave a talk on “A First Cut on IoT Security – A Cyber-Physical Perspective” on Feb 6, 2020.

PUBLICATIONS


- Wenqiang Li, Le Guan, Jingqiang Lin, Jiameng Shi, Fengjun Li, “From Library Portability to Para-rehosting: Natively Executing Microcontroller Software on Commodity Hardware,” in NDSS 2021. (accepted).
- Bo Luo, Razvan Beuran, and Yasuo Tan. Smart Grid Security: Attack Modeling from a CPS Perspective. In IEEE ComComAp, December 20-22, 2020.
- Lingjing Yu, Bo Luo, Jun Ma, Zhaoyu Zhou, and Qingyun Liu, “You Are What You Broadcast: Identification of Mobile and IoT Devices from (Public) WiFi,” In *USENIX Security Symposium*, August 12-14, 2020.
- Abdulmalik Humayed, Fengjun Li, Jingqiang Lin, and Bo Luo, CANSentry: Securing CAN-Based Cyber-Physical Systems against Denial and Spoofing Attacks. In *European Symposium on Research in Computer Security (ESORICS)*, September 14-18, 2020.
- Sohaib Kiani, Sana Awan, Fengjun Li, Bo Luo and Jun Huan, “WOLF: Automated Machine Learning Workflow Management Framework for Malware Detection and Other Applications,” in Hot Topics in the Science of Security (HotSoS), Sept. 22-14, 2020.
- Sana Awan, Fengjun Li and Bo Luo, “A Federated Transfer Learning Framework for IoT Applications,” in HotSoS WIP session, September 22-14, 2020.
- Qiang Zhou, Chengliang Tian, Hanlin Zhang, Jia Yu, Fengjun Li, “How to Securely Outsource the Extended Euclidean Algorithm for Large-scale Polynomials over Finite Fields,” in *Information Sciences*, 512, 641-660, 2020.

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 PROBLEM

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

John Symons and his colleagues Ramón Alvarado (University of Oregon) and Kamuran Osmanoglu (Koç University) organized a six-month seminar series on Data Ethics that covers several Science of Security topics in privacy and resiliency. The webinar is free and open to all. Presentations have included:

- Reid Blackman, PhD - What is data ethics and what can it do?
- Tom (Xiaowei) Wang, Renmin University of China - The Confucian and the Californian ideology: contrasting approaches to the development of the internet and their role in the growth of a data economy.

- Ramon Alvarez, University of Oregon and John Symons, University of Kansas - Epistemic Injustice in Data Ethics: identifying the harms particular to data science
- John Symons presented an invited talk, The Metaphysics of Resilience, for the Semanário em Metafísica da Ciência, CFCUL at the University of Lisbon, March 9 2020
- John Symons hosted a talk by Tom Wang from Renmin University on the Confucian Critique of the Open Internet (by skype), Feb11 2020.

PUBLICATIONS

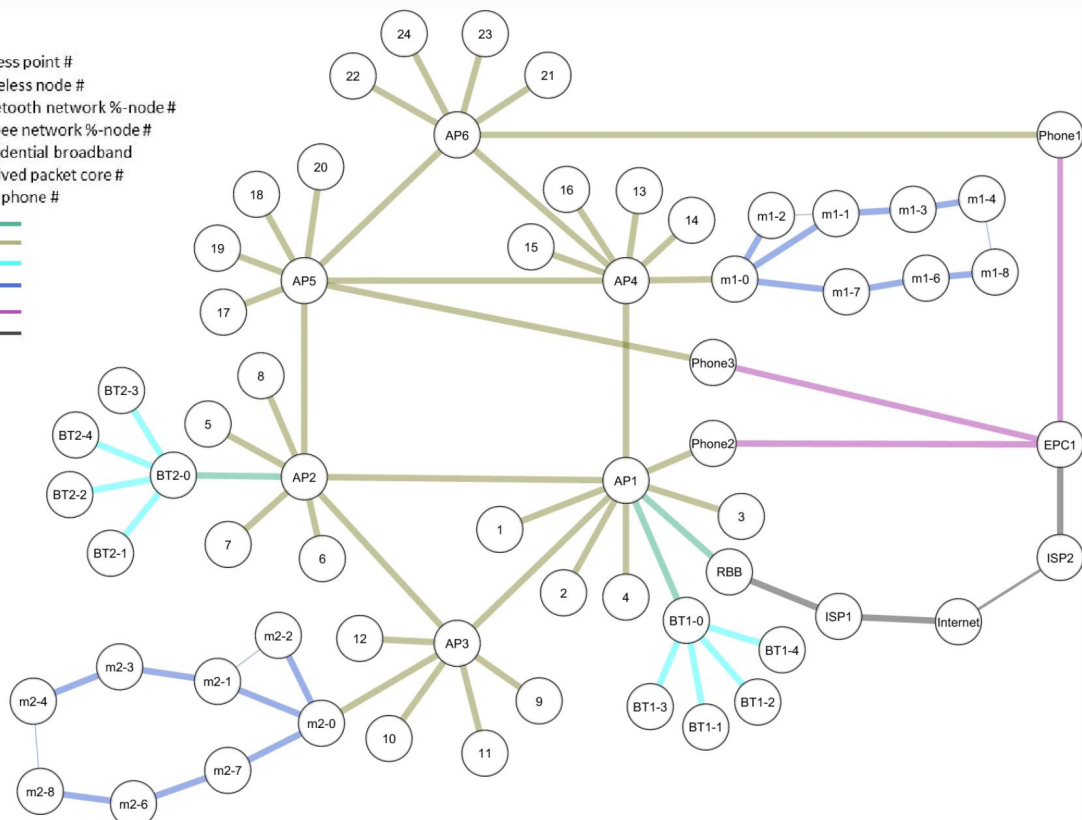
- Amir Modarresi and John Symons, "Technological heterogeneity and path diversity in smart home resilience: A simulation approach," *Procedia Computer Science*, 170, 177-186.
- Amir Modarresi and John Symons, "Resilience and technological diversity in smart homes. A Graph-Theoretic Approach to Modeling IoT Systems with Integrated Heterogeneous Networks," *Journal of Ambient Intelligence and Humanized Computing*, 1-19.
- John Symons, "Meaningfulness and kinds of normative reasons," *Philosophia*, 1-13.
- Jack Horner and John Symons, "Software Engineering Standards for Epidemiological Modeling." arXiv preprint arXiv:2009.09295 Forthcoming in *History and Philosophy of the Life Sciences*.
- Jack Horner and John Symons, "What Have Google's Random Quantum Circuit Simulation Experiments Demonstrated about Quantum Supremacy?" arXiv preprint arXiv:2009.07359.

Fig. 6

From: [Resilience and technological diversity in smart homes](#)

Legends

AP#:	Access point #
#:	Wireless node #
BT%-#:	Bluetooth network %-node #
m%-#:	Zigbee network %-node #
RBB:	Residential broadband
EPC#:	Evolved packet core #
Phone#:	Cell phone #
LAN	
WLAN	
BT	
ZB	
LTE	
WAN	



An instance of a smart home with six access points

Scalable Trust Semantics and Infrastructure





Perry Alexander



Garrett Morris

PI:	Perry Alexander	CO-PI: Garrett Morris (through mid-July 2020)
URL:	https://cps-vo.org/KU_ScalableTrustSemanticsandInfrastructure	
HARD PROBLEM:	Scalability and Composability, Policy-Governed Secure Collaboration	



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 and appraises evidence 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.

Progress on formalizing the Attestation Monad resulted in significant changes in the overall architecture. Each place in a Copland phrase has consistently been defined as an attestation manager modeled as an Attestation Monad. As we began modeling attestation patterns and examples, we realized that a place is actually a collection of such monads for negotiation, attestation, and appraisal. Thus, each attestation manager is a collection of these negotiation, attestation, and appraisal services. Decomposing definition of place and attestation manager in this way maximizes flexibility allowing an attestation manager to only perform a subset of attestation and appraisal operations or use a collection of Attestation Monads to provide flexible infrastructure. Initial models of attestation patterns suggest this approach has great promise in achieving our goal of scalable trust. We extended the attestation manager design to include late launch via key release where attestation manager boot releases keys as trust is established.

KU continued the development and verification of the

Attestation Virtual Machine and a compiler from Copland to the AVM. The basic compiler is largely complete and verified in Coq.

Working with MITRE, JHUAPL and NSA we have identified an initial set of attestation architecture constructs that include: a) delegated appraisal; b) cached; c) layered; and d) mutual attestation. We have been working with a set of specific architectures and are moving towards a collection of architecture building-blocks that assemble to define architectures. For example, a layered attestation involves an appraiser delegating appraisal tasks to its subsystems. A cached attestation involves evidence gathered prior to the attestation request. Composing layered and cached attestation results in an architecture where subcomponents update an attestation agent rather than waiting for requests. A layered mutual attestation allows for several mutual attestation steps to build a full attestation result. We are working to identify and specify a set of such architecture constructions that allow defining attestation systems over large systems.

We made significant progress on attestation protocol negotiation and policies for defining negotiation and selection. The goal is assuring that a best protocol is negotiated for both appraiser and target that satisfies their local policy. We began defining a lattice of protocols where the lattice partial ordering defines preference for the appraiser or target. Given an attestation request, the best protocol is the local maxima of the intersection of the sublattices formed by those protocols that

satisfy the attestation request while respecting local policy. This is important because it captures formally what “best” means providing a verification condition as we move forward to a more complete model.

Our Science of Security research contributes to the CakeML/seL4 implementation being co-developed for DARPA and AFRL. We continued to refine the implementation and made progress towards verification. We will be using this implementation to demonstrate our attestation patterns during the remainder of the calendar year. attestation request, the best protocol is the local maxima of the intersection of the sublattices formed

by those protocols that satisfy the attestation request while respecting local policy. This is important because it captures formally what “best” means providing a verification condition as we move forward to a more complete model.

Our Science of Security research contributes to the CakeML/seL4 implementation being co-developed for DARPA and AFRL. We continued to refine the implementation and made progress towards verification. We will be using this implementation to demonstrate our attestation patterns during the remainder of the calendar year.

IMPACT ON HARD PROBLEM

- 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

- Technology exchange with Missouri University of Science & Technology in support of the Kansas City National Security Campus (KCNSC), a DoE research and development center. We are prototyping IoT trust infrastructure in support of their model-based manufacturing center.
- Poster presentation, “An Infrastructure for Faithful Execution of Remote Attestation Protocols”, Adam Petz, Hot Topics in Science of Security (HoTSoS 2020), September 22-23, 2020.
- Poster presentation, “An seL4-based Architecture for Layered Attestation,” Grant Jurgensen, Adam Petz, Michael Neises, and Perry Alexander, Hot Topics in the Science of Security (HoTSoS 2020), September 22-24, 2020.

PUBLICATIONS

- Multiple papers have been submitted for publication.

An Infrastructure for Faithful Execution of Remote Attestation Protocols

Adam Petz
The University of Kansas - ITTC

Objective: Design, implement, and prove correct a collection of software components that provide a sound infrastructure for remote attestation of layered systems.

Copland Language and Reference Semantics [1,2]

- Copland Phrases
 - System Measurements (local and cross-domain)
 - Cryptographic bundling of evidence
 - Remote Requests
 - Measurement Sequencing
- Copland Evidence
 - Precise cryptographic structure
 - Concrete measurement values
- Reference Semantics
 - Characterization of attestation-relevant system events
 - Evidence Shape
 - Ideal for comparing protocol alternatives [3]

Copland Compiler + Attestation Virtual Machine

- Copland Compiler
 - Phrase → Attestation Instructions
 - Maps abstract measurement specifications and cryptographic operations to concrete services
- Attestation Virtual Machine (AVM)
 - Attestation instructions → Evidence
 - Functional program in monadic style
 - AVM Monad
 - Invokes attestation services (measurements + crypto)
 - Principled updates to evidence bundle
 - Protects evidence (tampering, disclosure)

The diagram illustrates the attestation process. It starts with 'System Events' from both 'Appraiser' and 'Target'. These events are processed through 'Event Correspondence' and 'AM Semantics' to produce 'Evidence Type'. This evidence is then processed by 'Appraisal Term Synthesis' to generate 'Copland Evidence'. This evidence is fed into the 'Attestation VM (AVM)', which also receives 'Attestation Instructions' from the 'Copland Compiler'. The AVM produces 'Evidence' which is then processed by the 'Attestation Manager' to reach a 'Trust Decision'. The 'Attestation Manager' also interacts with the 'Copland Compiler' and 'Attestation VM (AVM)'.

Attestation Manager Monad + Appraisal Term Synthesis

- AM Monad Environment
 - Nonce generation
 - Composing evidence from multiple Copland phrase runs
- Appraisal Configuration
 - Golden measurement values
 - Public keys
 - Mapping from measurement to appraisal routines
- Appraisal Term Synthesis
 - Attestation phrase + Evidence → Appraisal phrase
 - Leverages existing Copland Compiler + AVM
 - Less error-prone than manually constructing appraisal routines per-protocol

Formal Verification

- Evidence Semantics (Completed)
 - Shape of AVM-produced evidence respects Copland ref. semantics
- System Event Correspondence (Nearly Complete)
 - AVM respects event orderings of Copland reference semantics
- Appraisal Completeness and Soundness (Ongoing)
 - Every part of the evidence is appraised
 - What does a successful appraisal say about the target platform (and its configuration)?

[1] J. D. Ramsdell, P. D. Rowe, P. Alexander, S. C. Helble, P. Losocco, A. J. Pendergrass, and A. Petz. “Orchestrating Layered Attestations”. *POST 2019*, 2019.

[2] A. Petz and P. Alexander. “A Copland Attestation Manager”. *HotSoS 2019*, 2019.


[3] P. D. Rowe. “Confining adversary actions via measurement”. *Third International Workshop on Graphical Models for Security*, pages 150–166, 2016.

Secure Native Binary Execution



Prasad Kulkarni

PI:	Prasad Kulkarni
SUB-LABELT	University of Tennessee
URL:	https://cps-vo.org/KU_SecureNativeBinaryExecution
HARD PROBLEM:	Security 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 binary 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

There are many avenues for developers to harden their software against security threats, including a) using secure programming languages or constructs, b) using recommended best coding practices and mechanisms, c) manually inserting programmer identified security checks in the source code, and d) inserting compiler provided security checks during code generation. Our current focus is to identify the presence of compiler added security checks in any given binary code.

Our approach uses two unique insights, firstly that compiler checks are added uniformly to the code, and secondly that compiler-added security checks perform tasks that are orthogonal to the primary function of the program. For instance, the Stackguard technique inserts code at the start of most/all functions to insert a canary on the stack, and then, before the "return" instruction, to test its integrity. Likewise, Control-Flow Integrity (CFI) checks are added before most/all indirect calls or jumps. These code sequences appear disconnected in terms of control-flow and data-flow from the other parts of the code that compute the main algorithm of the program.

Our technique conducts the following steps to identify compiler-inserted security checks in the binary: a) employ Ghidra to detect interesting instructions, like returns, indirect calls/jumps, loads/stores, etc.; b) fetch and dump disassembled code from predecessor and successor blocks around the interesting instructions; c) process instruction trace to normalize constants, register numbers, labels, etc.; and d) find common instruction patterns across collected traces. The presence of common instruction sequences/patterns across traces indicates the likelihood of compiler checks to protect against attacks related to that code construct.

Our immediate next steps include developing logic to confirm the data- and control-flow disconnect of the security check code and to better reason that the common subsequences indicate the presence of explicit security mechanisms in the binary code.

We also continued our work to develop a hybrid framework to detect and prevent memory attacks on unmodified client-side binaries. We have developed one of the very few decoupled

binary-level techniques that offer complete memory safety for binary programs. Our current implementation is the only one that uses static analysis to determine relevant program information. We assessed reasons when even the complete availability of debug symbol information is not sufficient to replace the semantic information lost during program translation and prevent buffer overflows at run-time. We evaluated the effectiveness of (our) static analysis-based techniques for binaries that are stripped of debug information. We assessed whether advanced static reverse engineering and type inference algorithms can regenerate or predict symbol information with sufficient accuracy to improve effectiveness of memory protection techniques for stripped binaries.

The technique we implemented is inspired by SoftBound, which is a compiler-level mechanism to detect and prevent all spatial and temporal memory errors during program execution. With our binary-level implementation the input binary is statically analyzed using the Ghidra and IDA Pro reverse engineering

framework, which outputs information regarding the buffer bounds and the type referenced by each memory access (read/write) and pointer assignment instruction (called the owner) in the binary. At run-time, we employ the Pin virtual machine to keep track of owner information and check relevant buffer reads/writes to ensure fine-grained memory safety. We employ our framework to assess the effectiveness of binary-level memory safety techniques in different situations.

Dr. Michael Jantz and his team at the University of Tennessee have continued work on a static analysis and binary rewriting tool that inserts checks for spatial memory vulnerabilities during program startup. They have developed, tested, and validated their tool with multiple benchmarks (including xz and leela from SPEC CPU 2017 and several test cases from the SARD benchmark suite). They are currently working to test their tool with other benchmarks from SPEC CPU 2017 and are beginning studies to compare this tool with existing approaches.

IMPACT ON HARD PROBLEM

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.

PUBLICATIONS

- A paper based on our research to detect and prevent memory attacks on binaries is ready for submission.



Problem

- *Client* wants to run some binary software
 - How secure is the software?
 - Does it incorporate language/compiler level security checks?
 - How to update software to enforce desired *level* of security from attacks?
 - and maintain tolerable level of performance
- Current technology
 - Up to the software developer to secure binary
 - No feedback to client about security level of distributed software
 - Not easy for the client to add security to distributed binary software
 - *Trust* the software developer/distributor


Goal: Develop high-performance framework for client-side security assessment and enforcement for COTS binary software

Side-Channel Attack Resilience

PI: Heechul Yun

URL: https://cps-vo.org/KU_Side-ChannelAttackResistance

HARD PROBLEM: Resilient Architectures



Heechul Yun

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

We have continued to develop effective micro-architectural attacks on commercial off-the-shelf computing platforms. We discovered a new contention-based covert channel on Intel and AMD processors, and we are evaluating viable attack scenarios with the covert channel. We also developed an improved denial-of-service attack methodology on shared cache in embedded multicore processors, which significantly outperform the state-of-the-art. We ported and evaluated our cache partitioning and memory bandwidth regulation software solutions on Jetson Nano, Jetson Xavier and RPi4 platforms. We developed optimal and heuristic scheduling algorithms for our OS-level gang scheduling framework.

We continue to develop software and hardware level defense mechanisms to mitigate micro-architectural covert/side-channel attacks. We developed a holistic real-time scheduling framework in Linux that integrates a novel OS level scheduler and state-of-the-art resource isolation mechanisms to limit or eliminate micro-architectural attacks. We developed a light-weight bandwidth regulator IP on a RISC-V SoC.

We have continued to develop hardware-based memory performance attack protection mechanisms. Specifically, we are extending our prior BRU work to better support both cache and memory bandwidth regulation to protect critical applications from potential micro-architectural performance attacks.

IMPACT ON HARD PROBLEM

This project is developing OS and architecture techniques to defend against potential microarchitectural side-channel attacks on embedded computing platforms for safety-critical

systems. The project covers the hardware problems of (1) resilient architectures (primary) and (2) security-metrics-driven evaluation, design, development and deployment.

EDUCATION AND OUTREACH

- Heechul Yun served on the program committees for ACM LCTES 2020, IEEE RTSS 2020, ACM/IEEE DAC 2020, and IEEE RTAS 2020.
- Heechul Yun is serving as editor ACM SIGBED Blog.
- Heechul Yun was invited to serve on the program committees of IEEE RTAS 2021, ACM/IEEE DAC 2021, ECRTS 2021.
- Poster presentation, "Exploiting DRAM Bank Mapping and HugePages for Effective Denial-of-Service Attacks on Shared Cache in Multicore" Michael Bechtel and Heechul Yun, Hot Topics in the Science of Security (HotSoS), September, 2020.

PUBLICATIONS

- Jacob Fustos, Michael Bechtel, and Heechul Yun, "SpectreRewind: Leaking Secrets to Past Instructions," *Workshop on Attacks and Solutions in Hardware Security (ASHES)*, November 2020.
- Farzad Farshchi, Qijing Huang, and Heechul Yun, "BRU: Bandwidth Regulation Unit for Real-Time Multicore Processors," *IEEE International Conference on Real-Time and Embedded Technology and Applications Symposium (RTAS)*, April 2020.
- Waqar Ali, Rodolfo Pellizzoni, and Heechul Yun, "Virtual Gang based Scheduling of Parallel Real-Time Tasks," *Design, Automation and Test in Europe Conference (DATE)*, February 2021. (to appear)
- Michael Bechtel and Heechul Yun, "Memory-Aware Denial-of-Service Attacks on Shared Cache in Multicore Real-Time Systems," arXiv preprint, arXiv:2005.10864, 2020.





Fig. 3. A simplified view of a multicore processor with shared memory resources. BRU regulates per-core bandwidth at source.

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, University of California, Berkeley, and University of Texas at Dallas 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.



Xenofon Koutsoukos

Despite the pandemic, the VU Lablet participated in multiple online conferences and workshops and had multiple research papers published.

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

- Policy 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




Policy Analytics for Cybersecurity of Cyber-Physical Systems



Nazli Choucri

PI:	Nazli Choucri
URL:	https://cps-vo.org/VU_AnalyticsforCPS-Cybersecurity
SUB-LABELT:	Massachusetts Institute of Technology
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, our goal is to provide analytics for cybersecurity policies and guidelines designed specifically to (a) overcome the limitations of the conventional text-based form, (b) extract knowledge embedded in policy guidelines,

and (c) assist the user community –analysts and operators – in implementation. Strategically, our goal is to construct a platform of new tools for application to policy directives, regulations, and guidelines across diverse domains and issue areas. The platform, and tools, are designed to enable users to explore mission-related system properties, issues, concerns, or contingencies

ABSTRACT

Cyber-Physical Systems (CPS) are embedded in an increasingly complex ecosystem of cybersecurity policies, guidelines, and compliance measures designed to support all aspects of operation during all phases of system's life cycle. By definition, such guidelines and policies are written in linear and sequential text form – word after word – often with different parts presented in different documents. This situation makes it difficult to integrate or understand policy-technology-security interactions. As a result, it also impedes effective risk assessment. Individually or collectively, these features inevitably undermine initiatives for cybersecurity. Missing are fundamental policy analytics to support CPS cybersecurity and facilitate policy implementation. This project is designed to develop a set of text-to-analytics methods and tools, with a "proof of concept" focused on the smart grid of electric power systems. The challenge is to develop a structured system model from text-based policy guidelines and directives in order to (a) identify major policy-defined system-wide parameters, (b) situate vulnerabilities and

impacts, (c) map security requirements to security objectives; and (d) advance research on the responses of multiple system features to diverse policy controls – all of which are necessary to strengthen the fundamentals of cybersecurity for cyber-physical systems.

Our "raw" data base consists of major reports prepared by the National Institute for Standards and Technology (NIST). Clearly, considerable efforts are always being made to "mine" NIST materials; however, few initiatives explore the potential value-added of drawing on multi-methods for knowledge extraction and/or of developing analytical tools to support user understanding of policy directives, analysis, and eventually to enable action. While our approach appreciates and is informed by such efforts, it transcends them by developing a platform for multi-method cybersecurity policy analytics – based entirely on the contents of policy documents. The case application, as "proof of concept," focuses on cybersecurity of the smart grid for electric power systems.

ACCOMPLISHMENTS

Thus far, we have aligned the project vision and mission to the National Cybersecurity Policy and identified the overall policy-relevant ecosystem. By focusing on national cybersecurity policies for securing the nation's critical infrastructure, we identified the core policy documents regarding smart grid CPS for the overall research design. We have extracted data from these documents and created a metric-based Dependency Structure Matrix (DSM) of the "as-is" smart grid NIST reference model. We also completed the design and operational strategy

for our data extraction and linkage method. This involves developing the method for moving *from* "policy-as-text" *to* "text-as-data". More specifically, we created a metric-based system-representation of policy documents and guidelines, and then captured the *value-added* in metric form. This enabled us to pre-test the operational framework for the next step.

We developed rules and methods for extracting data from key documents and constructed the necessary issue and

policy specific linked database. Once completed, this step allowed us to create (a) the initial exploratory tools for analysis of system information, and (b) a core DSM of the CPS based on the identification of first-level information dependencies. The dependency matrix must then be examined closely and validated and further transformed as needed into clusters and partitions of structure and process in order to explore system properties that reveal interconnections and “hidden features.” It is also the basis upon which added policy imperatives – also in text form – are incorporated, at a later stage in expanded DSM forms.

Throughout this process, we have also addressed critical research steps, including: a) highlighting potential correctives; b) replicating the core structured DSM model; c) extensions of the core DSM; d) general applications for a new method;

e) validation of initial structure model; and f) approaches to automation of the human research steps.

The following websites are all related to this research:

cyberpolitics.mit.edu (under development) focuses on the research work conducted under this research grant and its extensions based on discussion with NSA leadership and of relevance to US National Security and Cybersecurity.

cyberIRworld.mit.edu (operational) is a dynamic, interactive ontology-based knowledge system focused on the evolving, diverse and complex interconnections of cyberspace and international relations.

ecir.mit.edu (operational) provides the information on the foundations of the project that were completed under a separate DoD support via Minerva Project (2009-2014)

IMPACT ON HARD PROBLEM

Our major contribution to the specific hard problem we examine is the value of “text-as-data” in a complex cyber-physical system 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 cybersecurity directives and guidelines directly applicable to the system in question; (b) information

about relative vulnerability pathways throughout the whole or parts of the system-network, as delineated by the guidelines documents; (c) insights from contingency investigations, that is, “what...if...”; (d) design framework for information management within the organization; and (e) ways to facilitate information flows bearing on decision-making for cybersecurity.

EDUCATION AND OUTREACH

- Nazli Choucri, “Analytics for Cybersecurity of Cyber-Physical Systems.” Presented at *2020 Winter Science of Security and Privacy Quarterly Meeting*, Raleigh, North Carolina, January 15–16, 2020.
- At MIT, we began to explore logistics of organizing a summer program for senior SoS leadership on the topic of Cybersecurity in collaboration with MIT Professional Education. Due to COVID-19, we are working on a special online class on Cybersecurity. The class will be offered first to MIT students, and if requested, SoS team members will be invited to audit the class. The new course focuses

on multidisciplinary approaches to sources, types, and impacts of cyber threats. Attention is given to operations and solution strategies in technical, economic, political and strategic contexts.

- We presented the results of our research to date to the industry partner of the MIT Sloan School Cybersecurity CAMS Program addressing the problems that the text form creates for implementation policies and guidelines. We reviewed the research design and illustrated ways in which problems of policy-as-text can be addressed in order to facilitate policy implementation.

PUBLICATIONS


- Thomas Klemas, Sean Atkins, Nazli Choucri, and Rebecca Lively, “Accelerating Cyber Acquisitions: Introducing a Time-Driven Approach to Manage Risk with Less Delay.” *Research Paper No. 2020-7*. Massachusetts Institute of Technology: Political Science Department. <http://dx.doi.org/10.2139/ssrn.3703183>
- Keman Huang, Stuart Madnick, and Nazli Choucri, “Building Global Digital Supply Chain Hub by Cybersecurity Commitment: Singapore’s Strategic Role in the Digital Age,” *Global Policy*

Foundations of 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

Integrated moving target defense and control reconfiguration for securing CPS

With the increasingly connected nature of Cyber-Physical Systems (CPS), new attack vectors are emerging that were previously not considered in the design process. Moving Target Defense (MTD) techniques such as Instruction Set Randomization (ISR), and Address Space Randomization (ASR) have been shown to be effective against code injection and code reuse attacks. We addressed the problem of maintaining system and security properties of a CPS under attack by integrating MTD techniques, as well as detection, and recovery mechanisms to ensure safe, reliable, and predictable system operation. By using MTD such as ISR, and ASR, our approach provides the advantage of preventing attackers from obtaining the reconnaissance knowledge necessary to perform code injection and code reuse attacks, making sure attackers can't find vulnerabilities in the first place.

The rowhammer bug belongs to software-induced hardware faults, and has been exploited to form a wide range of powerful rowhammer attacks. Yet, how to effectively detect such attacks

remains a challenging problem. We propose a novel approach named RADAR (Rowhammer Attack Detection via A Radio) that leverages certain electromagnetic (EM) signals to detect rowhammer attacks. Using a common classification technique, we can achieve both effective and robust detection-based defense against rowhammer attacks, as evaluated on a RADAR prototype under various scenarios. In addition, our RADAR does not impose any performance overhead on the protected system. There has been little prior work that uses physical side-channel information to perform rowhammer defenses, and to the best of our knowledge, this is the first investigation on leveraging EM side-channel information for this purpose.

We studied the resilient diffusion problem in a network of robots aiming to perform a task by optimizing a global cost function in a cooperative manner. We propose a resilient aggregation rule based on the notion of centerpoint, which is a generalization of the median in the higher dimensional Euclidean space. We also showed that commonly used aggregation rules based on the coordinate-wise median and geometric median are, in fact, not resilient to certain attacks.

IMPACT ON HARD PROBLEM

CPS 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 attack

EDUCATION AND OUTREACH

- We presented our research in the “Anomaly Detection of Cyber-Physical Systems (ADCPS)” team meeting, USNA, January 29-30, 2020.
- Keynote talk, Systems Science of Secure and Resilient Cyber-physical Systems, International Conference on Contemporary Computing and Applications, (IC3A 2020), Lucknow, India, February 5-7, 2020
- Poster presentation, “Resilient Multi-Robot Target Pursuit,” by Jiani Li, Waseem Abbas, Mudassir Shabbir, and Xenofon Koutsoukos, Hot Topics in the Science of Security (HotSoS 2020), September 22-24, 2020v

PUBLICATIONS

- Jiani Li, Waseem Abbas, and Xenofon Koutsoukos, “Byzantine Resilient Distributed Multi-Task Learning,” *Thirty-fourth Annual Conference on Neural Information Processing Systems (NeurIPS 2020)*. Virtual conference. Dec 6-12, 2020.
- Mudassir Shabbir, Jiani Li, Waseem Abbas, and Xenofon Koutsoukos, “Resilient Vector Consensus in Multi-Agent Networks Using Centerpoints,” *American Control Conference 2020*, July, 1-3, 2020.
- Waseem Abbas, Mudassir Shabbir, Hassan Jaleel, and Xenofon Koutsoukos, “Improving Network Robustness through Edge Augmentation While Preserving Strong Structural Controllability,” *American Control Conference 2020*, July, 1-3, 2020.
- Jiani Li, Waseem Abbas, Mudassir Shabbir, and Xenofon Koutsoukos, “Resilient Distributed Diffusion for Multi-Robot Systems Using Centerpoint,” *Robotics: Science and Systems*, July 12-16, 2020
- Jiani Li, Waseem Abbas, and Xenofon Koutsoukos, “Resilient Distributed Diffusion in Networks with Adversaries. *IEEE Transactions on Signal and Information Processing over Networks*, vol. 6, pp. 1-17, 2020.
- Feiyang Cai and Xenofon Koutsoukos. “Real-time Out-of-distribution Detection in Learning-Enabled Cyber-Physical Systems,” *11th IEEE/ACM Conference on Cyber-Physical Systems (ICCPs’20)*, Sydney, Australia, April 22-24, 2020. (Best Paper Award Finalist)
- Carlos Barreto, Himanshu Neema, and Xenofon Koutsoukos, “Attacking Electricity Markets Through IoT,” *IEEE Computer*, Special Issue on Cybersecurity for the Smart Grid, vol. 53, no. 5, pp. 55-62, May 2020.
- Zhenkai Zhang, Zihao Zhan, Daniel Balasubramanian, Bo Li, Peter Volgyesi, Xenofon Koutsoukos, “Leveraging EM Side-Channel Information to Detect Rowhammer Attacks,” *IEEE Symposium on Security and Privacy (S&P 2020)*, pp. 862-879. May 18-20, 2020.
- Feiyang Cai, Jiani Li, and Xenofon Koutsoukos, “Detecting Adversarial Examples in Learning-Enabled Cyber-Physical Systems using Variational Autoencoder for Regression,” *Workshop on Assured Autonomous Systems (WAAS 2020)*, Held in conjunction with *IEEE S&P*, May 21, 2020.
- Dimitrios Boursinos and Xenofon Koutsoukos, “Trusted Confidence Bounds for Learning Enabled Cyber-Physical Systems,” *Workshop on Assured Autonomous Systems (WAAS 2020)*, Held in conjunction with *IEEE S&P*, May 21, 2020. (Best Paper Award)
- Bradley Potteiger, Feiyang Cai, Abhishek Dubey, Zhenkai Zhang, and Xenofon Koutsoukos, “Security in Mixed Time and Event Triggered Cyber-Physical Systems using Moving Target Defense,” *IEEE International Symposium On Real-Time Distributed Computing 2020 (ISORC 2020)*, May 19-21, 2020. (Best Paper Award Nomination)
- Aron Laszka, Waseem Abbas, Yevgeniy Vorobeychik, and Xenofon Koutsoukos, “Integrating redundancy, diversity, and hardening to improve security of industrial internet of things,” *Cyber-Physical Systems*, 6:1, 1-32, 2020.
- Saqib Hasan, Abhishek Dubey, Gabor Karsai, and Xenofon Koutsoukos, “A game-theoretic approach for power systems defense against dynamic cyber-attacks,” *International Journal of Electrical Power & Energy Systems*, Volume 115, 2020.
- Bradley Potteiger, Zhenkai Zhang, and Xenofon Koutsoukos, “Integrated moving target defense and control reconfiguration for securing Cyber-Physical systems,” *Microprocessors and Microsystems*, Volume 73, 2020.
- Xenofon Koutsoukos, “Systems Science of Secure and Resilient Cyber-physical Systems,” *Computer*, vol. 53, no. 3, pp. 57-61, March 2020.

Mixed Initiative and Collaborative Learning in Adversarial Environments




Claire Tomlin



Shankar Sastry

PI:	Claire Tomlin CO-PI: Shankar Sastry
URL:	https://cps-vo.org/VU_MixedInitiativeCollaborativeLearning
SUB-LABEL:	University of California, Berkeley
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.

ACCOMPLISHMENTS

In many settings where multiple agents interact, the optimal choices for each agent depend heavily on the choices of the others. These coupled interactions are well-described by a general-sum differential game, in which players have differing objectives, the state evolves in continuous time, and optimal play is characterized by Nash equilibria. Often, problems admit multiple Nash equilibria. From the perspective of a single agent in such a game, this multiplicity of solutions can introduce uncertainty about how other agents will behave. We propose a general framework for resolving ambiguity between Nash equilibria by reasoning about the equilibrium other agents are aiming for. We demonstrate this framework in simulations of a multi-player human-robot navigation problem that yields two main conclusions: First, by inferring which equilibrium humans are operating at, the robot is able to predict trajectories more accurately; and second, by discovering and aligning itself to this equilibrium the robot is able to reduce the cost for all players.

Many problems in robotics involve multiple decision-making agents. To operate efficiently in such settings, a robot must reason about the impact of its decisions on the behavior of other agents. Differential games offer an expressive theoretical framework for formulating these types of multi-agent problems. Unfortunately, most numerical solution techniques scale poorly with state dimension and are rarely used in real-time applications. For this reason, it is common to predict the future

decisions of other agents and solve the resulting decoupled, i.e., single-agent, optimal control problem. This decoupling neglects the underlying interactive nature of the problem; however, efficient solution techniques do exist for broad classes of optimal control problems. We take inspiration from one such technique, the Iterative Linear-Quadratic Regulator (ILQR), which solves repeated approximations with linear dynamics and quadratic costs. Similarly, our proposed algorithm solves repeated linear-quadratic games. We experimentally benchmark our algorithm in several examples with a variety of initial conditions and show that the resulting strategies exhibit complex interactive behavior. Our results indicate that our algorithm converges reliably and runs in real-time. In a three-player, 14-state simulated intersection problem, our algorithm initially converges in <0.25s. Receding horizon invocations converge in <50 ms in a hardware collision-avoidance test.

Partially Observable Markov Decision Processes (POMDPs) with continuous state and observation spaces have powerful flexibility for representing real-world decision and control problems but are notoriously difficult to solve. While recent online sampling-based algorithms that use observation likelihood weighting have shown unprecedented effectiveness in domains with continuous observation spaces, there has been no formal theoretical justification for this technique. This research offers such a justification, proving that a simplified algorithm, Partially

Observable Weighted Sparse Sampling (POWSS), will estimate Q-values accurately with high probability and can be made to perform arbitrarily near the optimal solution by increasing computational power.

When the pandemic hit, we began thinking about the resilience of CPS systems to attack. While there has been a great deal of SoS work on how to have cyber systems operate through attack, there has been almost no work on what it takes to restart a shut-down societal system. However, a large part of resilience is the ability to restart. Although we do not as yet have results on the main topic of resilience that we are working on, the highlight

thus far has been a development of an understanding of the linking between infection models for disease spread (referred to as SEIR models (Susceptible, Exposed, Infected, Recovered models), social distancing, partial shut downs, contact tracing with models of economic activity. We are formulating models of optimal decision making to open up sectors of the economy while keeping acceptable bounds on disease. The methods being developed involve a complex mixture of AI/ML applied to large data sets which are publicly available and model parameter estimation. The techniques are ones which will have widespread applicability to other classes of networks under cyber (or natural) attacks.

IMPACT ON HARD PROBLEM

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

Claire Tomlin ran the 6th installment of Berkeley Girls in Engineering (GiE), a program held at UC Berkeley for middle school students, in Summer 2020. The program has traditionally run for 4 weeks, with 30 students participating per week, but this year, it was offered in a virtual format; in addition to providing a kit, we loaned a chromebook and wifi hotspot access to each participant.

Shankar Sastry launched a new Institute entitled the C3 Digital Transformation Institute (<https://c3dti.ai>) a partnership of Berkeley and UIUC (co-leads) along with the University of Chicago, CMU, MIT, Princeton, and Stanford, to develop the science and technology of Digital Transformation. The

private philanthropy that supports this institute was very much leveraged on the support of Federal research such as this SoS lablet.

We developed a new course in systems theory at Berkeley for upper-level undergraduates and first and second year graduate students, on a rapprochement between control theory and reinforcement learning. The course focused 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 was taught by Shankar Sastry in Spring 2020 (ending in May 2020). This course was notable for the rich work it featured in multi-agent systems.

PUBLICATIONS

- David Fridovich-Keil, Ellis Ratner, Lasse Peters, Anca Dragan, and Claire Tomlin, "Efficient iterative linear-quadratic approximations for nonlinear multi-player general-sum differential games," *International Conference on Robotics and Automation (ICRA)*, 2020.
- Lasse Peters, David Fridovich-Keil, Claire Tomlin, and Zachary Sunberg, "Inference-based strategy alignment for general-sum differential games," *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2020.
- Michael Lim, Claire Tomlin, and Zachary Sunberg, "Sparse tree search optimality guarantees in POMDPs with continuous observation spaces," *29th International Joint Conference on Artificial Intelligence and the 17th Pacific Rim International Conference on Artificial Intelligence (IJCAI-PRICAI 2020)*.
- Somil Bansal, Andrea Bajcsy, Ellis Ratner, Anca Dragan, and Claire Tomlin, "A Hamilton-Jacobi Reachability-Based Framework for Predicting and Analyzing Human Motion for Safe Planning," *International Conference on Robotics and Automation (ICRA)*, 2020.

Multi-Model Test Bed 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 Metrics 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

In order to provide a richer developer/analysis environment within the developed testbed, we have been working on a Jupyter Notebook integration capability in WebGME, which would allow Python-based data exploration and/or model modifications to be implemented. We completed the first working prototype of this feature. The current notebook-based workflow and the key architectural elements are as follows: (1) A (Javascript-based) WebGME plugin can programmatically generate Jupyter Notebooks, based on the contents of the model; (2) The Jupyter Notebook server (co-hosted with the WebGME server) is accessed with a simple iFrame-based visualizer inside the WebGME interface; (3) With generated notebook code developers can implement custom analysis algorithms and may send data back to WebGME (modify the model); and (4) The model modification is supported by another (Javascript-based) plugin that has direct access to the model.

We studied the vulnerabilities of protection systems that can detect cyber-attacks in power grid systems and showed that machine learning-based discriminators are not resilient against

Denial-of-Service (DoS) attacks. We demonstrated that an adversarial actor can launch DoS attacks on specific sensors, render their measurements useless and cause the attack detector to classify a more sophisticated cyber-attack as a normal event. We proposed a defense model that improves the resilience of neural networks against DoS through adversarial training.

Experimentation with Blockchain-based Transactive Energy Systems (TES) has been an ongoing research effort, involving the design and modeling of such systems, developing a novel class of attacks against TES, analyzing the system-level effects and potential mitigation strategies. To support this work, an experimentation platform has been built by integrating GridLAB-D - which simulates the power generation, distribution, and consumption aspects of the smart-grid - with TRANSAX, an in-house developed open-source framework for solving and executing market bids and contracts in a Blockchain-based decentralized energy exchange.

IMPACT ON HARD PROBLEM

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

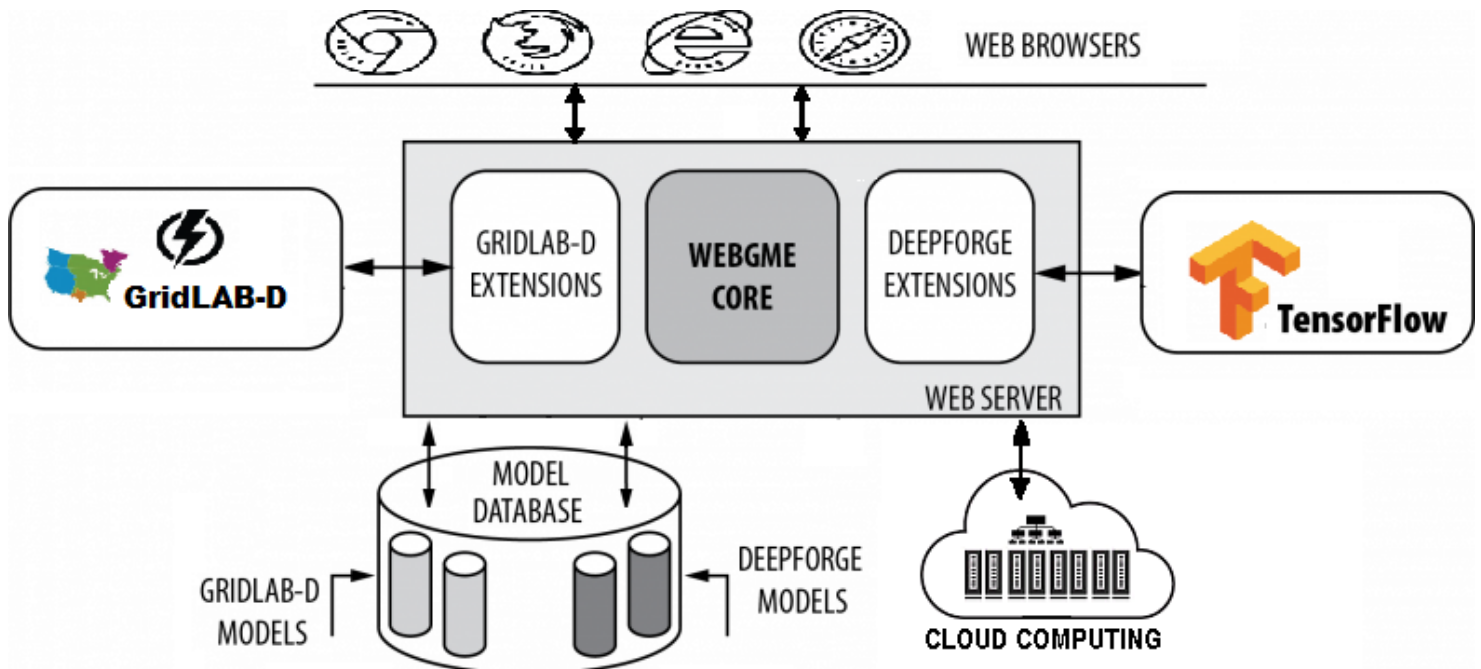
cloned for testing alternative algorithms or parameters. The current modeling strategy – which allows for custom Python-based processing blocks to be implemented through the web interface – provides a more open and pragmatic approach for experienced researchers to develop re-usable components for the 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 held a virtual technology workshop with them in August. Discussion topics included the Fujitsu Cyber Range Builder, Fujitsu Cyber Threat Intelligence Tools, MITRE CALDERA* for Cyber exercise coupled to Cyber Range Builder, and Vanderbilt WebGME and Deepforge integration with Fujitsu security testbeds
- We are collaborating with NIST on threat modeling and risk analysis in ICS. Discussion topics included: Threat modeling in Railway ICS; Risk Analysis; Quantitative Risk Evaluation; and Integration with Simulation-Based Evaluation.

PUBLICATIONS

- Scott Eisele, Carlos Barreto, Abhishek Dubey, Xenofon Koutsoukos, Taha Eghtesad, Aron Laszka, and Anastasia Mavridou, “Blockchains for Transactive Energy Systems: Opportunities, Challenges, and Approaches,” *IEEE Computer, Special Issue on Blockchain & Cyber-physical Systems*. vol. 53, no. 9, pp. 66-76, Sept. 2020.
- Ali Ozdagli, Carlos Barreto, and Xenofon Koutsoukos. “@PAD: Adversarial Training of Power Systems Against Denial-of-Service Attacks,” *Hot Topics in the Science of Security (HotSoS 2020)*. September 22-24, 2020.
- Carlos Barreto, Taha Eghtesad, Scott Eisele, Aron Laszka, Abhishek Dubey, and Xenofon Koutsoukos, “Cyber-Attacks and Mitigation in Blockchain Based Transactive Energy Systems,” *3rd IEEE International Conference on Industrial Cyber-Physical Systems (ICPS 2020)*, Tampere, Finland, June 10-12, 2020.
- Himanshu Neema, Peter Volgyesi, Xenofon Koutsoukos, Thomas Roth, and Cuong Nguyen, “Online Testbed for Evaluating Vulnerability of Deep Learning Based Power Grid Load Forecasters,” *8th Workshop on Modeling and Simulation of Cyber-Physical Energy Systems (MSPES 2020)*, Sydney, Australia, April 21, 2020.
- Himanshu Neema, Xenofon Koutsoukos, Bradley Potteiger, Cheeyee Tang, and Keith Stouffer, “Simulation Testbed for Railway Infrastructure Security and Resilience Evaluation,” *Hot Topics in the Science of Security (HotSoS 2020)*, September 22-24, 2020. Best paper award.



Quarterly Meetings

Winter 2020 Labet Quarterly meeting



The Winter 2020 Science of Security and Privacy (SoS) Quarterly Labet meeting was held at North Carolina State University (NCSU) on January 15-16, 2020. The Quarterly meeting was hosted by NCSU Principal Investigators (PIs) Laurie Williams and Munindar Singh and featured invited presentations by representatives from government, academia, and industry as well as a mix of Labet project presentations and a poster session. This Labet Quarterly also introduced PechaKucha presentations, providing each speaker, all NCSU students, five minutes to cover their topic.

Invited Presentations

Neal Ziring, Technical Director for NSA's Cybersecurity Directorate, gave a presentation entitled " 'Deep' Thoughts on Information Sharing." He provided an overview of the Cybersecurity Directorate's key missions, and emphasized that information sharing will help empower all stakeholders. The four 'deep' topics he cited were: 1) assessing information sharing/ the axes of assessment; 2) a simple capability-centric model for sharing participants; 3) the importance of standardization; and 4) a vision for shared awareness and coordinated action.

Mike Bender, Director of the joint NSA-NCSU Laboratory for Analytic Sciences, spoke about data breaches and their impact in a presentation entitled "The Weakest Link." He noted that protecting privacy, shielding intellectual property, and securing IT infrastructure are becoming more difficult and discussed the analytic challenges to addressing these problems, including how to optimize algorithms in the analysis since you never have all the data you need.

In a presentation entitled "Zero Trust 101: An Evolution in Enterprise Cybersecurity," Alper Kerman and Scott Rose of NIST discussed the NIST Special Publication Zero Trust Architecture (ZTA) which is the product of a multi-agency collaboration overseen by the Federal CIO Council. While not intended to be a single deployment plan for ZTA, the publication describes ZTA strategies for enterprise network architectures and provides a roadmap to migrate and deploy ZTA concepts to an enterprise network.

Anthony Grieco, Trust Strategy Officer at Cisco spoke on "Building Trust into the Future," and addressed what Cisco is doing to enhance trust relationships with its customers. He described the way Cisco is building resilience into their solutions based on the foundation of trust and noted that they want to make sure they have shared goals with executives that are not security people.

"Privacy in a Decentralized World: Crypto Tools for blockchains, Applications, and Governance" was the title of the presentation given by Alessandra Scafuro of NCSU. Key tenets of a decentralized world include public verifiability, distributed trust, smart contracts, interoperability, and distributed governance, and she described cryptologic primitives that can enable privacy-preserving blockchain governance.



Alessandra Scafuro, NCSU, spoke on blockchains and privacy.

Lablet Project Presentations

Jonathan Aldrich, Carnegie Mellon University, spoke on “Usability Evaluation of the Obsidian Smart Contract Language” in which he provided the results from experiments that evaluated the usability of the final Obsidian language design, evaluating how quickly developers can learn Obsidian and whether it helps them avoid making errors when writing smart contracts. The results showed that experienced developers can learn Obsidian in about 90 minutes; developers can use Obsidian to do interesting tasks; and developers write contracts with fewer vulnerabilities than in Solidity, leading the researchers to conclude that an interdisciplinary language design approach can provide both assurance and usability. [More information on this project can be found here.](#)

Perry Alexander, Anna Fritz, and Adam Petz of University of Kansas gave a presentation entitled “The Attestation Monad – A Principal Architecture for Remote Attestation.” The research goals for the project include formal semantics of trust; verified remote attestation infrastructure; enterprise attestation and appraisal; and sufficiency and soundness of measurement.

Andy Meneely of Rochester Institute of Technology, a Sub-Lablet of NCSU, spoke on “Discovery and Attacker Behavior in a Penetration Testing Competition” in which he described the work being done analyzing over 9 TB of data collected during collegiate penetration testing competitions. The goal of the analysis is to assess discoverability via behavior, and they are in the process of developing a stochastic Markov chain model. [More information on this project can be found here.](#)

In her presentation “Analytics for Cybersecurity of Cyber-Physical Systems,” Nazli Choucri of Massachusetts Institute of Technology, a Sub-Lablet of Vanderbilt University, provided an update on the project which is focused on introducing analytics for CPS cybersecurity to enhance value of guidelines and directives. The expected product of the research is a platform for cybersecurity analytics with customized tools to support user needs. [More information on this project can be found here.](#)

David Nicol, University of Illinois at Urbana-Champaign, spoke on “Efficient Estimation of the Cyber-Attack Loss Distribution” in which he updated the work he has been doing on developing the

mathematical basis for quantifying and including uncertainty into system security analysis. The current research seeks to estimate the cost of cyber attacks given uncertainty in the presence of vulnerabilities and the routing and application of access control. [More information on this project can be found here.](#)

Nick Doty, University of California, Berkeley, a Sub-Lablet of the International Computer Sciences Institute, gave a presentation entitled “Finding Solutions for Privacy Problems: Privacy Design Patterns.” In this presentation he provided a definition of design patterns and why they were important as well as the challenges associated with identifying more design patterns. [More information on this project can be found here.](#)



Nick Doty, UC Berkeley, addressed privacy design patterns.

The complete agenda and selected talks, including the PechaKucha presentations, can be found [HERE](#).

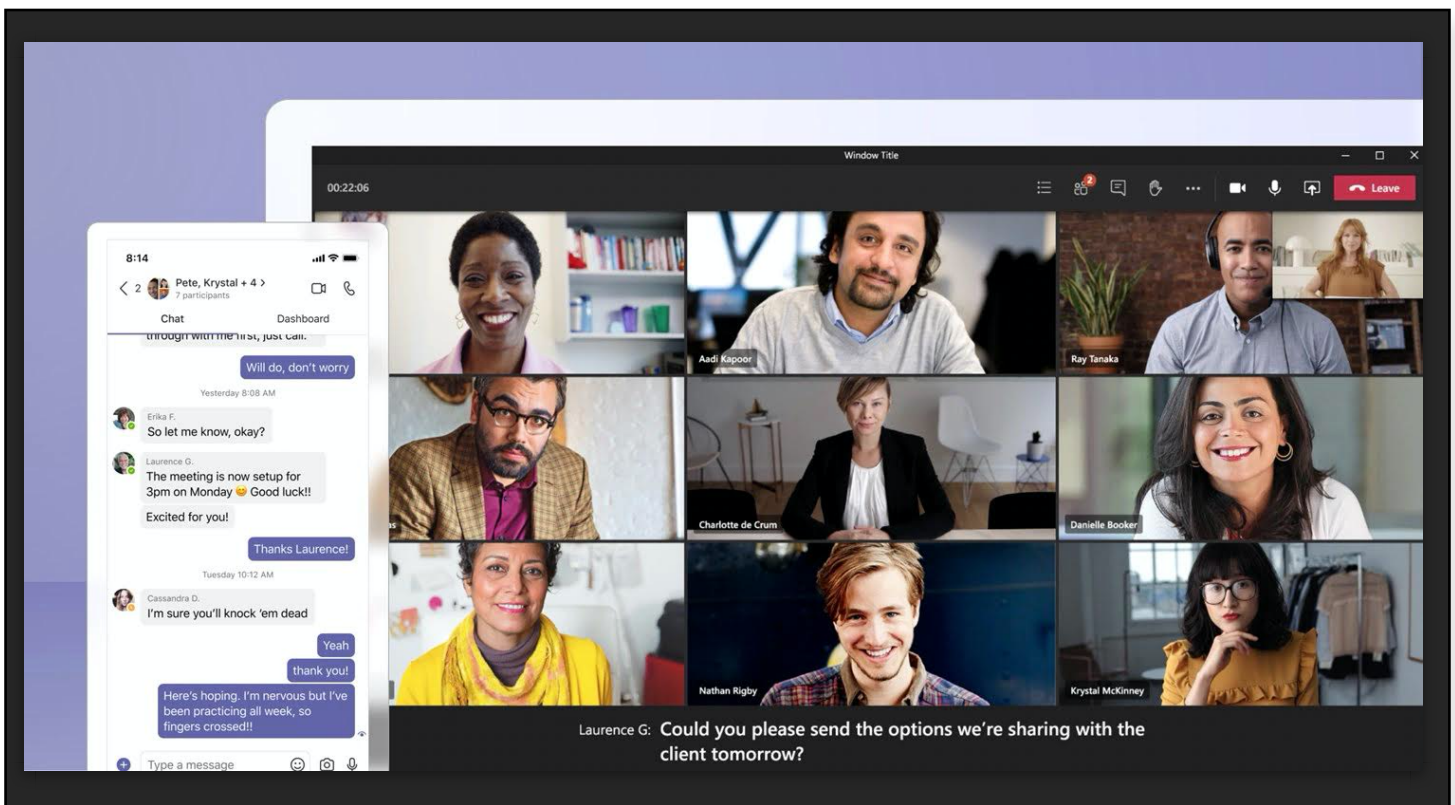
SoS Virtual Seminar Series

In response to NSA's transition to telework, the SoS team launched the SoS Virtual Seminar Series in November of 2020. The seminars, held using Microsoft Teams on NSA's Unclassified Work Environment (UWE), allow NSA personnel who are teleworking to join other teleworkers, as well as those in the office, to hear about Lablet research and interact with Lablet community members. One of the goals of the seminar series is to broaden the exposure of Lablet research to additional NSA organizations, as well as to showcase the breadth of Lablet research concentrations. For the presenters, the seminar series is envisioned as a practice environment for students sponsored by the SoS Lablets, often a dissertation prior to a thesis defense or before a conference presentation.

The first seminar, held in late November, was attended by 18 people. The speakers were Bill Sanders, Dean of the College of Engineering at CMU and formerly PI at the UIUC Lablet, and Ben Ujcich, a recent graduate of the UIUC Lablet and now Assistant Professor in the Department of Computer Science at Georgetown University. Dean Sanders gave an overview of the project and Professor Ujcich presented on his thesis work. The project, "Monitoring, Fusion, and Response for Cyber Resilience" seeks to improve the resilience of a system's intrusion and incident detection infrastructure against monitor compromise.

Abstract of the talk: Software-defined networking (SDN) has seen rapid growth because of its unparalleled flexibility in network configuration. This flexibility results from programmable control planes in which developers can write network apps and operators can execute them. Network operating systems oversee shared control plane resources. Although SDN has been touted as a solution to long-standing challenges of consistent network security policy enforcement, SDN is vulnerable to new attack vectors that undermine such policies. This talk provides an overview of the author's efforts to explore attack vectors that impact the SDN control plane and network operating system security. These efforts have led to the design and implementation of information flow control in the SDN control plane, the systematic discovery of 15 novel control plane vulnerabilities in popular network operating systems used by telecommunications and cloud providers, and the use of data provenance and program analysis techniques to analyze and record cross-layer data dependencies.

This research was presented at the 2020 Internet Society's Network and Distributed System Security Symposium (NDSS '20).



Section 2

Promoting Rigorous Scientific Principles



In 2020 the Science of Security and Privacy Initiative (SoS) promoted rigorous scientific principles through their funding of fundamental research at the Lablet universities as well as through the 8th Annual Best Scientific Cybersecurity Paper Competition. The SoS initiative also sponsored a Best Paper Award at the Hot Topics in the Science of Security Symposium (HotSoS 2020).

The National Security Agency's Research Directorate selected "Spectre Attacks: Exploiting Speculative Execution" as the winner of its 8th Annual Best Cybersecurity Research Paper competition. Originally published at the 2019 IEEE Security & Privacy Symposium, the winning paper, in combination with Meltdown, a 7th Annual Paper Competition Honorable Mention by the same researchers, launched a global effort to mitigate critical vulnerabilities in processors.

The winner of the HotSoS Best Paper competition, "Simulation Testbed for Railway Infrastructure Security and Resilience Evaluation" was written by researchers from Vanderbilt University, Johns Hopkins Applied Physics Laboratory, and the National Institute of Standards and Technology. The winning paper was automatically submitted to the 9th Annual Best Scientific Cybersecurity Paper Competition.

The SoS initiative had sponsored awards at the Intel International Science and Engineering Fair (ISEF) for five consecutive years, but the competition was cancelled this year due to Covid.

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



Congratulations to the authors of "Spectre Attacks: Exploiting Speculative Execution," for winning @NSAGov's 8th Annual Best Scientific #Cybersecurity Research Paper Competition. Learn more about how this research makes processors & computers more secure:



Winner of NSA's 8th Annual Best Scientific Cybersecurity Research Paper
The National Security Agency's Research Directorate selected "Spectre Attacks: Exploiting Speculative Execution" as the winner of its 8th Annual Best ...
nsa.gov

2:45 PM · Nov 30, 2020 · Sprout Social

29 Retweets 1 Quote Tweet 65 Likes

8TH Annual Paper Competition



The National Security Agency and Science of Security selected “Spectre Attacks: Exploiting Speculative Execution” as the winner of its 8th Annual Best Scientific Cybersecurity Paper competition.

Originally published at the 2019 IEEE Security & Privacy Symposium, the winning paper, in combination with “Meltdown: Reading Kernel Memory from User Space” a 7th Annual Honorable Mention winner by the same researchers, launched a global effort to mitigate critical vulnerabilities in processors.

Although the Spectre vulnerability was independently discovered and reported by various teams of researchers from Europe and the United States, twelve researchers collaborated to document their findings within one paper. The twelve researchers and their institutions are:

- Paul Kocher, Independent
- Jann Horn, Google Project Zero
- Anders Fogh, Intel Corporation
- Daniel Genkin, University of Michigan
- Daniel Gruss, Graz University of Technology
- Werner Haas, Cyberus Technology
- Mike Hamburg, Rambus Labs
- Moritz Lipp, Graz University of Technology
- Stefan Mangard, Graz University of Technology
- Thomas Prescher, Cyberus Technology
- Michael Schwarz, CISPA Helmholtz Center for Information Security
- Yuval Yarom, University of Adelaide and Data61

The landmark Spectre research uncovered how a performance feature of modern computer and mobile device processors is vulnerable to leaking private and sensitive data. Specifically, when idle, modern processors predict what will be needed to be computed next and then use this predictive result if the prediction is correct and discard the result if incorrect. The researchers found an opportunity to leak data when they tricked the processor in computing a prediction that would be found to be incorrect and a violation of security protections — a vulnerability internationally known as Spectre.

The winning research team demonstrated that a long-held bedrock assumption about computing security was not valid. Their efforts visibly displayed key tenets of science including the importance of reviewing past results, testing assumptions, employing rigorous methodologies, and verifying and documenting results. Already having been cited in 1,000+ subsequent research papers, this study is spawning a review of previous research and launching new inquiries. Additionally, the paper had broad scope because the researchers tested many platforms to understand the nature of the issue. This research will have a profound impact on how future processors and computers are built.

Due to the Covid-19 Pandemic, the 8th Competition did not feature a research presentation at NSA.

About the Competition:

The Best Scientific Cybersecurity Paper Competition is sponsored yearly by NSA’s Research Directorate and reflects the Agency’s desire to increase scientific rigor in the cybersecurity field. The competition was initiated in 2013 to encourage



the development of scientific foundations in cybersecurity and support enhancement of cybersecurity within devices, computers, and systems through rigorous research, solid scientific methodology, documentation, and publishing. The competition recognizes current research that exemplifies the development of scientific rigor in cybersecurity research. Papers published in peer-reviewed journals, magazines, or technical conferences are eligible for nomination.

SoS is a broad enterprise, involving both theoretical and empirical work across a diverse set of topics. While there can only be one best paper, no single paper can span the full breadth of SoS topics. Nevertheless, work in all facets of security science is both needed and encouraged.

For this year's paper competition, a group of ten internationally renowned cybersecurity experts along with NSA experts collectively reviewed 52 nominated papers. After review and ranking, the Distinguished Experts forwarded their recommendations to NSA for final selection. This year's winning research received high praise from the experts including, "This paper is hugely influential, well written, and well done scientifically," and "This paper is top notch, pure and simple."

The Distinguished Experts were:

- Dr. Whitfield Diffie, Unaffiliated
- Prof Kathleen Fisher, Tufts University
- Dr. Dan Geer, In-Q-Tel
- Dr. Eric Grosse, Unaffiliated
- Dr. John Launchbury, Galois Inc
- Dr. Sean Peisert, Lawrence Berkeley National Laboratory
- Prof Stefan Savage, University of California, San Diego
- Mr. Phil Venables, Goldman Sachs
- Dr. Arun Vishwanath, Unaffiliated
- Ms. Mary Ellen Zurko, MIT Lincoln Laboratory

The 9th Annual Best Scientific Cybersecurity Paper Competition opened for nominations on December 15, 2020 for papers published during calendar year 2020 in peer-reviewed journals, magazines, or technical conferences that show an outstanding contribution to cybersecurity science.



Home » NSA announces winner of 8th Annual Best Scientific Cybersecurity Research Paper Competition

- Technologies
- Management
- Cyber
- Security Newswire
- Security Enterprise Services
- Security Leadership and Management
- Logical Security
- Security & Business Resilience
- Cyber Security News

NSA announces winner of 8th Annual Best Scientific Cybersecurity Research Paper Competition



December 1, 2020

KEYWORDS cyber security / mobile device security / national security
Order Reprints

NSA's Best Scientific Cybersecurity Research Paper Competition was initiated in 2013 with the intent to encourage the development of scientific foundations in cybersecurity and support enhancement of cybersecurity within devices, computers, and systems through rigorous research, solid scientific methodology, documentation, and publishing. Papers published in



SoS-VO.org
@SoS_VO_org

Congrats to the 2020 Best Scientific Cybersecurity Paper Competition Winners! Each year is competitive, and this last cycle was no exception!

Stay tuned for the 9th Annual Competition announcement!

More Info: cps-vo.org/node/72048

[#CyberSecurity](#) [#ScienceofSecurity](#)



Phil Venables
@philvenables

Perhaps unsurprisingly, but no less deservingly, the Spectre Attacks paper won the NSA's 8th Annual Best Scientific Cybersecurity Research Paper Competition.

nsa.gov/News-Features/...



Yuval Yarom
@yuvalyarom

The Spectre paper wins the NSA's 8th Annual Best Scientific Cybersecurity Research Paper Competition.
nsa.gov/News-Features/...

Work with Paul Kocher, [@tehjh](#), [@anders_fogh](#), Daniel Genkin, [@lavados](#), Werner Haas, Mike Hamburg, [@mlqxyz](#), [@StefanMangard](#), Thomas Prescher, and [@misc0110](#).

Section 3

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 2020, despite the pandemic, the SoS continued to grow the Science of Security community of interest.

For the seventh year, SoS sponsored the Hot Topics in the Science of Security Symposium (HotSoS) which was held virtually and hosted by the University of Kansas on September 22-24, 2020. HotSoS is held in cooperation with the Association for Computer Machinery (ACM), and over 400 participants from government, industry, and academia attended HotSoS 2020, the largest HotSoS attendance ever.

The SoS-Virtual Organization (SoS-VO), a longstanding initiative designed to grow the Science of Security community, reached over 1850 members in 2020. It continues to provide a centralized location for cybersecurity research, events, and news, and was critical to maintaining awareness of SoS research, activities, and initiatives during the pandemic.

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 reaches over 1700 subscribers. This year also saw a growing SoS presence on social media, with close to 350 Facebook postings to nearly 100 members. 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.



Science of Security's Annual Security Conference Goes Virtu:

<https://www.nsa.gov/News-Features/Feature-Stories/ArticleView/Article/2396500/science-of-securitys-annual-security-conference-goes-virtual-and-get-r...> Photo of Adam Tagert behind computer monitor FT. MEADE, Md., Oct. 28, 2020 -- October is National Cybersecurity Awarene Month (NCSAM), so we have been celebrating...

Hot SoS



Hot Topics in the Science of Security (HotSoS)

The University of Kansas virtually hosted the 7th Annual Symposium on the Science of Security (HoTSoS), from 22-24 September 2020. HotSoS brings together researchers from diverse disciplines to promote the advancement of work related to the Science of Security and features a mix of invited talks, panels, tutorials, and refereed papers to be published by ACM. HotSoS was established several years ago to set standards for the SoS community as well as to provide a venue for publishing foundational scientific work, which is often underappreciated in larger conferences. The proceedings of the seven HotSoS conferences have been published by ACM and, as a result, other conferences and symposia have begun to appreciate the value of Science of Security research. This year's virtual event, originally scheduled for in-person attendance in April, provided the opportunity for those who might otherwise be unable to attend to engage with other SoS researchers. Perry Alexander (KU) served as General Chair with Drew Davidson (KU) and Baek-Young Choi (UMKC) as Program Co-Chairs. Virtual Hotsos was run on a new online platform, Hopin and due to the conference being virtual, registration fees were waived for everyone. Over 400 participants representing 126 institutions, including government, industry, and academia registered; a new attendance record. In addition to keynote and paper presentations, participants also reviewed Works-in-Progress (WiP) and posters.

Keynote presentations

1. Access Control Verification for Everyone

Andrew Gacek of Amazon Web Services (AWS)

Dr. Gacek's presentation focused on AWS' Identity and Access Management (IAM) Access Analyzer, an automated reasoning service for auditing permissions to cloud resources. Dr. Pacek noted that few customers have the requisite skills to formally specify and verify security properties, thus limiting the number of customers who may be able to take advantage of increased security capabilities. The AWS Access Analyzer verification approach focuses on taking a policy and breaking it down into a specification so that findings are sound, precise, and compact. The Access Analyzer quickly and automatically discovers security properties and asks customers which ones are intended, thus eliminating the skill barrier and setup costs associated with traditional formal verification. In response to a question concerning whether automated reasoning has the

ability to highlight or recommend potential or known security vulnerabilities for the intended access policy, Dr. Gacek noted that the findings can help customers find misconfigurations and the customer can then change it and rescan immediately to determine if the problem has been fixed, thus making a connection between the finding and the policy.

2. Is Hardware Root of Trust hard to do, and Trustworthy?

Lyle Paczkowski, Sprint

Mr. Paczkowski began his presentation by discussing 5G, its impact on the growth of IoT, and the resultant security concerns that must be addressed. He addressed communication, physical, lifecycle, and software vulnerabilities, but suggested that the most urgent system that requires trust and secure data is the supply chain, and that hardware Root of Trust is part of the answer. He suggested using a Root of Trust to provide assurance of a device's origin, thus establishing the foundation for a Trusted Supply Chain, with Blockchain adding an additional layer of trust in the overall system supply chain. He concluded his presentation by discussing emerging problems and technology trends which, combined with the impact of 5G and IoT, will cause supply chain to be the most significantly affected business process.

3. Trust Engineering via Cryptographic Protocols

Joshua Guttman, Worcester Polytechnic Institute/MITRE

Dr. Guttman noted that trust engineering requires a system design with each decision based on definite assumptions and reliable conclusions about peers, that peer answers are cryptographically protected, and there is a relationship between protocol analysis and trust. He spoke of designing the protocol to include device behaviors, controller behavior and protocol considerations, and also addressed rules and SGX.

4. Evaluating Fuzz Testing (and other technologies)

Michael Hicks, University of Maryland and Correct Computation Inc.

The paper "Evaluating Fuzz Testing" was the winner of the SoS 7th Annual Best Scientific Cybersecurity Paper Competition. Traditionally, the winner of the previous Best Cybersecurity Paper Competition is invited to HoTSoS to present the winning work.

Dr. Hicks described the research process, noting that the research team investigated the evaluation process of fuzz testing tools,



which are used to 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. In addition to presenting the findings in the paper, Dr. Hicks spoke about the impact of the paper since its publication, including whether things have changed, and provided some guidelines for fuzz testing.

Paper Sessions

Three paper sessions dealt with CPS and Industrial Control, Modeling, and Systems. The 12 papers selected for presentation represented the work of 46 authors from 16 institutions, 13 of which are universities. The Best Paper award was given to the presentation, part of the first paper session, entitled “Simulation Testbed for Railway Infrastructure Security and Resilience Evaluation” by researchers from Vanderbilt University, Johns Hopkins Applied Physics Laboratory, and the National Institute of Standards and Technology.

Paper Session

1: CPS and Industrial Control

1. Simulation Testbed for Railway Infrastructure Security and Resilience Evaluation

Himanshu Neema and Xenofon Koutsoukos, Vanderbilt University

Bradley Potteiger, Johns Hopkins Applied Physics Laboratory
CheeYee Tang and Keith Stouffer, National Institute of Standards and Technology

This paper, the winner of the HotSoS 2020 Best Paper Award, describes a domain-specific framework for simulations in the railway domain. The framework allows analyzing the resilience of railway operations in the presence of cyber-attacks. In particular, this simulation framework allows modeling the railway network as well as the railway transportation. It provides an online graphical modeling environment that allows multiple users to collaborate, through a web-based interface, over the same model for the railway infrastructure as well as network attacks. The framework also allows the user to configure and run experiments through the web-interface and to visualize the key operational metrics from the railway domain as the experiment is running. The framework also supports executing large simulations in the cloud. In addition, it supports Hardware-in-the-Loop (HIL) simulation for incorporating physical effects and network attacks that can only be realized realistically in the hardware. A detailed case study is provided to demonstrate the

framework’s capabilities. The testbed has been successfully transitioned to NIST’s Engineering Laboratory where it is being further developed and refined for real-world use-cases. Researchers plan to apply the testbed to other transportation applications such as self-driving vehicles, and to extend model libraries with more reusable cyber-attacks and security solutions.

2. @PAD: Adversarial Training of Power Systems Against Denial-of-Service Attacks

Ali Ozdagli, Carlos Barreto, and Xenofon Koutsoukos Vanderbilt University

This paper builds upon the authors’ work dealing with vulnerabilities of systems that protect against cyber attacks in power grid systems. Earlier research presented an optimization problem to determine which sensors to attack within a given budget such that the existing classifier can be deceived. This research extends that optimization problem to find attacks for more complex classifiers such as neural networks. Their research demonstrates that a neural network, in particular, with RELU activation functions, can be represented as a set of logic formulas using Disjunctive Normal Form, and the optimization problem can be used to efficiently compute a DoS attack. In addition, they propose a defense model that improves the resilience of neural networks against DoS through adversarial training, and, finally, they evaluate the efficiency of the approach using a dataset for classification in power systems.

3. The More the Merrier: Adding Hidden Measurements for Anomaly Detection and Mitigation in Industrial Control Systems

Jairo Giraldo, University of Utah

David Urbina, University of Texas at Dallas

CheeYee Tang, National Institute of Standards and Technology

Alvaro Cardenas, The University of California, Santa Cruz

Industrial Control Systems (ICS) collect information from a variety of sensors and then use that information to control some physical components. For economic reasons and to achieve efficiency, there are typically only a small number of sensors, but as ICS attacks continue to grow there is a need to develop a systemic way to add sensors to the system. The authors propose enhancing system security by the addition of hidden sensor measurements, which they define as measurements that were not considered in the original design of the system, and are not used for any operational reason—they are only added to improve system security and are used in anomaly detection and mitigation. This paper shows that the addition of these new, independent, but correlated measurements to the system makes it harder for adversaries to launch false-data injection stealthy attacks and, even if they do, it is possible to limit the impact caused by those attacks. When an attack is detected, the compromised sensor measurements are replaced with estimated ones from the new sensors, improving the risky open-loop simulations proposed by previous work.

4. RUCKUS: A Cybersecurity Engine for Performing Autonomous Cyber-Physical System Vulnerability Discovery at Scale

Bradley Potteiger, Jacob Mills, Daniel Cohen, and Paul Velez

Johns Hopkins Applied Physics Laboratory

While the 2016 Cyber Grand Challenge (CGC) demonstrated the need to move toward an autonomous cybersecurity approach and improvements have been noted since then,

most of the efforts have focused on enterprise systems, leaving a gap in the CPS domain. With a large amount of legacy software, CPS remain an extremely vulnerable target. The authors propose to apply the most beneficial concepts from the CGC to create more secure and resilient CPS, and this paper introduces a CPS security assessment architecture, RUCKUS, for autonomously identifying and analyzing CPS firmware, identifying vulnerabilities, and developing exploits. The approach also considers how to integrate graph analytics to extrapolate findings to firmware at scale, allowing for measuring the potential widespread impact of attacks. The architecture is demonstrated using an automotive case study, leveraging firmware from the most popular automotive and router manufacturers to assess the real-world potential impact of CPS attacks.

Paper Session 2: Modeling

1. Exploring Hackers Assets: Topics of Interest as Indicators of Compromise

Mohammad Alramahi, Izzat Alsmadi, and Joshua Davenport
Texas A&M, San Antonio

As hacking techniques continue to proliferate, cyber defenders need to adopt proactive and offensive approaches within hackers' territories. The authors propose a systematic approach to automatically extract Topics of Interest (ToI) from hackers' websites that can eventually be used as inputs to actionable security controls or Indicators of Compromise (IoC) collectors. The authors selected the hackers' news website "CrackingFire" as a test case. ToI can be integrated into IoC and, once correlated with other signs of attacks, those IoC will trigger further cybersecurity offense or defense actions. The researchers also developed their own dark web crawler and evaluated extracting ToIs while observing the types of challenges in both the crawling and the processing stages.

2. Cyber Threat Modeling and Validation: Port Scanning and Detection

Eric Vugrin, Jerry Cruz, Christian Reedy, Thomas Tarman, and Ali Pinar
Sandia National Laboratories

Because port scanning is a commonly applied technique in the discovery phase of cyber attacks, defending against them has long been the subject of many research and modeling efforts. Though modeling efforts can search large parameter spaces to find effective defensive parameter settings, confidence in modeling results can be hampered by limited or omitted validation efforts. This paper introduces a novel, mathematical model that describes port scanning progress by an attacker and intrusion detection by a defender. The paper further describes a set of emulation experiments that were conducted with a virtual testbed and used to validate the model. Results are presented for two scanning strategies: a slow, stealthy approach and a fast, loud approach. Estimates from the model fall within 95% confidence intervals on the means estimated from the experiments. Consequently, the model's predictive capability provides confidence in its use for evaluation and development of defensive strategies against port scanning.

3. Can We use Software Bug Reports to Identify Software Vulnerabilities Strategies

Farzana Ahamed Bhuiyan, Raunak Shakya, and Akond Rahman
Tennessee Technological University

The goal of this paper is to help cybersecurity researchers characterize vulnerabilities by conducting an empirical study of software bug reports. The researchers performed qualitative analysis on 729, 908, and 5336 Open Source Software (OSS) bug reports respectively, collected from Gentoo, LibreOffice, and Mozilla to investigate if bug reports include vulnerability discovery strategies i.e. sequences of computation and/or cognitive activities that an attacker performs to discover vulnerabilities, where the vulnerability is indexed by a credible source, such as the National Vulnerability Database (NVD). They evaluated two approaches, text feature-based and regular expression-based, to automatically identify bug reports that include vulnerability discovery strategies, and observed the Gentoo, LibreOffice, and Mozilla bug reports to include vulnerability discovery strategies. Using text feature-based prediction models, they observed the highest prediction performance for the Mozilla dataset with a recall of 0.78, and then observed a recall of 0.83 for the same dataset using the regular expression-based approach. The paper's findings provide the groundwork for cybersecurity researchers to use OSS bug reports as a data source for advancing the science of vulnerabilities.

4. Automated Influence and the Challenge of Cognitive Security

Sarah Rajtmajer and Daniel Susser
Pennsylvania State University

Advances in AI are powering increasingly precise and widespread computational propaganda, posing serious threats to national security. The military and intelligence communities are starting to discuss ways to engage in this space, but the path forward is still unclear. These developments raise pressing ethical questions, about which existing ethics frameworks are silent. The authors argue in this paper that understanding these challenges through the lens of "cognitive security," offers a promising approach.



Paper Session 3: Systems

1. Neutralizing Manipulation of Critical Data by Enforcing Data-Instruction Dependency

Chandra Sharma, Nathan Miller, and George Amariuca
Kansas State University

The authors propose a new approach to neutralize attacks that tamper with critical program data with a technique that uses a sequence of instructions as a trap against the illicit modification of the critical data. Specifically, they established a dependency such that the continued execution of the program is contingent upon the successful execution of the instruction sequence and the successful execution of the instruction sequence is contingent upon the integrity of the critical data. In particular, they present a specific implementation of the technique focusing on a critical data that is often subject to malicious manipulation: the return address of a function. The paper shows that their technique can be an effective countermeasure to defend against attacks that overwrite the return address to divert control to a malicious code, and they further demonstrated that their technique offers significant protection without resorting to complementary defenses such as ASLR, DEP or StackGuard.

2. Ghostbusting: Mitigating Spectre with Intraprocess Memory Isolation

Ira Jenkins, Prashant Anantharaman, Rebecca Shapiro, J. Peter Brady, Sergey Bratus, and Sean Smith Dartmouth College

Speculative execution creates so-called transient instructions, those whose results are ephemeral and not committed architecturally. However, various side-channels exist to extract these transient results from the microarchitecture, e.g., caches. Spectre Variant 1, the so-called Bounds Check Bypass, was the first such attack to be demonstrated. Leveraging transient read instructions and cache-timing effects, the adversary can read secret data. This paper explores the ability of intraprocess memory isolation to mitigate Spectre Variant 1 attacks, demonstrating this using Executable and Linkable Format-based access control (ELFbac), a technique for achieving intraprocess memory isolation at the application binary interface (ABI) level. The authors also consider Memory Protection Keys (MPKs), a recent extension to Intel processors, that partition virtual pages into security domains. Using the original Spectre Proof-of-Concept (POC) code, they show how ELFbac and MPKs can be used to thwart Spectre Variant 1 by constructing explicit policies to allow and disallow the exploit. They compare these techniques against the commonly suggested mitigation using serialized instructions, e.g., lfence, and consider other Spectre variants based on transient execution that intraprocess memory isolation would naturally mitigate.

3. WOLF: Automated Machine Learning Workflow Management Framework for Malware Detection and other Applications

Sohaib Kiani, Sana Awan, Fengjun Li, and Bo Luo University of Kansas

Applying ML techniques to solve real-world problems is a highly iterative process which may require up to thousands of iterations to find the optimum set of hyper-parameters. It is also difficult to find the best ML techniques for a given dataset. The WOLF framework presented in this paper has been designed to simultaneously automate the process of selecting the best algorithm and searching for the optimum hyper-parameters. It can be useful to both to those who are novices in ML and

just want to find the best algorithm for their dataset, and also to those who are experts in the field and want to compare their new features or algorithm with state-of-the-art techniques. By incorporating the WOLF framework in their designs, it is easier for

novices to apply ML techniques on their dataset. With a wide range of evaluation metrics provided, WOLF also helps data scientists develop better intuition towards ML techniques and speed up the process of algorithm development. Another main feature of the WOLF framework is that users can easily integrate new algorithms at any stage of the ML pipeline. This paper both presents the WOLF architecture and demonstrates how it could be used for standard ML datasets and for Android malware detection tasks. Experimental results show the flexibility and performance of WOLF.

4. A Formal Security Analysis of Zigbee

Li Li and Endadul Hoque, Syracuse University
Proyash Podder, Florida International University


The rapid increase in the number of IoT devices in recent years, and the competition to develop them, has given rise to poor quality IoT devices that are prone to adversarial attacks that can compromise both security and safety. Many recently reported attacks are due to the insecurity present in the underlying communication protocol stacks, and ZigBee is one of them. Considering the emergence and adoption of ZigBee 3.0 and the current market share of ZigBee 1.0, it is essential to study and analyze these protocol stacks at their specification level so that any insecurity at the specification level should be identified and fixed before they go into production. With that goal in mind, this paper presents a model for ZigBee (1.0 and 3.0) and reasons about its security properties using a security protocol verification tool (named Tamarin). This model of ZigBee closely follows the ZigBee specification, and the security properties are derived from the ZigBee specification. The authors use Tamarin to verify these properties on the model and report their findings on ZigBee 1.0 and ZigBee 3.0.



7TH ANNUAL

HOT TOPICS *in the* SCIENCE OF SECURITY

SEPT 22-23, 2020 | UNIVERSITY OF KANSAS | LAWRENCE, KS



Works-in Progress (WiP) Sessions

The WiP sessions were introduced at last year's HotSoS and were again part of HotSoS 2020. These sessions assist authors in writing high quality research papers by having a community discussion on the research early in the research process. This early feedback promotes faster and easier publication by enabling researchers to adjust on-going research to respond to concerns traditionally raised after the paper is completed and it is being peer-reviewed. During WiP sessions 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. The ultimate goal for each WiP session is 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.

The papers shown below were part of the HotSoS 2019 WiP process and were subsequently published:

- "Privacy Attitudes of Smart Speaker Users" by Nathan Malkin, Joe Deatrack, Allen Tong, Primal Wiiesekera, David Wagner, and Serge Egleman; in Proceedings on Privacy Enhancing Technologies (PoPETS), 2019. The paper won the Privacy Papers for Policymakers Best Student Paper Award.
- "How does Misconfiguration of Analytic Services Compromise Mobile Privacy?" by Xueling Zhang, Xiaoyin Wang, Rocky Slavin, Travis Breaux, and Jianwei Niu; International Conference on Software Engineering, 2020.
- "Impact of Ad-Blockers on Consumer Behavior: A Lab Research Experiment" by Alisa Frik, Amelia Haviland, and Alessandro Acquisiti; Workshop on Economics of Information Security (WEIS) 2019, and USENIX Security 2020.

At HotSoS 2020, six papers were presented and discussed at three separate sessions to allow adequate time for a full exchange of ideas.

The following authors presented and discussed their research:

- Hao Xue, Qiaozhi Wang, Bo Lou, Chao Lan, and Fenjun Li University of Kansas and University of Wyoming
- Jacob Fustos and Heechul Yun University of Kansas
- Sohaib Kiani, Fengjun Li, Chao Lan, and Bo Luo University of Kansas and University of Wyoming
- Ryan Karl, Jonathan Takeshita, and Taeho Jung University of Notre Dame
- TSION Yimer, Md Tanvir Arafin, and Kevin Kornegay Morgan State University
- Sana Awan, Fenjun Li, Bo Luo University of Kansas



Posters

Twenty posters were submitted to HotSoS 2020, all of which were selected for presentation. The selected posters were from 17 universities and represented the work of 59 authors. Two Best Poster Awards were presented. 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 posters are described below:

Best Poster

1. Decentralized Backup and Recovery of TOTP Secrets

Conor Gilsenan, Noura Alomar, and Andrew Huang University of California, Berkeley

The Time-based One-Time Password (TOTP) algorithm is a widely deployed method of two-factor authentication (2FA). This research presents usability design requirements dealing with security, privacy, and usability for the backup and recovery systems of apps implementing the algorithm.

Best Undergraduate Poster

2. A Raspberry Pi Sensor Network for Wildlife Conservation

Andrew Arnold, Paul Corapi, Michael Nasta, Kevin Wolgast, and Thomas Babbitt United States Military Academy

This is the third consecutive HotSoS that USMA cadets have presented a poster on the Raspberry Pi Sensor Network. This poster depicts the progress they have made on the network to be used for wildlife detection and monitoring. Raspberry Pi sensor nodes collect and store data and transfer it over a mesh network to an android app interface.

3. A Curated Dataset of Security Defects in Scientific Software Projects

Justin Murphy, Elias T. Brady, Shazibul Islam Shamim, and Akond Rahman Tennessee Technological University

Julia is a programming language used to efficiently develop scientific software, software used to explore and analyze data to investigate unanswered research questions. This study uses Julia to address the research question dealing with how frequently security defects appear in scientific software projects.

4. A Preliminary Taxonomy of Techniques Used in Software Fuzzing

Raunak Shakya and Akond Rahman

Software fuzzing is an automated testing technique which provides invalid, unexpected, or random data as inputs to a software program so that the software of interest can be monitored for exceptions such as crashes. This study seeks to help researchers performing software fuzzing by categorizing different techniques used in software fuzzing literature.

5. Accelerating Block Propagation in PoW Blockchain Networks with Pipelining and Chunking (PiChu)

Kaushik Ayinala, Baek-Young Choi, and Sejun Song University of Missouri

Blockchain is an open, verifiable, and distributed consensus of transactions among different parties, relying on P2P technology for connectivity between nodes, but the time needed for block propagation limits inceptions of another consensus. This research presents a proposed method to accelerate block propagation in PoW blockchain networks by pipelining message transaction and verifications in parallel over a network with chunks of a block (PiChu). Simulation results demonstrate significantly less latency of block propagation than traditional method as the size of a P2P network increases.

6. An Infrastructure for Faithful Execution of Remote Attestation Protocols

Adam Petz University of Kansas

An emerging technology for gaining trust in a remote computing system is remote attestation, the activity of making a claim about properties of a target by supplying evidence to an appraiser over a network. This study shows the research that aims to design, implement and prove correct a collection of software components that provide a sound infrastructure for remote attestation of layered systems.

7. An seL4-based Architecture for Layered Attestation

Grant Jurgensen, Michael Neises, and Perry Alexander University of Kansas

Remote attestation is an essential tool for identifying malicious or compromised actors but attestation evidence is only as trustworthy as the architecture it was collected on. This study describes how the researchers seek to incorporate key management into the chain-of-trust process where a layer's private key is only delivered to its attestation component after it has been measured.

8. An Uncertain Graph-based Approach for Cyber-security Risk Assessment

Hoang Hai Nguyen University of Illinois at Urbana-Champaign

This study presents a risk assessment approach for quantifying the security risk of lateral movement attacks whereby the attack propagation is modeled as an uncertain graph and the attack impact is a function of the set of compromised devices.

9. Application of the Armament Cyber Assessment Framework

Aidan McCarthy, Liam Furey, Keagan Smith, Daniel Hawthorne, and Raymond Blaine The United States Military Academy

This research introduces the Armament Cyber Assessment Framework (ACAF), the goal of which is to introduce a security-oriented mindset into the solution prior to release, and to provide meaningful results at every level by integrating security assessment workflows into the design process.

10. Building a Conceptual Framework for Ethical Hacking

Danielle Alexandre, Rebecca Labitt, and Asher Rodriguez Simmons University

This study addresses research questions dealing with the current research trends including, current terminology and concepts, used in ethical hacking; current challenges and best practices in ethical hacking; and how the findings of the researchers' multiple case-study, relate to each of three industry case studies in ethical hacking. The results of their research suggest an improved framework for research that encompasses a multitude of factors and attributes of major attacks that threaten computer security.

11. Do Configuration Management Tools Make Systems More Secure? An Empirical Research Plan

Md Rayhanur Rahman, William Enck, and Laurie Williams North Carolina State University

While Configuration Management Tools (CMT) help developers manage the system and installed application in an automated and efficient manner, misconfiguration can make a system vulnerable to compromises. This study proposes an empirical research plan on the impact of CMT on systems where these tools have been applied.

12. Exploiting DRAM Bank Mapping and HugePages for Effective Denial-of-Service Attacks on Shared Cache in Multicore

Michael Bechtel and Heechul Yun University of Kansas

This research proposes memory-aware cache DoS attacks that can induce more effective cache blocking by taking advantage of information of the underlying memory hardware.

13. How to Swap Instructions Midstream: An Embedding Algorithm for Program Steganography

Ryan Gabrys, Luis Martinez, and Sunny Fugate Naval Information Warfare Center

The poster presents research that proposes an encoding/decoding algorithm for program executable steganography which, unlike previous work, does not require the introduction of new instructions which may be detectable. The approach also does not require storing the locations where changes in the program executable are made.

14. Improving Architectures for Automating Network Security Using Specification-Based Protocols

Khair Henderson and Kevin Kornegay Morgan State University

The research presented in this poster leverages a specification-based protocol called Manufacturer Usage Description (MUD) that is designed to automate access control at the network edge where IoT devices reside. The research presented in this poster leverages a specification-based protocol called Manufacturer Usage Description (MUD) that is designed to automate access control at network edge where IoT devices reside. This approach leads to improved network security by



underlining inherent weaknesses and key research areas to create a sustainable and scalable resilient architecture.

15. Resilient Multi-Robot Target Pursuit

Jiani Li, Waseem Abbas, and Xenofon Koutsoukos, Vanderbilt University, Mudassir Shabbir, Information Technology University

The research presented in this poster addresses resilient distributed diffusion in multi-robot networks where robots, some of which might be adversarial, move in a cooperative manner to pursue some target. The researchers consider an adapt-then-combine diffusion algorithm whereby each non-adversarial robot computes its estimate about the target by optimizing a local cost function (adapt) and then aggregates the estimates received from neighbors (combine) that the robot uses in optimizing the local cost function.

16. Time Series Anomaly Detection in Medical Break the Glass

Qais Tasali, Nikesh Gyawali, and Eugene Y. Vasserman Kansas State University

This poster presents research on the effectiveness of anomaly detection to ease the human burden of post-hoc audits in the context of medical emergencies, aka Break-the-Glass (BTG) situations. The researchers use two different prediction models to perform real-time and post-BTG statistical analysis on time-series session log data for flagging anomalous user sessions and actions, an approach that combines a real-time fast analysis engine working on a partial feature set, as well as a post-hoc, slower analysis tool which works with the complete times series data of everything which occurred during the entire time of the emergency.

17. Tokens of Interaction: Psycho-physiological Signals, A Potential Source of Evidence of Digital Incidents

Nancy Mogire University of Hawaii at Manoa

This research investigates how the electrical signals generated by the body respond to human interaction with digital incidents. The research posits that if response-related signal changes are consistently notable and can be located within a recorded signal with an accuracy that is greater than chance, then it can be claimed that psycho-physiological signals contain markers of digital incidents.

18. Toward Just-in-Time Patching for Containerized Applications

Olufogorehan Tunde-Onadele, Yuhang Lin, Jingzhu He, and Xiaohui Gu North Carolina State University

Recent studies have shown that containerized applications are susceptible to various security attacks, and this research presents new on-demand targeted patching framework for containerized applications. This approach combines dynamic vulnerability exploit identification and targeted vulnerability patching to achieve more efficient security attack containment.

19. Using Intel SGX to Improve Private Neural Network Training and Inference

Ryan Karl, Jonathan Takeshita, Taeho Jung University of Notre Dame

This paper investigates a new technique utilizing a Trusted Execution Environment (TEE) to allow for the high-performance training and execution of Deep Neural Networks (DNNs), an ML algorithm that has recently been used with great success in a variety of challenging tasks, including speech and face recognition.

20. Vulnerability Trends in Web Servers and Browsers

M S Raunak and Richard Kogut, Loyola University
Richard Kuhn and Raghu Kacker, National Institute of Standards and Technology

The key questions researchers sought to address in this study are the trends in vulnerabilities for web servers and browsers and the magnitude of their impact on users; whether web browsers and servers are becoming more secure over time as vulnerabilities are discovered and programmers become more experienced; and how trends vary by vulnerability type?

The 2020 program agenda and presentations can be found at: <https://cps-vo.org/hotsos2020/agenda>.

The 2020 proceedings are now available in the ACM Digital Library and can be found here: <https://dl.acm.org/doi/proceedings/10.1145/3384217>



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 2020. Whether within their own organizations, in the classroom, 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 1850 members in 2020, 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.

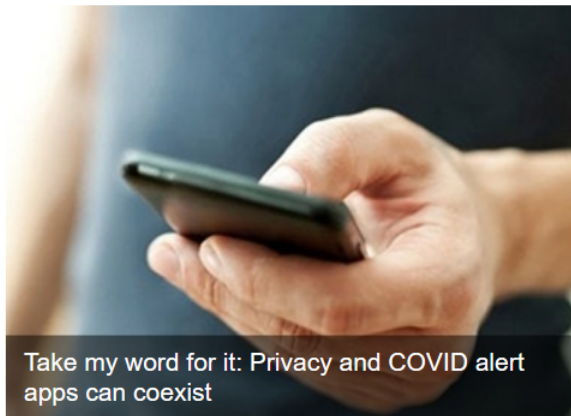
In 2020 the SoS-VO continued to provide information on SoS activities, to include Lablet research and meetings, HotSoS, the annual Best Scientific Cybersecurity Paper Competition, and other SoS activities. The SoS-VO enables its members to post research findings and publications done elsewhere, advertise community events, host chats, blog, create and participate in forums. It also provides information on upcoming events, position openings, calls for papers for conferences, and general cybersecurity news.

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 12 editions in 2020 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 over 3000 Pub Crawl items published in 2020 covering over 330 topics and representing the curated work of over 8000 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 2020 approximately 770 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



Take my word for it: Privacy and COVID alert apps can coexist

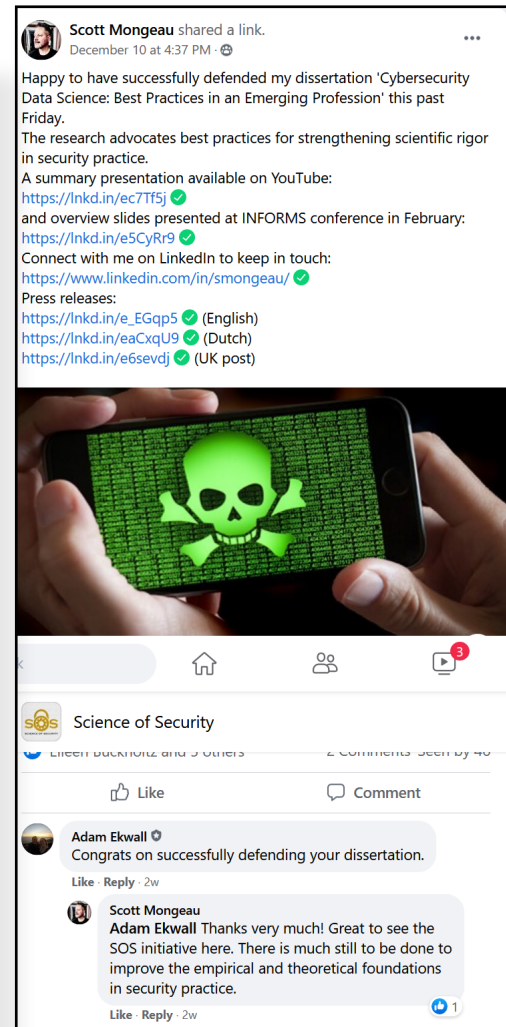

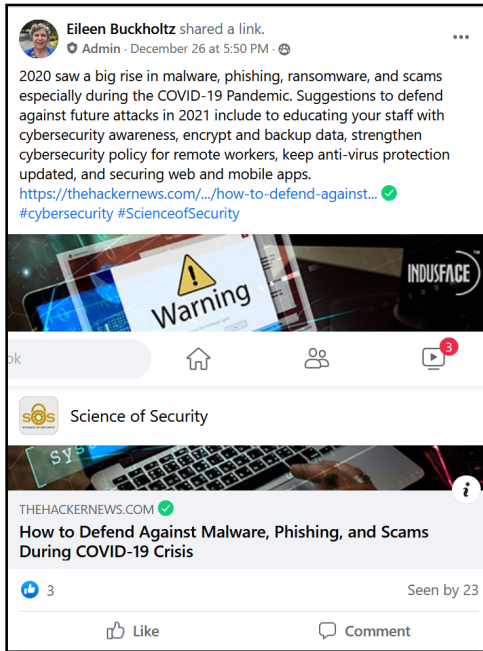
BY LORRIE CRANOR, OPINION CONTRIBUTOR -- 11/10/20 09:30 AM EST Since the COVID-19 pandemic began, technologists across the country have rushed to develop digital apps for contact tracing and exposure notifications. New York, New Jersey, Pennsylvania, and Delaware have all recently announced the launch of such apps,...



The Science of Security also maintains a Facebook presence, and there were close to 350 Facebook postings in 2020.

At the outset of the pandemic when lockdowns were instituted and in-person activities were curtailed, the SoS initiative worked to ensure that the SoS community remained connected. From April through September 2020, for example, there were 520 news items, 70 events, 1400 curated bibliographies, and 35 unique entries posted to the VO. In the six R&O newsletters that were published during that period, there were 630 articles dealing with Science of Security; of those, 120 were related to the pandemic. Also during that time, there were 120 postings on the SoS Facebook group, 30 of which were pandemic-related.

The SoS outreach efforts increase the likelihood that ad hoc and common practice approaches to security will be replaced by scientifically supported methods. By developing strategic rather than tactical methods of approaching cybersecurity, the practice of cybersecurity can be transformed to become efficient and proactive in both attack and defense.





produced by cyber pack ventures, inc.



SOS-VO.org