

SCIENCE OF SECURITY



ANNUAL REPORT

VISION

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

Science of Security and Privacy

Initiative 2022



Table of Contents

Executive Summary	5
Section 1: Engaging the Academic Community for Foundational Research	7
Hard Problems	8
Science of Security Lablets	
Carnegie Mellon University	11
International Computer Science Institute	20
North Carolina State University	29
University of Illinois at Urbana-Champaign	38
University of Kansas	47
Vanderbilt University	57
Science of Security Quarterly Meetings	66
Section 2: Promoting Rigorous Scientific Principles	73
Best Scientific Cybersecurity Paper Competition	74
Section 3: Growing the Science of Security	76
HotSoS 2021	77
Outreach	84
Cyber Makerspace	85

EXECUTIVE SUMMARY

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 ongoing 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 2021.

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

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

The six SoS Lablets are Carnegie Mellon University (CMU), the International Computer Science Institute (ICSI), North Carolina State University (NCSSU), 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 2021, 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 continuing its 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:

- 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
- Reasoning about Accidental and Malicious Misuse via Formal Methods

University of Illinois at Urbana-Champaign:

- 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:

- 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 2021 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 other activities undertaken by SoS that specifically reinforce that effort. The SoS initiative sponsored the 9th Annual Best Scientific Cybersecurity Paper Competition. There were 34 papers nominated in the 9th Annual Paper Competition bringing the total submissions to over 325 during the competition's nine years. This year's winning paper, "On One-way Functions and Kolmogorov Complexity," written by Yanyi Liu, Cornell University and Rafael Pass, Cornell Tech, was published at the 2020 IEEE Symposium on Foundations of Computer Science. Receiving honorable mention was the paper "Retrofitting Fine Grain Isolation in the Firefox Renderer" written by Shravan Narayan, Craig Disselhoen, Tal Garfinkel, Nathan Froyd, Sorin Lerner, Hovav Shacham and Deian Stefan. This paper was originally published at the USENIX Security Conference 2020.

Details on SoS activities to promote rigorous scientific principles in 2021 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. For the eighth year, SoS sponsored the Hot Topics in the Science of Security Symposium (HotSoS) which was held virtually and hosted by the NSA from 13-15 April 2021. Over 1200 individuals registered for HotSoS '21, and more

than 625 participated over the three days. The participants, a mix of government, academia, and industry, came from 36 countries. HotSoS 2021 was designed to encourage interaction among presenters and attendees and focus on WiP, posters, and student presentations.

The SoS-Virtual Organization (SoS-VO), a longstanding initiative designed to grow the Science of Security community, reached over 2000 members in 2021. 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 nearly 1800 subscribers. There also continued to be over 300 Facebook postings to more than 100 members.

The Vanderbilt University Lablet created a new program to grow the SoS community in 2021 with its Cyber Makerspace initiative. Cyber Makerspace is a simulated, networked environment that facilitates instruction on cyber-physical systems, their security and related topics while reducing cost and complexity. The approach will facilitate reaching audiences from traditionally underrepresented groups.

Details on what the SoS initiative did in 2021 to grow the Science of Security community can be found in Section 3.

The year 2022 marks the beginning of the second decade of the SoS initiative. Over the past ten years the SoS initiative has achieved significant successes and 2022 promises to bring more of the same. In 2022, the SoS initiative will continue to sponsor foundational research at the Lablets and seek to increase the impact of Lablet research on cybersecurity nationwide. SoS personnel will also initiate the process for sponsoring new foundational research projects at academic institutions. Technical areas of interest for the research will be on critical needs that will provide a strategic change in cybersecurity. The Best Scientific Cybersecurity Paper Competition, the 10th Annual, will recognize papers published in 2021 that exemplify the development of scientific rigor in cybersecurity research. HotSoS 2022 will be hosted by the University of Illinois at Urbana-Champaign and held virtually from April 5-7, 2022. HotSoS 2022 will bring together academic, government, and industry researchers from across the growing SoS community and focus on examining scientific foundation of trustworthy systems. The SoS initiative will tackle new challenges in 2022 as it seeks to raise awareness of the need for foundational cybersecurity science to ensure a mature and reliable cyberdefense.



Engaging the Academic Community for Foundational Research

In 2021 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 ongoing pandemic, the Lablets and their Sub-Lablets performed substantive research on 24 projects in 2021 and published 63 peer-reviewed articles or papers bringing to approximately 800 the number of papers published by Lablet researchers since the SoS initiative was established in 2011. The papers have addressed multiple aspects of the five Hard Problems, and have been presented at conferences, symposia, and workshops around the world. As was the case last year, most of the presentations were again virtual. The foundational research embodied by the papers has contributed significantly to enhancing the scientific rigor of research into cybersecurity.

The Principal Investigators (PIs) of the Science of Security Lablets, along with the NSA Research organization, developed five Hard Problems when the SoS initiative was established. These Hard Problems serve as a means of establishing challenging and critical research goals

and 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. This year, SoS community influencers revisited the SoS Hard Problems and their definitions in preparation for a second decade of the NSA SoS Program. HotSoS 2021 included a special breakout discussion session centered around what should constitute the Science of Security Hard Problems going forward. The Hard Problems session consisted of small discussion groups followed by a joint session with summaries from the discussion group moderators. The Summer Lablet Quarterly kicked off with a panel discussion on the Hard Problems, one of the purposes of which was to elicit input from the SoS community. Because AI and ML weren't part of the original Hard Problems and have developed significantly in the past ten years, the SoS community needs to consider how AI and ML can be applied to cybersecurity challenges.

Lablet researchers were again 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 met virtually at their Winter, Summer, and Fall Lablet Quarterlies to hear presentations from invited speakers, present technical papers, and exchange ideas designed to further the Science of Security.

The Hard Problems, Lablet projects and activities, and the Quarterly meetings are described in this section.



Science of Security Hard Problems

Early in 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 for the community. 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 were 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 the 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 current Hard Problems, see the following two documents:

[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>.

[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>.

HotSoS 2021 included a special breakout discussion session centered around what should constitute the Science of Security Hard Problems going forward. SoS community influencers revisited the SoS Hard Problems and their definitions in preparation for a second decade of the NSA SoS Program. The Hard Problems session consisted of small discussion groups followed by a joint session with summaries from the discussion group moderators.

The SoS Lablet Initiative




Leads:

Travis Breaux
Jonathan Aldrich

Sub-Lablets:

George Mason University
Duke University

Hard Problems:

-  Human Behavior
-  Metrics
-  Resilient Architectures

Projects:

- p. 12* Characterizing user behavior and anticipating its effects on computer security with a Security Behavior Observatory
- p. 14* Model-Based Explanation For Human-in-the-Loop Security
- p. 16* Obsidian: A Language for Secure-by-Construction Blockchain Programs
- p. 18* Securing Safety-Critical Machine Learning Algorithms




Lead:

Perry Alexander

Sub-Lablet:

University of Tennessee

Hard Problems:

-  Metrics
-  Scalability and Composability
-  Resilient Architectures

Projects:

- p. 49* Formal Approaches to the Ontology & Epistemology of Resilience
- p. 51* Scalable Trust Semantics & Infrastructure
- p. 53* Secure Native Binary Execution
- p. 55* Side-Channel Attack Resistance





Lead:

Serge Egleman

Sub-Lablets:

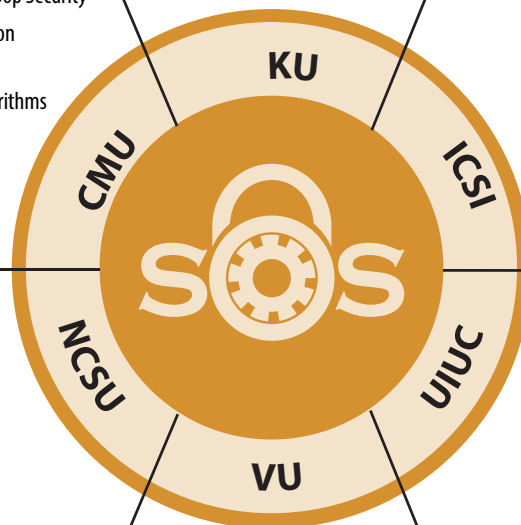
Cornell Tech
UC Berkeley

Hard Problems:

-  Human Behavior
-  Metrics
-  Scalability and Composability
-  Secure Collaboration

Projects:

- p. 21* Designing for Privacy
- p. 23* Governance for Big Data
- p. 25* Operationalizing Contextual Integrity
- p. 27* 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. 30* Coordinated Machine Learning-Based Vulnerability & Security Patching for Resilient Virtual Computing Infrastructure
- p. 32* Development of Methodology Guidelines for Security Research
- p. 34* Predicting the Difficulty of Compromise through How Attackers Discover Vulnerabilities
- p. 36* 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. 58* Analytics for Cyber-Physical System Cybersecurity
- p. 60* Foundations of a CPS Resilience
- p. 62* Mixed Initiative and Collaborative Learning in Adversarial Environments
- p. 64* Multi-model Test Bed for the Simulation-based Evaluation of Resilience





Leads:

Sayan Mitra
David M. Nicol

Sub-Lablets:

University of Arkansas
UT Austin

Hard Problems:

-  Metrics
-  Resilient Architectures
-  Scalability and Composability
-  Secure Collaboration

Projects:

- p. 39* Automated Synthesis Framework For Network Security and Resilience
- p. 41* Monitoring, Fusion, and Response for Cyber Resilience
- p. 43* Resilient Control of Cyber-Physical Systems with Distributed Learning
- p. 45* Uncertainty in Security Analysis


Carnegie Mellon University

In 2021 the Carnegie Mellon University (CMU) Science of Security Lablet continued to focus 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. The CMU Lablet again sponsored the Security Science Research Experience for Undergraduates (REU), funding four students to work with Carnegie Mellon Researchers in Summer 2021.

Principal Investigator (PI) Travis Breaux and Co-PI Jonathan Aldrich lead a team of faculty, postdoctoral and PhD student researchers from CMU and two Sub-Lablets, George Mason University and Duke University. Many of the CMU Lablet faculty and graduate students are affiliated with the CMU CyLab, the coordinating entity for cybersecurity research.

The four REU students were:

- Benito Geordie, Rice University: “Democratizing and Decentralizing Collaborative Web Apps” Advisor: Heather Miller
- Crystal Li, University of Pittsburgh, “User Awareness of Social Media Algorithms” Advisor: Daniel Klug
- Megan Li, Harvey Mudd College, “Usable Consent Interfaces” Advisor: Lorrie Cranor
- Sophia Roshal, Cornell University, “Wyvern: Designing a Next-Generation Programming Language” Advisor: Jonathan Aldrich




Additionally, another six REU students spoke to NSA researchers and associates in August about their summer work, and gave the following talks:

- “picoCTF Cybersecurity and Education Research through Online Gaming” Jenna Bustami, Rachel Nguyen, and Xinyue Lai
- “Enhancing Flexible Science, Technology, Engineering, and Mathematics Thinking by Generating Interactive Diagrams at Scale” Hwei-Shin Harriman
- “Data Structures for Distributed, Federated Machine Learning”, James Flemmin
- “Privacy-Preserving Deep Learning” Allen Marquez, California State University, Los Angeles, Steven Wu
- “Polymorphic Memory Hierarchy” Jennifer Seibert, SUNY Binghamton, Nathan Beckmann
- “User Awareness of Social Media Algorithms” Maya De Los Santos, Northeastern University, Daniel Klug

In mid-July Adam Tagert, SoS Technical Lead, spoke to REU students about careers at NSA in research, and all REU students presented their work to the CMU community at a culminating poster session on August 5.

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

PI:	Lorrie Cranor CO-PI: Nicolas Christin
HARD PROBLEMS:	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 used a multi-purpose observational resource, the Security Behavior Observatory (SBO),

which was developed to collect data from Windows home computers. The SBO collected a wide array of system, network, and browser data from over 500 home Windows computer users (who participated 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. While data collection for this project ended in 2019, the team continues to analyze the dataset and conduct ongoing work on a number of research questions. The research team is also committed to keep the dataset accessible to researchers.

ACCOMPLISHMENTS

For the first paper cited under Publications, we used SBO data to examine 1) how often people read about security incidents online; 2) whether and to what extent they then follow up and take action; and 3) what influences the likelihood that they will read about an incident and take some action.

The second paper cited below used data collected through the SBO to provide new insights into how users browse the internet. We first compared our data to previous studies conducted over the past two decades and identified changes in user browsing and navigation. Most notably, we observed a substantial increase in the use of multiple browser tabs to switch between pages. Using the more detailed information provided by the SBO, we identified and quantified a critical measurement error inherent in previous server-side measurements that do not capture when users switch between browser tabs. This issue leads to an incomplete picture of user

browsing behavior and an inaccurate measurement of user navigation and dwell time. In addition, we observed that users exhibit a wide range of browsing habits that do not easily cluster into different categories, a common assumption made in research study design and software development. We found that browsing the web consumes the majority of users' time spent on their computer, eclipsing the use of all other software on their machine. We showed that users spend the majority of their time browsing a few popular websites, but also spend a disproportionate amount of time on low-visited websites on the edges of the internet. We found that users navigating to these low-visited sites are much more likely to interact with riskier content like adware, alternative health and science information, and potentially illegal streaming and gambling sites. Finally, we identified the primary gateways that are used to navigate to these low-visited sites and discussed the implications for future research.

IMPACT ON HARD PROBLEMS


The Security Behavior Observatory addresses the hard problem of “Understanding and Accounting for Human Behavior” by collecting data directly from people’s own home computers, thereby capturing people’s computing behavior “in the wild”. This data is the closest to the ground truth of the users’ everyday security and privacy challenges

that the research community has ever collected. We expect the insights discovered by analyzing this data will profoundly impact multiple research domains, including but not limited to behavioral sciences, computer security and privacy, economics, and human-computer interaction.

PUBLICATIONS

- Sruti Bhagavatula, Lujó Bauer, and Apu Kapadia, “What breach? Measuring online awareness of security incidents by studying real-world browsing behavior”, European Symposium on Usable Security (EuroUSEC 2021), Online, October 11-12, 2021, and *IEEE Workshop on Technology and Consumer Protection (ConPro 2021)*, Virtual Conference, May 27, 2021.
- Kyle Crichton, Nicolas Christin, and Lorrie Cranor, “How Do Home Computer Users Browse the Web?” To appear in the Feb 2022 issue of the *ACM Transactions on the Web* journal. <https://dl.acm.org/doi/10.1145/3473343>

Model-Based Explanation for Human-in-the-Loop Security

PI:	David Garlan
HARD PROBLEMS:	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

In this year, we worked on the following thrusts:

Game Theoretic Approaches to Self-Protection: Game theory approaches have been explored in security to model malicious behaviors and design reliable defense for the system in a mathematically grounded manner. However, modeling the system as a single player, as done in prior works, is insufficient for the system under partial compromise and for the design of fine-grained defensive strategies where the rest of the system with autonomy can cooperate to mitigate the impact of attacks. To deal with such issues, we propose a new self-adaptive framework incorporating Bayesian game theory and model the defender (i.e., the system) at the granularity of components. Under security attacks, the architecture model of the system is translated into a Bayesian multi-player game, where each component is explicitly modeled as an independent player while security attacks are encoded as variant types for the components. The optimal defensive strategy for the system is dynamically computed by solving the pure equilibrium (i.e., adaptation response) to achieve the best possible system utility, improving the resiliency of the system against security attacks.

To provide exploration capabilities for game-theoretic approaches to self-protection we developed a tool, xGames, that allows operators to (a) visualize and explore games by selecting nodes in the game tree and understanding the state of the game at that point, (b) ask "why", "why not", and "what if" questions about alternative courses of action, (c) understand the impact of games on the system that is affected by moves in the game, and (d) be customizable to arbitrary games and systems. We published a talk on this at: https://www.youtube.com/watch?v=WihcNjPA_fo.

Preparing humans: Informed by work in cognitive science on human attention and context management, we extended our formal framework on reasoning about human-in-the-loop adaptation, to reason about using preparatory notifications in self-adaptive systems involving human operators. Our framework characterizes the effects of managing attention via task notification in terms of task context comprehension. We also built on our framework to develop an automated probabilistic reasoning technique able to determine when and in what form a preparatory notification tactic should be used to optimize system goals.

Explainability of Trade-offs: While recent developments in architectural analysis techniques can assist architects in exploring the satisfaction of quantitative guarantees across the design space, existing approaches in software design are limited because they do not explicitly link design decisions to satisfaction of quality requirements. Furthermore, the amount of information they yield can be overwhelming to a human designer, making it difficult to distinguish the forest through the trees. We developed an approach to analyzing architectural design spaces that addresses these limitations and provides a basis to enable the explainability of design tradeoffs. Our approach combines dimensionality reduction techniques

employed in machine learning pipelines with quantitative verification to enable architects to understand how design decisions contribute to the satisfaction of strict quantitative guarantees under uncertainty across the design space. Our results show feasibility of the approach in two case studies and evidence that dimensionality reduction is a viable approach to facilitate comprehension of tradeoffs in poorly-understood design spaces. This is foundational work that, while focused on software design, is also applicable to explaining run-time decisions when the decision space of possible actions is large, by focusing on the key elements that influence the decision made.

IMPACT ON HARD PROBLEMS

We are addressing resilience by providing defense plans that are automatically generated as the system runs and accounting for current context, system state, observable properties of the attacker, and potential observable operations of the defense. We are addressing

human behavior by providing understandable explanations 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

- Nianyu Li, Mingyue Zhang, Eunsuk Kang, and David Garlan, "Engineering Secure Self-adaptive Systems with Bayesian Games," In *Proceedings of the 24th International Conference on Fundamental Approaches to Software Engineering*, 27 March - 1 April 2021.
- Rebekka Wohlrab and David Garlan, "Defining Utility Functions for Multi-Stakeholder Self-Adaptive Systems," In *Proceedings of the 27th International Working Conference on Requirements Engineering: Foundation for Software Quality*, (Virtual), Essen, Germany 12-15 April 2021.
- Danny Weyns, Bradley Schmerl, Masako Kishida, Alberto Leva, Marin Litoiu, Necmiye Ozay, Colin Paterson, and Kenji Tei, "Towards Better Adaptive Systems by Combining MAPE, Control Theory, and Machine Learning," In *Proceedings of the 16th Symposium on Software Engineering for Adaptive and Self-Managing Systems*, (Virtual), 17-24 May 2021.
- Nianyu Li, Javier Camara, David Garlan, Bradley Schmerl, and Zhi Jin, "Hey! Preparing Humans to do Tasks in Self-adaptive Systems," In *Proceedings of the 16th Symposium on Software Engineering for Adaptive and Self-Managing Systems*, (Virtual), 18-21 May 2021. **Awarded Best Student Paper.**
- David Garlan, "The Unknown Unknowns are not Totally Unknown," In *Proceedings of the 16th Symposium on Software Engineering for Adaptive and Self-Managing Systems*, Virtual, 18-21 May 2021.
- Danny Weyns, Tomas Bures, Radu Calinescu, Barnaby Craggs, John Fitzgerald, David Garlan, Bashar Nuseibeh, Liliana Pasquale, Awais Rashid, Ivan Ruchkin, and Bradley Schmerl, "Six Software Engineering Principles for Smarter Cyber-Physical Systems," In *Proceedings of the Workshop on Self-Improving System Integration*, 27 September 2021.
- Javier Camara, Mariana Silva, David Garlan and Bradley Schmerl, "Explaining Architectural Design Tradeoff Spaces: a Machine Learning Approach," In *Proceedings of the 15th European Conference on Software Architecture*, (Virtual; originally, Vaxjo Sweden), 13-17 September 2021.
- Mohammed Alharbi, Shihong Huang and David Garlan, "A Probabilistic Model for Personality Trait Focused Explainability," In *Proceedings of the 4th International Workshop on Context-aware, Autonomous and Smart Architecture (CASA 2021)*, co-located with the *15th European Conference on Software Architecture*, (Virtual; originally Vaxjo Sweden), 13-17 September 2021.
- Maria Casimiro, Paolo Romano, David Garlan, Gabriel Moreno, Eunsuk Kang and Mark Klein, "Self-Adaptation for Machine Learning Based Systems," In *Proceedings of the 1st International Workshop on Software Architecture and Machine Learning (SAML)*, Springer, (Virtual; originally Vaxjo, Sweden), 14 September 2021.
- Changjian Zhang, Ryan Wagner, Pedro Orvalho, David Garlan, Vasco Manquinho, Ruben Martins, and Eunsuk Kang, "AlloyMax: Bringing Maximum Satisfaction to Relational Specifications," In the *ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE) 2021*, (Virtual), 23-28 August 2021. **Received a Distinguished Paper designation.**

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

PI:	Jonathan Aldrich	CO-PI: Brad Myers
SUB-LABELT:	George Mason University	
HARD PROBLEMS:	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 also developed formative and summative methods for user-centered design of programming languages and applied them to make Obsidian usable. 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. On a technical level, Obsidian applies typestate to express 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 added a second target to the Obsidian compiler. We had previously compiled Obsidian code to Hyperledger fabric to run on open-source Hyperledger blockchains. We are now adding the ability to compile to the Ethereum Virtual Machine (EVM) to run on the more widely used Ethereum blockchains. Our work is in partnership with the Ethereum foundation. Ethereum tracks contract running cost with a unique mechanism called “gas.” Tracking “gas” usage is an interesting and challenging use case for Obsidian’s resource tracking mechanisms. We are currently focusing on allocating objects, which is a key step toward evaluating the gas costs of using Obsidian for smart contracts on Ethereum. After we are able to allocate objects in memory, we will arrange for them to be archived in storage so the ones that need to persist can do so.

can be applied. In addition, it will help us design new language features in a data-driven way.

We built a detector for one known-vulnerable protocol usage pattern. In that pattern, a lockable contract can have its unlock function called when the contract is already unlocked, resulting in a state where ownership of the contract could be improperly reclaimed. Of 4,500 contracts that we ran the detector on, the detector identified 55 as potentially buggy. Manual review showed that 36 of them were vulnerable. So, although we did not identify a general-purpose mechanism for finding bugs in general, we showed that one can develop checkers for specific protocols to find real vulnerabilities, and that some known-vulnerable protocols are used repeatedly on Ethereum.

We are also conducting a study of contracts running on the Ethereum Virtual Machine written in the Solidity language. Our goal is to understand how these contracts track resources, own or share state, and enforce ordering constraints. This will help us understand when and how Obsidian’s existing mechanisms

We compared the gas costs between Obsidian and Solidity. Because the Obsidian compiler is still limited, the tests focus on low-level operations, such as arithmetic, so the results should be considered exploratory. We found

that the Obsidian versions of contracts cost about 10k gas less to deploy (~ \$2 US) and the Obsidian functions cost about 1k gas less to execute (~\$0.20 US). We find this a little surprising, and it merits some additional investigation regarding why, but it is promising so far. It could be that the Yul optimizer, which we use, does a better job than the optimizer that Solidity uses when emitting EVM bytecode.

Obsidian uses a technique called tpestate to capture rules for protocols that users of objects must follow. For example, a file has to be opened before it can be read, and it can only be closed after it

is first opened. We were interested in whether this kind of technique could detect or prevent bugs in real-world smart contracts. To that end, we used the Slithr static analysis framework to build a detector for protocols like the ones Obsidian's type system can capture. On a sample of 100 Solidity contracts from the SmartBugs Wild dataset, the detector identified 71 protocols in 37 contracts. A manual review found that two of these protocols were false positives, leaving 69 protocols among 37 contracts. We think it is promising that over a third of the contracts we analyzed include protocols that Obsidian's type system could capture.

IMPACT ON HARD PROBLEMS

Scalability and composability. Obsidian is designed to enable composition of mutually distrusting programs on scalable blockchain platforms. It has a specific focus on preventing specific composition-related vulnerabilities like invalid reentrancy 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


The Obsidian project has partnered with support from the Ethereum Foundation. Obsidian currently supports the Hyperledger Fabric blockchain platform. We are working on building a proof-of-concept

version of Obsidian for Ethereum. The ultimate goal is to make Obsidian a viable alternative to Solidity for Ethereum developers so that Ethereum users can obtain the usability and security benefits of using Obsidian.

PUBLICATIONS

- Michael Coblenz, Jonathan Aldrich, Brad Myers, and Joshua Sunshine, "Can Advanced Type Systems be Usable? An Empirical Study of Ownership, Assets, and Tpestate in Obsidian," In Proceedings of the ACM on Programming Languages (OOPSLA), 2020. **Distinguished Artifact Award**
- Michael Coblenz, Reed Oei, Tyler Etzel, Paulette Koronkevich, Miles Baker, Yannick Bloem, Brad Myers, Joshua Sunshine, and Jonathan Aldrich, "Obsidian: Tpestate and Assets for Safer Blockchain Programming," ACM Transactions on Programming Languages and Systems (TOPLAS), 2020.
- Michael Coblenz, Gauri Kambhatla, Paulette Koronkevich, Jenna Wise, Celeste Barnaby, Jonathan Aldrich, Joshua Sunshine, and Brad A. Myers, "PLIERS: A Process that Integrates User-Centered Methods into Programming Language Design," ACM Transactions on Computer-Human Interaction (TOCHI), 2021.

Securing Safety-Critical Machine Learning Algorithms

PI:	Lujo Bauer CO-PI: Matt Fredrikson
SUB-LABELT:	Duke University
HARD PROBLEMS:	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

The paper presented at Asia CCS extends our previous arXiv preprint with some new results: we had previously shown that malware binaries could often be transformed so that they evaded correct classification by anti-virus programs (i.e., they would be incorrectly classified as benign). Leveraging an expanded experimental infrastructure, we more recently showed that such attacks can ultimately succeed even when attempting to transform binaries that initially appear resistant to attack. Specifically, we recognized that previous attacks are sufficiently stochastic that even when they usually fail, a determined adversary who attempts enough attacks will eventually succeed with high probability.

We continued along the same directions as previously, including investigating more powerful attacks on malware classifiers, and explaining malware classifier decisions to better understand the potential strengths and weaknesses of different malware classifiers.

We began studying safe ordering properties, a class of non-relational safety properties that capture a broad range of safety concerns for classification models. Our current work in this direction is aimed at transforming existing classifiers into variants that provably satisfy a given set of safe ordering properties. Key concerns that we aim to address are runtime efficiency and preservation of accuracy.

We have set up an experimental infrastructure and have been experimenting with using Fast Gradient Sign Method (FGSM) for adversarial training for malware detectors. Although using FGSM to create malware variants wouldn't produce functional malware (because the attack is not aware of the semantics of a binary), we are exploring whether it could nevertheless lead classifiers to become more robust. We have also been developing additional experimental infrastructure to use our previous, semantics-aware attacks for adversarial training of malware classifiers.

IMPACT ON HARD PROBLEMS

Both of the Hard Problems this project addresses are tackled in the context of deep neural networks, which are a particularly popular and performant type of machine learning algorithm. This project develops metrics that characterize the degree to which a neural-network-based classifier can be evaded through practically realizable, inconspicuous attacks. The project also develops architectures for neural networks that would make them robust to adversarial examples. The framework for explaining the predictions made by deep neural networks may

identify the network-internal factors that cause misclassifications, and we leverage this capability to make progress on the Hard Problems. Finally, as we examine classifiers' robustness to these attacks, we are analyzing the effect that increased robustness may have on other information security metrics, such as those that characterize the confidentiality of training data. Our results have begun to shed new light on the tradeoffs that emerge when certain defensive tactics are employed.

EDUCATION AND OUTREACH

- Project participants gave the following presentations:
 - "On the practical risks and benefits of AI to security," keynote, 5th Italian Conference on Cybersecurity (ITASEC'21). Apr. 8, 2021
 - "Beyond Ip balls: Attacks on real-world uses of machine learning," Workshop on Adversarial Machine Learning in Real-World Computer Vision Systems / CVPR workshop, Jun. 19, 2021 and the SoS Summer Quarterly Label meeting, July 13, 2021.
- PI Lujo Bauer participated in the International School on Foundations of Security Analysis and Design (FOSAD), in Bertinoro, Italy, and presented "Attacks on real-world uses of machine learning", which included content developed under SoS sponsorship, Sep. 2-3, 2021

PUBLICATIONS

- Keane Lucas, Mahmood Sharif, Lujo Bauer, Michael Reiter, and Saurabh Shintre, "Malware Makeover: Breaking ML-based Static Analysis by Modifying Executable Bytes," In Proceedings Asia CCS (Virtual) June 7-11 2021.

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 four research projects.


The four research projects are identified below:

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



PI Serge Egelman

Designing for Privacy

PI:	Serge Egelman
HARD PROBLEMS:	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

We conducted a narrative literature review to understand the factors that have impact on developers’ adoption of privacy and security practices in product design and software development. On the basis of this survey, we developed a model that categorizes the factors that affect developers’ decision-making: from the environmental level (context outside of organization or team) to organizational, product and development process-related levels, and finally, personal levels.

With students at UC Berkeley, we evaluated the familiarity of smartphone users with privacy and security settings, their expectations about their ability to configure those settings, their understanding of the privacy and security threats against which the settings are supposed to protect them, and their expectations about the effectiveness of the settings. In an online survey with 178 users with diverse backgrounds and demographics, we found that many people were not aware of smartphone privacy/security settings and their defaults, and had not configured them in the past, though they expressed willingness to do it in the future. Some participants perceived low self-efficacy and expected difficulties and usability issues with configuring those settings. We compared the findings across various socio-economic groups of participants to draw conclusions about what groups are especially vulnerable to the identified issues. We found that, compared to so-called “average users”, certain user groups, such as older adults, racial/ethnic minorities, and females, were less concerned about online privacy and security, engaged less in configuring smartphone privacy and security settings, and expected more difficulties with configuring them and more negative impacts on user experience. But even the “average users” in that survey expressed low levels of awareness and

engagement, and some concerns about their expected difficulties with configuring smartphone privacy/security settings. Building on these findings, we are conducting a new study to investigate whether those self-reported expectations are confirmed by behavioral observations. The goals of the new study are to better understand what difficulties users actually face when configuring smartphone privacy settings (specifically, finding and enabling them), where negative user experiences may stem from, what users understand about the implications of changing those settings, and how changes to the design of settings interfaces could make configuration easier. As with the survey, we are particularly interested in examining potential differences among users with different socio-economic backgrounds, and the implications of those differences for design.

We conducted cognitive walkthrough interviews with a demographically diverse sample of iOS and Android users. We asked the participants to talk through configuring three smartphone privacy settings, then asked follow-up questions about difficulty and clarity. We are currently analyzing the data from the walkthroughs and follow-up interviews.

We are continuing to analyze and thematically code interviews with nannies and au pairs about their experiences with, and views of, working in homes with cameras and smart home devices. Our analysis particularly focuses on identifying the most feasible points of intervention for improving domestic employees’ control over the privacy effects of such devices, whether technical controls or in terms of negotiations between nannies and employers, such as guidelines and education (see Education and Outreach). For more details on this research, see the project Operationalizing Contextual Integrity.

IMPACT ON HARD PROBLEMS

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

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


EDUCATION AND OUTREACH

- Based on the studies with nannies mentioned above, we are working with colleagues at the University of Oxford to help them develop workshop content for events for domestic workers about their privacy rights and options.
- The study on smartphone settings was conducted in collaboration with several student researchers at UC Berkeley (Masters students and undergrads), providing research experiences and contributing to a publication track record early in their academic careers.

PUBLICATIONS

- Several papers based on the studies and surveys described above are in preparation

Governance for Big Data

PI:	Serge Egelman
SUB-LABELT:	University of California, Berkeley
HARD PROBLEMS:	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

Software development teams are responsible for making and implementing software design decisions that directly impact end-user privacy, a challenging task to do well. Privacy Champions within organizations--people who strongly care about advocating for privacy--play a useful role in supporting privacy-respecting development cultures. To understand their motivations, challenges, and strategies for protecting end-user privacy, we conducted 12 interviews with Privacy Champions in software development teams. We found that common barriers to implementing privacy in software design include: negative privacy culture, internal prioritization tensions, limited tool support, unclear evaluation metrics, and technical complexity. To promote privacy, Privacy Champions regularly use informal discussions, management support, communication among stakeholders, and documentation and guidelines. They perceive code reviews and practical training as more instructive than general privacy awareness and on-boarding training. Our study is a first step towards understanding how Privacy Champions work to improve their organization's privacy approaches and improve the privacy of end-user products.

With colleagues at the University of Edinburgh, we designed and conducted a study of how the privacy information and choices presented to mobile app developers influence their choices about integrating personalized vs. non-personalized ads into their apps. Mobile advertising networks present personalized advertisements to developers as a way to increase revenue. These types of ads use data about users to select potentially more relevant content. However, choice framing also impacts app developers' decisions which in turn impacts their users' privacy. Currently, ad networks provide choices in developer-facing dashboards that control the types of information collected by the ad network as well as how users will be asked for consent. Framing and nudging have been shown to impact users'

choices about privacy; we anticipate that they have a similar impact on choices made by developers. We conducted a survey-based online experiment with 400 participants with experience in mobile app development. Across six conditions, we varied the choice framing of options around ad personalization. Participants in the condition where privacy consequences of ads personalization are highlighted in the options are significantly (11.06 times) more likely to choose non-personalized ads compared to participants in the Control condition with no information about privacy. Participants' choice of ad type is driven by impact on revenue, user privacy, and relevance to users. Our findings suggest that developers are impacted by interfaces and need transparent options.

With colleagues at University of Bristol, we are initiating a research study to examine how developers of mobile health apps approach privacy for those apps, including app developers' expectations and beliefs about legal protections for health data, technical possibilities for protecting privacy, and users' expectations and preferences. As a first step, we are analyzing posts on StackOverflow to develop an initial sketch of some of the technical concerns and issues developers of health apps have about the privacy of the data they are dealing with, and who or what is driving those issues or concerns. For example, many questions in the last few years are related to the permissions requirements for the health-specific APIs provided by mobile platforms (Apple's HealthKit, Google Fit, Samsung Health). Our next step will be to conduct interviews with app developers to broaden our view of how they approach health data. We are continuing our analysis of how developers of mobile health apps approach privacy for those apps, wrapping up our analysis of StackOverflow posts, examining other textual sources, and turning to designing a human subjects study with members of health app product teams.

We are designing a vignette-style survey study to analyze users' privacy expectations about consumer health apps and what those expectations are based on, comparing across different types of apps that users may perceive as being more and less likely to be constrained by laws that regulate collection and sharing of health data. (For example, telemedicine apps vs. workout trackers vs. dating apps for people with certain medical conditions.)

With colleagues at Aalto University in Finland, we are preparing to conduct a related study with healthcare professionals on their views of consumer health apps. We have developed the study design

and submitted it for ethics review by Aalto University. Our research questions include how these views and expectations on collection of health data by consumer health apps are related to policies and governance structures for health data in medical settings. We therefore plan to interview healthcare professionals in several countries (the U.S., Finland, Sweden, Sri Lanka, and Singapore) that are quite different in terms of legal and ethical protections for medical data, the degree of centralization and control in the systems where patients' medical data resides, and the availability of government-sponsored health self-management platforms and apps.

IMPACT ON HARD PROBLEMS


A new data governance approach focuses on accountability and relates more to accounting and auditing. The first step is to develop a design methodology from all different approaches and mechanisms, and then

validate the design methodology by working with practitioners and building case studies for generalizable design patterns.

PUBLICATIONS

- Mohammad Tahaei, Alisa Frik, and Kami Vaniea, "Privacy Champions in Software Teams: Understanding Their Motivations, Strategies, and Challenges," In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, ACM, May 2021.
- Mohammad Tahaei, Alisa Frik, and Kami Vaniea, "Deciding on Personalized Ads: Nudging Developers About User Privacy," Symposium on Usable Privacy and Security (SOUPS), USENIX, August 9-10, 2021. <https://doi.org/10.7488/DS/3045>.

Operationalizing Contextual Integrity

PI:	Serge Egelman CO-PI: Helen Nissenbaum
SUB-LABEL:	Cornell Tech
HARD PROBLEMS:	Scalability and Composability, Policy-Governed Secure Collaboration, Security Metrics and Models, Human Behavior 

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

The sharing of information between older adults and their friends, families, caregivers, and doctors promotes a collaborative approach to managing their emotional, mental, and physical well-being and health, prolonging independent living and improving care quality and quality of life in general. However, information flow in collaborative systems is complex, not always transparent to elderly users, and may raise privacy and security concerns. Because older adults' decisions about whether to engage in information exchange affects interpersonal communications and delivery of care, it is important to understand the factors that influence those decisions. While a body of existing literature has explored the information sharing expectations and preferences of the general population, specific research on the perspectives of older adults is less comprehensive. Our work contributes empirical evidence and suggests a systematic approach. We presented the results of semi-structured interviews with 46 older adults age 65+ about their views on information collection, transmission, and sharing using traditional ICT and emerging technologies (such as smart speakers, wearable health trackers, etc.). Based on analysis of this qualitative data, we developed a detailed model of the contextual factors that combine in complex ways to affect older adults' decision-making about information sharing. We also discussed how our comprehensive model compares

to existing frameworks for analyzing information sharing expectations and preferences. Finally, we suggested directions for future research and described practical implications of our model for the design and evaluation of collaborative information-sharing systems, as well as for policy and consumer protection. Specifically, we proposed ways that product developers could use the model to identify and mitigate potential privacy impacts of their products, such as prompts based on the individual factors. The increasing use of smart home devices affects the privacy not only of device owners, but also of individuals who did not choose to deploy them, and may not even be aware of them. Some smart home devices and systems, especially those with cameras, can be used for remote surveillance of, for example, domestic employees. Domestic workers represent a special case of bystanders' privacy, due to the blending of home, work, and care contexts, and employer-employee power differentials. To examine the experiences, perspectives, and privacy concerns of domestic workers, we begin with a case study of nannies and of parents who employ nannies. We are analyzing the transcripts of 26 interviews with nannies and au pairs about their experiences with smart devices in their employers' homes, and their attitudes, expectations, and preferences with regard to data collection in smart homes. Our goal in this case study is to examine what factors in this combined home/workplace/

caregiving context impact employers' and employees' data-sharing choices, and how such choices and attitudes reflect or change power dynamics in their relationships. We also aim to identify potential points of intervention (technical and social) for better respecting bystanders' privacy preferences. We are planning a research agenda to integrate the concerns of bystanders into our work with smart home product developers, including experimental interventions to prompt more attention to the issue in design.

We started designing a study of in-home personal assistants, privacy, and the types of features that people would like to use (while balancing those privacy concerns). We decided to apply to the "Experience Sampling Method" (ESM) to periodically survey people about the conversations that they just had within their homes, and their level of comfort with having those conversations disclosed to devices and services. This involved creating a software prototype to randomly survey people on their phones. A related study involves showing people snippets of human conversations and asking them their level of comfort of having either humans or machines access those conversations for various purposes. The goal is to identify the types of topics and cues that could be used by classifiers to automatically apply privacy protections for in-home data capture devices (e.g., future smart assistants). We are working to use this data as a training set: in this stage, we're using human coders to code the ground truth data, that will later be used by the classifier.

We performed a study that collected people's perceptions of passive listening, their privacy preferences for it, their reactions to different modalities of permission requests, and their suggestions

for other privacy controls. Based on our results, we created a set of recommendations for how users should be presented with privacy decisions for these and other future in-home data capture devices. Our study of passive-listening devices used an interactive app store experience that provided a unique means of measuring consumer sentiment in a scenario modeling real life. Using both quantitative and qualitative analysis, we determined people's views on privacy models for always-listening voice assistants, which generally ranged from an outright rejection of the voice assistant described in our survey to preferring one model for its increased privacy protections. Based on our prior work on providing privacy controls for always-listening devices that leverage CI theory, we observed that users desire audit mechanisms, so that they can examine what decisions have been previously made about their privacy. We showed that providing users with feedback and examples of the types of data apps may collect is an effective method for helping them detect malicious apps that may cause privacy violations. We also show that these techniques are likely applicable to other domains that rely on machine learning and which offer opportunities to decompose larger problems into smaller, self-contained tasks that are amenable to human verification. To find out how people would react to different kinds of runtime permission requests, we are asking participants in this study to hold conversations while getting ambient suggestions (and plenty of permission requests) from a passive listening assistant, which we will simulate in real time using the Wizard of Oz technique. Most of our participants seem to be excited about passive listening, but want control over the assistant's actions and their own data. They generally seem to prioritize an interruption-free experience above more fine-grained control over what the device is allowed to record.


IMPACT ON HARD PROBLEMS

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

PUBLICATIONS

- Alisa Frik, Julia Bernd, Noura Alomar, and Serge Egelman, "A Qualitative Model of Older Adults' Contextual Decision-Making About Information Sharing," *Workshop on the Economics of Information Security (WEIS'20)*, December 2020.
- Julia Bernd, Ruba Abu-Salma, and Alisa Frik, "Bystanders' Privacy: The Perspectives of Nannies on Smart Home Surveillance," *2020 USENIX Workshop on Free and Open Communications on the Internet (FOCI '20)*.

Scalable Privacy Analysis

PI:	Serge Egelman CO-PI: Narseo Vallina-Rodriguez
HARD PROBLEMS:	Scalability and Composability, Policy-Governed Secure Collaboration, Security Metrics and Models 

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

Mobile Operating Systems (OSes) are mostly locked down to ensure the interactions between the user, the apps, and the OS remain tightly controlled by the OS for security. Some mobile devices, however, can be altered, physically or through software means, to run altered OSes that allow users and some apps to have higher privileges than they normally would, breaking the locked-down isolated model. So-called "rooted" or "jailbroken" mobile devices are not only popular among users who like to tinker with and customize their devices outside the realm of what is sanctioned or even allowed by manufacturers, but also a staple in the arsenal of researchers and app store owners who use them to analyze mobile apps for a variety of purposes. These include

privacy and security research, app store policy enforcement, and anti-piracy and law enforcement. This is why it is a well-known and well-studied fact that various strains of mobile malware look for root access (perhaps to try and exploit it when found) or detect analysis tools to ensure they are not running on a monitored device or a honey-pot. Regardless of their intent, the use of root-detection and anti-analysis techniques hinders the ability of researchers to accurately study these apps and the ability of app stores to effectively enforce their policies, as both groups rely on analysis tools that are detected and flagged by apps that use root detection and anti-analysis techniques. Despite their importance and huge potential impact on mobile app research,

our knowledge of these techniques used in legitimate mobile apps (and not just mobile malware) is limited at best. To the best of our knowledge there is no comprehensive study of the usage of root detection and anti-analysis techniques in non-malware mobile apps. As a result, we set out to study these techniques and their popularity among mobile apps. Our work has so far focused on studying anti-analysis techniques in Android applications. These are programming methods, tools, and other developer choices that, intentionally or otherwise, make it more difficult for researchers to analyze mobile applications. Some notable examples are root detection, emulation detection, debugging environment detection, certificate pinning, using custom non-standard encryption, code obfuscation, and other techniques that end up negatively impacting legitimate privacy and security research. Our goal is to identify anti-analysis techniques and understand the reason behind their adoption within the context of the applications. Specifically, we have designed and implemented tools to taxonomize root, emulation, and debugging environment detection methods; devised counter-measures to these methods so that we can test apps in our dynamic analysis environment; and planned a measurement study around our findings to study how app behavior changes under different conditions. We are in the process of analyzing our testing data from ~10k apps.

Using our network traffic instrumentation, we're performing another study to examine how ad networks process personal information and whether they obey opt-out flags to disable behavioral targeting. Using our instrumentation, we've reverse-engineered several ad network APIs, which will allow us to conduct real-time actions that

can be used for controlled experiments (e.g., examining the impact of gender information by examining how ads and their associated pricing change as a function of the user's specified gender). In a similar manner, we plan to examine whether various ad networks are obeying flags for complying with various privacy laws or disabling behavioral targeting. Our instrumentation allows us to see all bid data from ~30 ad networks, and so by creating ad auctions by spoofing app traffic, we can perform controlled experiments to manipulate elements of the request to examine their impact on the dependent variables, which are the amounts of each bid and the types of ads offered.

We began several studies this year to examine developers' perspectives, since one of the key issues that we've uncovered is that most privacy issues in mobile apps are due to misbehaving or misconfigured Software Development Kits (SDKs). Thus, we're using AppCensus data to identify apps with questionable privacy behaviors (including specifically looking at children's apps, since that's a regulated area), and then inviting their developers to interviews/surveys about their compliance practices. We recently surveyed developers of kids' apps about their compliance with child privacy laws, and are in the process of analyzing the results. Our goal is to understand compliance processes. Ironically, while planning to study compliance processes, an IRB snafu halted our surveys for ~2 months. We've since resumed: we're surveying app developers about their privacy compliance processes. To date, we've collected >50 survey responses, and have now moved on to follow-up interviews. The main finding is that developers aren't aware of their compliance obligations and look to platforms for guidance; at the same time, they don't believe that the platforms are providing adequate guidance.

IMPACT ON HARD PROBLEMS

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

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

EDUCATION AND OUTREACH

- PI Egelman testified before the U.S. Senate on mobile privacy, specifically apps' compliance with COPPA and how COPPA could be improved. He has also been interviewed about online privacy issues.

PUBLICATIONS

- Several papers based on the research above are in preparation.

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 four research projects.

We continue to bring up the Science of Security in a variety of fora, including:

- Presentations at and discussions with colleagues at academic conferences, including the Deep Learning Cyberthreat Intelligence workshop at the International Conference on Data Mining.
- We hosted two distinguished lectures on Ethics and Safety of AI that involved themes relating to Science of Security; one lecture was on regulations and risk and the other on safety envelopes as a way to provide guarantees on behavior in a tractable manner.
- We regularly hold discussions with non-Lablet colleagues locally and at other universities.
- We conducted three Secure Software Supply Chain summits; two of the summits were with industrial organizations involving 24 organizations, and one summit was with five government organizations.



PI Laurie Williams




CO-PI 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
- Predicting the Difficulty of Compromise through How Attackers Discover Vulnerabilities
- Reasoning about Accidental and Malicious Misuse via Formal Methods

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

PI:	Xiaohui (Helen) Gu
HARD PROBLEMS:	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 continued to study a hybrid learning framework that combines our previous CDL (Classified Distributed Learning) solution with a supervised learning model to further improve our attack detection accuracy. We also further studied how to identify vulnerable code patterns automatically by analyzing code control flows to extract call

paths that can reach vulnerable Java library functions or infinite loops. We developed a Hybrid Machine Learning (HML) approach to detecting security attacks in containerized applications and are refining the system based on feedback.


IMPACT ON HARD PROBLEMS

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

PUBLICATIONS

- Several papers based on the studies described above are in preparation.

Development of Methodology Guidelines for Security Research

PI:	Jeff Carver
SUB-LABELT:	University of Alabama
HARD PROBLEMS:	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 cybersecurity community. To ensure that the guidelines are widely usable and accepted, it is important that we included experts from different parts of the cybersecurity 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 cybersecurity research community.

ACCOMPLISHMENTS

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 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 their analysis that would have otherwise been missed from the traditional cybersecurity settings. To this end, we continue to work towards a better and more versatile rubric or set of guidelines as they 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. The 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 challenges 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 the 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, and would like to get feedback on this initial draft either through a workshop or through some other types of interactions. The feedback would include 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 the conclusions and provide any additional perspectives that may have been missed.

We also did a new study that supplemented the interview data by analyzing the comments left by reviewers on submissions to the 2020 HotSoS symposium. We are integrating the results of this analysis into the guideline development process. Because the HoTSoS conference is focused on the science of security, the content of these reviews helps to better understand the types of information science of security experts (the HotSoS reviewers) find important when reviewing cyber

security research. This data will be used to further verify and support the guidelines created from the interview data. For the guidelines for scientific reporting in cyber security, we have refined the thematic groups which categorize the data gathered from the interviews. These thematic groups are integral to the work in finalizing the first version of the guidelines for improving scientific rigor and validity in cyber security reporting.

IMPACT ON HARD PROBLEMS


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.

PUBLICATIONS

- The Good Examples article is based on past SoS work and showcases papers listed as "Notable SoS Papers" (<https://cps-vo.org/node/61173>).

Predicting the Difficulty of Compromise through How Attackers Discover Vulnerabilities

PI:	Andy Meneely CO-PI: Laurie Williams
SUB-LABEL:	Rochester Institute of Technology
HARD PROBLEMS:	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

Based on our data analysis of the 2019 Collegiate Penetration Testing Competition, we found that: 1) vulnerabilities related to protection mechanism failure (e.g. lack of SSL/TLS) and improper neutralization (e.g. SQL injection) are discovered faster than others; 2) vulnerabilities related to protection mechanism failure and improper resource control (e.g. user sessions) are discovered more often and are exploited more easily than others; and 3) there is a clear process followed by penetration testers of discovery/collection to lateral movement/preattack.

We are performing comparison and evaluation of existing vulnerable dependency detection tools. The goal of this study is to aid security practitioners and researchers in understanding the current state of vulnerable dependency detection through a comparative study of existing tools. We ran 10 industry-leading dependency detection tools on a large web application composed of 44 Maven (Java) and npm (JavaScript) projects and found that the tools' results vary for both vulnerable dependencies and the unique vulnerabilities.

We are building on a natural language classifier to mine apology statements from software repositories to systematically discover self-admitted mistakes. This classifier is being applied to a random sampling of GitHub repositories, including language from commits, issues, and pull request conversations. Thus far, we have collected data

from 17,491 repositories, which are the top 1000 ranked repositories from 54 different programming languages. We are scaling up our data collection as well as working on apology results.

We completed a study comparing the output of Software Component Analysis (SCA) tools, a tool type getting increased attention with the Executive Order on Cybersecurity's emphasis on software supply chain. Our manual analysis of the tools' results suggests that the accuracy of the vulnerability database is a key differentiator for SCA tools. We recommend that practitioners should not rely on any single tool at the present as that can result in missing known vulnerabilities.

We conducted a study to aid software practitioners and researchers in understanding the current practice of releasing security fixes by open source packages through an empirical measurement study. Specifically, the study focused on the fix-to-release delay, code change size, and documentation of security releases over 4,377 advisories across seven open-source package ecosystems. We found that packages are typically fast and light in their security releases as the median release comes under 4 days of the corresponding fix and contains 402 lines of code change. Furthermore, we found that 61.5% of the releases come with a release note that documents the corresponding security fix, while 6.4% of these releases also mention a breaking change.

IMPACT ON HARD PROBLEMS

We are developing a new set of metrics for measuring exploitability using the attack surface. These metrics are based on the behavior observed by penetration testers in a competition environment. The intrusion detection data collected from the CTPC have provided us

with detailed timelines of how attackers find, exploit, and pivot with vulnerabilities. When studying how they work with the known attack surface, we will develop metrics that show which vulnerabilities are at highest risk based on the current deployment.


EDUCATION AND OUTREACH

- Jacob Woolcutt has put the NCSU research team in touch with Stephen Magill at Sonatype to establish further collaboration.

PUBLICATIONS

- Nasif Imtiaz, Seaver Thorn, and Laurie Williams, "A comparative study of vulnerability reporting by software composition analysis tools," *15th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM 2021)*.
- Nikolaos Alexopoulos, Max Muhlhauser, and Andrew Meneely, "Who are Vulnerability Reporters? A Large-scale Empirical Study on FOSS," *15th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM 2021)*.
- Saikath Bhattacharya, Munindar Singh, and Laurie Williams, "Software Security Readiness and Deployment," In *Proceedings, 2021 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW), ISSREW 2021*.

Reasoning about Accidental and Malicious Misuse via Formal Methods

PI:	Munindar Singh CO-PIs: William Enck, Laurie Williams
HARD PROBLEMS:	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 proposed Unexpected-Catch, a framework for identifying mobile apps that enable information access about users and others that violate user expectations. We term such apps UIA (Unexpected Information Access)-enabling apps. Our framework identified 83 UIA-enabling apps from the seed dataset and found an additional 48 UIA-enabling apps via snowballing. We manually verified that 53 of the initial 83 and 32 of the additional 48 apps are truly UIA-enabling. That is, we obtained a higher percentage of true positives than in an earlier study. We evaluated our framework based on app reviews. We created a gold dataset of 100 apps that contains 62 rogue apps. Our method an F1 score of 76% on this dataset, beating the previous study (which leveraged app descriptions), which achieved an F1 score of 63%. In this setting, false negatives carry a high risk and are thus costlier than false positives. Our approach achieved 89% recall at 66% precision, whereas the previous study achieved 52% recall at 82% precision.

We proposed a framework for analyzing story structures as sequential patterns of event types in app reviews. App users write about different user-app interaction stories in app reviews to achieve different purposes, such as reporting bugs and expressing unmet expectations. Our proposed framework enables analysts and developers to search for stories based on their structures.

Healthcare professionals and patients need HIPAA-compliant mobile apps to store, seek, or communicate about private health information. However, not all available apps that serve such purposes are HIPAA compliant, and information about apps' HIPAA compliance is not always available. We collected the app descriptions, privacy policies, and developers' webpages of more than 30,000 medical apps on Apple App Store, and are currently exploring a methodology for choosing trustworthy apps that healthcare professionals and patients can use.

We continued our analysis of Payment Service Provider (PSP) application programming interfaces (APIs), identifying a set of relevant criteria from the M-ASVS and mapping them to program analysis tasks. We also adapted an existing application-based data flow analysis framework to operate on a PSP library file. We built a dataflow-based static program analysis tool to study how Payment Service Provider (PSP) libraries for mobile Android apps store security critical information.

We conducted a comparison study of three Natural Language Processing (NLP)/Machine Learning (ML) models for extracting attacker techniques from CTI.

We enhanced our work on norm extraction from breach reports. Specifically, whereas our previous work leveraged crowdsourcing, we focused on automated extraction of useful actions and descriptive

phrases from breach reports. Based on the extracted information, we built a tool for action suggestions based on breach descriptions.

IMPACT ON HARD PROBLEMS

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.

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 University of Arkansas and the University of Texas at Austin are engaged in four 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 cyberattacks. The Lablet's work draws on several fundamental areas of computing research as well as ideas from other the mathematics and engineering disciplines.

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

- 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



An Automated Synthesis Framework for Network Security and Resilience

PI:	Matt Caesar CO-PI: Dong (Kevin) Jin
SUB-LABELT:	University of Arkansas
HARD PROBLEMS:	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. Veriflow is a startup company commercializing verification technology that came out of this project's SoS Labelt funding. The collaborations in this year targeted enhancement of our verification technology to operate on real-time traffic data, as well as development of a "high-speed" variant of our approach that can perform verification and quickly answer queries on large environments while requiring only small footprints in terms of memory and CPU. We worked on approaches to parallelize computations, making them amenable to deployment across clouds and other decentralized environments. We are in the process of conducting performance evaluation.

We studied the interdependence between the power system and the communication network to improve resilience in critical energy infrastructures, which addresses the resilient architecture hard problem. We proposed a two-layer distribution system model with both power and communication components. Based on the model, we formulated the Distribution Service Restoration (DSR) process as a routing problem and developed a simulation-based method to quantitatively evaluate the DSR process on large-scale power systems (e.g., IEEE 123-node system and Ckt-7 system). The experimental results show that our method improves the total restored energy up to 57.6% and reduces the recovery time up to 63% by considering the power-communication interdependency. We also added a dysfunctional switch model to ensure system stability under new operational constraints (such as voltage and capacity limits), and formulated a stochastic optimization model

to address communication uncertainties.

We developed a testing platform for cyber-physical system resilience and security evaluation, which addresses the resilient architecture hard problem. The platform consists of a container-based network emulator, network/power system simulator, and real hardware for rapid prototyping of network applications with high fidelity and scalable testing environment. To overcome the statistical error in virtual time advancement within the platform due to process waiting time including disk I/O time, network I/O time and GPU computational time, we formulated an analytical model of virtual time, proposed a time compensation mechanism, and implemented it in the Linux kernel to precisely control time advancement by considering the non-CPU task waiting time. We conducted extensive experiments for error analysis and system evaluation.

We developed a general and interpretable framework for analyzing PMU data in real-time; the proposed framework enables grid operators to understand changes to the current state and to identify anomalies in the PMU measurement data. We first learn an effective dynamical model to describe the current behavior of the system by applying statistical learning tools on the streaming PMU data. We use the probabilistic predictions of our learned model to principally define an efficient anomaly detection tool. Our framework produces real-time classification of the detected anomalies into common occurrence classes. We demonstrate the efficacy of our proposed framework through numerical experiments on real PMU data collected from a transmission operator.

We continued our collaboration with Boeing on constructing a resilient IoT platform for the battlefield. We are exploring an approach that leverages deep learning to dynamically relocate drone-mounted access points to evade the adversary. We formulated a placement algorithm that leverages Model-Agnostic Machine Learning (MAML) to construct an algorithm resilient to an adversary attempting to disrupt the learning process. We have developed new deep learning mechanisms that are resilient to data sets that are “constructed” by adversaries, and our early

simulation results show benefits to these approaches in practical settings.

We have developed a design and evaluation framework for a self-driving “service provider infrastructure” that leverages our prior work on verification and synthesis to automatically self-configure to become resilient to attacks. Our initial focus is on network and container orchestration systems, and our first implementation targets Kubernetes. Our platform leverages AI planning algorithms to synthesize steps the system needs to take to protect itself against incoming attacks from an intelligent adversary.

IMPACT ON HARD PROBLEMS

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

Scalability and Composability. Our unique emulation/simulation-integrated testing platform enables scalable

evaluation methodologies, and the experimental results indicate the approach scales to large operational environments while providing formal guarantees on correctness of the cyber domain.

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


EDUCATION AND OUTREACH

- Matthew Caesar co-founded and serves on the organizing committee of theNetworkingChannel, an online channel to discuss topics related to computer networking, systems, and security.
- Matthew Caesar served as a Juror in the ACM SIGCOMM Student Research Competition, 2021.
- Matthew Caesar served as an invited panelist in the 39th Brazilian Symposium on Computer Networks and Distributed Systems (2021).
- Matthew Caesar was elected as the Vice Chair for ACM SIGCOMM, and will serve a four-year term. In his position, he will be responsible for leading initiatives in the SIGCOMM community, with an emphasis on education and cybersecurity.
- Matthew Caesar serves as the Sponsor Chair for ACM SIGCOMM 2022.
- Matthew Caesar serves on the Program Committee for USENIX NSDI 2022.
- Kevin Jin organized a track on Dynamic Data-Driven Application Systems for the 2021 INFORMS Annual Meeting and served as a panelist in the track.
- Kevin Jin serves as the Program Co-chair for ACM SIGSIM-PADS 2022.
- Kevin Jin developed a new graduate-level class, CSCE5655 Network Security, for the University of Arkansas Global Campus. The class is being offered in Spring 2022.
- Yanfeng Qu and Kevin Jin gave an NSA seminar talk “Cyber-Resilience Enhancement of PMU Networks Using Software-Defined Networking” in March 2021.
- Matthew Caesar, Kevin Jin, and Gabriella Xue gave a talk “Automated Synthesis Framework for Network Security and Resilience” in the SoS Labet quarterly meeting, November 2021.
- Xiaoliang Wu, a Ph.D. student of Kevin Jin, graduated in December 2021, and will join Facebook, working on data center network design and performance evaluation.

PUBLICATIONS

- Otto Piramuthu and Matthew Caesar, “How Effective are Identification Technologies in Autonomous Self-Driving Vehicles?” IEEE CommNet, December 2021.
- Bingzhe Liu, Kuan-Yen Chou, Pramod Jamkhedkar, Bilal Anwer, Rakesh Sinha, Kostas Oikonomou, Matthew Caesar, and Brighten Godfrey, “Practical Automation for Management Planes of Service Provider Infrastructure,” Workshop on Flexible Networks (FlexNets), August 2021.
- Christopher Hannon, Deepjyoti Deka, Dong Jin, Marc Vuffray, and Andrey Lokhov, “Real-time Anomaly Detection and Classification in Streaming PMU Data,” IEEE PowerTech Conference, June 2021.
- Christopher Hannon, Jiaqi Yan, and Dong Jin, “Distributed Virtual Time Based Synchronization for Simulation of Cyber-Physical Systems,” ACM Transactions on Modeling and Computer Simulation (TOMACS), April, 2021.

Monitoring, Fusion, and Response for Cyber Resilience

PI:	William Sanders
HARD PROBLEMS:	Security Metrics Driven Evaluation, Design, Development, and Deployment
	

GOAL

The goal of this project is to facilitate faster Sensitivity Analysis (SA) and Uncertainty Quantification (UQ) of slow-running cybersecurity models using a novel stacked metamodel approach.

ABSTRACT

Realistic state-based discrete-event cybersecurity simulation models are often quite complex. The complexity can manifest in models that (a) contain many input variables whose values are difficult to determine precisely, and (b) take a relatively long time to execute. Sensitivity Analysis and Uncertainty Quantification are used to understand and manage the uncertainty in the inputs. Unfortunately, the long execution times of models may make traditional SA and UQ prohibitively time

consuming. In this project, we developed a novel approach for performing faster SA and UQ by using a metamodel composed of a stacked ensemble of regressors that emulates the behavior of the base model. We demonstrate its use on a number of previously-published models, which we use as test cases. We find that our metamodels run hundreds or thousands of times faster than the base models, and are more accurate than state-of-practice metamodels.

ACCOMPLISHMENTS

We began to pursue two main directions in our research at the beginning of the year. First, we investigated whether adaptive sampling could be used to collect higher-quality training data which could be used to build more accurate metamodels for SA and UQ. Our adaptive sampling approach was not successful; the metamodels trained using the data collected by the adaptive sampling method were less accurate than the metamodels trained on non-adaptive sampling methods. We may return to adaptive sampling in the future, but we decided to shift our focus to a new direction.

The second approach we tried was to determine if the metamodeling approach generalized. Prior to this year, we had initially only tried it on two models, but we found six more models to use as test cases. We have found that our metamodeling approach works well on these six additional models. We created metamodels that were thousands of times faster than the base models, and more accurate than state-of-the-practice metamodels.

Name	Base Model Execution (seconds)	Metamodel Execution (seconds)	Metamodel Training (seconds)
Botnet	1115.5	0.3	137.4
Circadian	18291.6	0.5	204.7
Cluster	2034.9	0.6	66.4
Cyclin	2668.2	0.5	76.9
Embedded	684.7	0.2	60.6
Kanban	5737.6	0.3	64.5
Molecules	3714.5	1.1	67.1

Figure 1: Speed comparison of base model vs. metamodel for 7 test case models, and metamodel training times. Each base model and corresponding metamodel was run with the same 200 inputs.

In addition, we designed and tested several different variations on the base stacked metamodel architecture to determine the most effective architectures given our test cases. We also developed a plan to make the metamodeling tool we developed more widely available to academics.

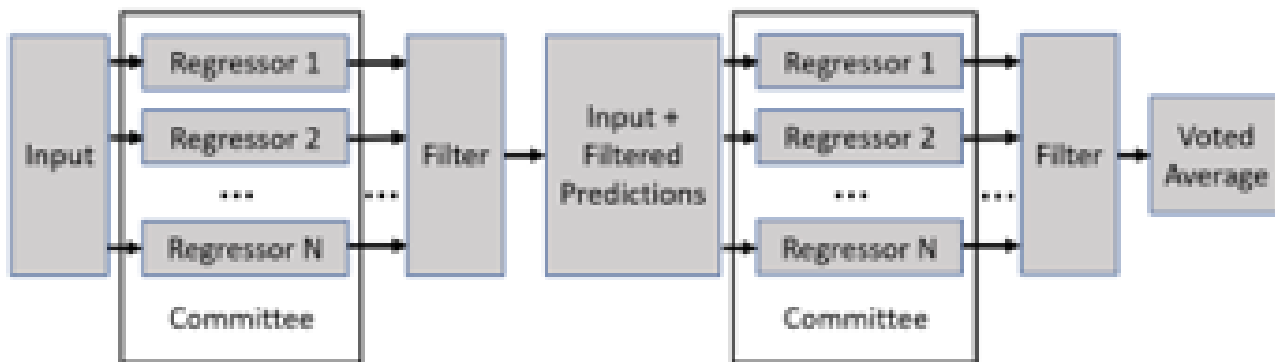


Figure 2: Base stacked metamodel architecture.

IMPACT ON HARD PROBLEMS


Security Metrics Driven Evaluation, Design, Development, and Deployment – Quantitative cybersecurity models produce metrics that are used in the evaluation, design, development, and deployment of critical cyber assets. The metrics the models produce must be of high-quality, and their limitations must be well-understood. Our stacked metamodel approach enables SA and UQ for models that run

too slowly for traditional SA and UQ. This improves the usefulness of quantitative cybersecurity models because the uncertainty in the models' outputs/metrics can be better understood and managed using the techniques we developed.

PUBLICATIONS

- Michael Rausch and William Sanders, "Evaluating the Effectiveness of Metamodels in Emulating Quantitative Models," In *Proceedings of the International Conference on Quantitative Evaluating of SysTems (QEST)*, Paris, France, August 23-27, 2021.
- Michael Rausch, "Addressing Challenges to Quantitative Security Modeling," PhD Dissertation, University of Illinois Urbana-Champaign, 2021.

Resilient Control of Cyber-Physical Systems with Distributed Learning

PI:	Sayan Mitra CO-PIs: Geir Dullerud, Sanjay Shakkottai
SUB-LABEL:	University of Texas at Austin (UT Austin)
HARD PROBLEMS:	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 are developing safety and security analysis approaches for real-life autonomous and Cyber-Physical Systems using statistical and Machine Learning techniques. Our approaches rely 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 and hybrid systems. 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.

IMPACT ON HARD PROBLEMS

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

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


EDUCATION AND OUTREACH

- Sayan Mitra presented a lecture on “Towards verified robot code” at Semiautonomous Systems Seminar, UC Berkeley, February 12th, 2021.
- Sanjay Shakkottai was co-organizer of and speaker at the Edison Lecture 2021 titled “Failing Well: Big Engineering Failures that Led to Big Successes,” February 2021. The Edison Lecture is a presentation to engage more than 1000 middle and high school students across Central Texas on STEM topics.
- Sayan Mitra participated in a panel discussion at the AFCEA Ideation and Innovation Virtual Event. March 10, 2021 on the state of the art and challenges in implementing autonomy.
- Sayan Mitra is serving as the General Chair of HoTSoS 22.
- Sayan Mitra gave the following invited lectures:
 - “Interfaces for models and data in verification and synthesis,” Workshop on Learning and Control, and seminars co-located with CPS-IoTWeek 21, May 18, 2021.
 - “Data requirements for estimation and verification,” Simons Institute, Theoretical Foundations of Computer Science Seminar, May 11, 2021.
 - “Verified autonomy code” for the SAE G-34/EUROCAE WG-114, AI in Aviation standard committee
 - “Data requirements for estimation and verification,” Simons Institute, Theoretical Foundations of Computer Science Seminar, May 11, 2021.
- Sayan Mitra participated in a virtual roundtable on “[Formal methods for cyber-physical systems](#)”, that appeared in the IEEE Computer magazine, 2021.
- We developed and organized the [GRAIC](#) Autonomous Racing Competition which was co-located with CPSWeek 2021. The live event had more than 80 registered members and 30+ attendees. The software framework has been made available to the community for research.

PUBLICATIONS

- Yu Wang, Nima Roohi, Matt West, Mahesh Viswanathan, and Geir Dullerud, “Verifying Stochastic Hybrid Systems with Temporal Logic Specifications via Model Reduction,” *Transactions on Embedded Computing Systems*, May 2021.
- Sung Woo Jeon and Sayan Mitra, “Egocentric abstractions for verification of distributed cyber-physical systems,” *IEEE Workshop on the Internet of Safe Things (SafeThings’21)*, co-located with IEEE S&P, May, 2021. **Best Paper Award.**
- Minghao Jiang, Kristina Miller, Dawei Sun, Zexiang Liu, Yixuan Jia, Arnab Datta, Necmiye Ozay and Sayan Mitra, “Continuous Integration and Testing for Autonomous Racing Software: An Experience Report from GRAIC,” Contributed paper in *ICRA 21 Workshop on Opportunities and Challenges with Autonomous Racing*, 31 May, 2021.
- Negin Musavi, Dawei Sun, Sayan Mitra, Sanjay Shakkottai, and Geir Dullerud, “Verification and Parameter Synthesis for Stochastic Systems using Optimistic Optimization,” In *Proceedings of IEEE Conference on Control Technology and Applications (CCTA)*, September 2021.
- Nard Strijbosch, Geir Dullerud, Andrew Teel, and Maurice Heemels, “L2-gain Analysis of Periodic, Event-Triggered Control and Self-Triggered Control using Lifting,” *IEEE Transactions on Automatic Control*, November, 2021.
- Hussein Darir, Hussein Sibai, Chin-Yu Cheng, Nikita Borisov, Geir Dullerud, and Sayan Mitra, “MLEFlow: Learning from History to Improve Load Balancing in Tor,” *Privacy Enhancing Technologies Symposium (PETS)*, 2022, To appear.

Uncertainty in Security Analysis

PI:	David Nicol
HARD PROBLEMS:	Security Metrics and Models, Resilient Architectures 

GOAL

The goal of this project is to develop a mathematical basis for describing and analyzing the ability of an adversary to laterally traverse networks in the presence of uncertainty about connections and uncertainty

about exploitable vulnerabilities. We will use this basis to develop algorithms for quantified risk analysis of cyber-physical systems.

ABSTRACT

Cyber-security vulnerabilities in Cyber-Physical Systems (CPS) allow an adversary to remotely reach and damage physical infrastructure. Following the initial point of entry, the adversary may move laterally through the computer network using connections that are allowed by the access control but which give access to services with exploitable weaknesses and 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 severe 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 are not. 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 considerable damage in the presence of uncertainty with respect to the connections and the vulnerabilities which enable lateral movement.

ACCOMPLISHMENTS

1. Attack loss quantification

We introduced the attack loss, a function that quantifies the loss to the network 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 analyses 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 loss distribution contains the “worst-case scenarios” under which the attack inflicts the largest amounts of loss. Understanding the risks of these low-probability but high-impact security events allow organizations to make better defense decisions. However, due to the extremely low probabilities, these events are difficult to quantify using standard Monte Carlo (MC) techniques.

To this end, we proposed two advanced MC approaches based on the idea of Importance Sampling (IS). The first approach is inspired by the Zero-Variance Approximation (ZVA). Our accomplishments were to (i) derive the analytically optimal IS scheme for estimating the loss tail probability and to (ii) propose an approximation to the optimal IS scheme that has the assurance of bounded relative error. (Recall the loss tail probability with respect to a threshold T measures the probability of the event of an attack causing a loss that is greater than T .) While the approximation scheme requires solving an NP-hard problem, we use a search procedure that becomes more efficient as the attack loss threshold increases.

While the ZVA-based approach enjoys the bounded relative error property and does not assume any property of the network under study, the method suffers from the curse of dimensionality as the number of links and hosts in the network increases. This issue limits the applicability of IS-based rare event simulation to safety and security analysis of large-scale systems. To address this issue, we proposed another IS approach based on the idea of Maximum Weight Minimization (MWM). To alleviate the effect of the curse of dimensionality, MWM exploits two problem-specific structures that are commonly found in engineered systems, namely, monotonicity and symmetry. Our simulation results demonstrated that MWM can produce good IS estimators in a short amount of time for large systems that consist of hundreds of variables.

2. Cybersecurity incident response

No matter how much resource an organization spends on cyber-defense, their computer network will eventually be breached. As a result, incident response is an integral part of the organization’s risk management strategy to reduce the residual risk of damage due to cyber-attacks. To this end, we extended risk analysis to incident response with a focus on understanding and quantifying (i) the impact of uncertainty on detecting and predicting the current and future evolution of an ongoing attack and (ii) the tradeoff between timeliness (i.e., being able to act quickly and effectively despite the lack of information) and observability (i.e., being able to grasp the full scope

of the attack, which takes a considerable amount of time). While this is an ongoing effort, our existing Bayesian framework can combine incomplete and potentially conflicting evidence about the state of an ongoing attack that is collected during the investigation phase. Using Sample Average Approximation (SAA), a stochastic

optimization method based on MC simulation, we can compute the optimal containment strategy that minimizes the net impact of false positive (i.e., containing a suspected but otherwise not compromised host) and false negative (i.e., not containing a compromised host).

IMPACT ON HARD PROBLEMS

This research intersects the predictive security metric problem since we are attempting to predict uncertainty associated with a system model. It also intersects with resilience as a system's resilience will

be established by analysis of some model and decisions (e.g., how significant the breach may be, whether to interdict and where, where to focus recovery activity) will be made as a result.


University of Kansas

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

The University of Kansas Lablet continued work on four projects targeting resiliency, preventing side channel communication, developing semantics and infrastructure for trust, and secure native binary execution. Specifically, KU investigators are: a) reducing micro-architectural side-channels by introducing new OS abstractions while minimally modifying micro-architecture and OS; b) developing an epistemology and ontology for framing resilience; c) formalizing remote attestation and experimenting with scalable prototypes; and d) developing a framework for client-side security assessment and enforcement for COTS software.

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 in addition to Dr. Alexander are John Symons, Heechul Yun, and Prasad Kulkarni. The KU Sub-Lablet at the University of Tennessee is led by Michael Jantz. The 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, and the Department of Sociology.

The Lablet team meets regularly with an Industrial Advisory Board that provides input on research directions, facilitates outreach, and helps establish a regional, midwestern cyber security community. The IAB has members from T-Mobile, Garmin, Eriksson, Kansas City National Security Campus, Ft. Leavenworth, Cboe, Collins Aerospace and The Kansas City Federal Reserve, among others. The Advisory board helps us track real-world cyber-security issues and evangelizes our research among their companies and the wider community.



PI Perry Alexander

KU PIs participated in an informational meeting with the KU Chancellor, Provost, Assessment and Achievement Institute (AAI) and LTG James Rainey, US Army CAC commander where we presented Lablet research projects. LTG Rainey is responsible for cybersecurity training across the US Army and is interested in our research for purposes of enhancing their training subjects and delivering training to a diverse military. Our Lablet and KU's AAI has significant potential for delivering cybersecurity education to soldiers in the field. AAI provides experiential learning and micro certification capabilities that can significantly enhance CAC's training capabilities.

Perry Alexander and John Symons helped organize a Red Hot Research symposium at KU on blockchain with emphasis on security and privacy. The symposium featured a number of student and faculty presentations targeting a general audience. A similar presentation is planned for SUNY Buffalo.

Lablet PIs continue to support our Center for High-Assurance Secure Systems and IoT (CHASSI) effort with Syracuse, Indiana, Minnesota and Case Western Reserve. We are defining projects that are of interest to our industrial partners. Several projects are closely related to our Lablet research projects and could result in significant technology transfer interest. Several KU Lablet PIs supported CHASSI Planning Workshop defining a research agenda in high assurance secure systems. The workshop included over 30 companies in support of the joint Industry/University Research Partnership. The workshop featured work associated with all of our Lablet research projects.

Last year, the KU Innovation Park, representing Ad Astra Integrity Measurement Systems negotiated rights to commercialize LKIM in cooperation with NSA. This year Ad Astra pursued hiring executive leadership and demonstrated software for several major players in

the IoT and telecommunications industry. While not directly funded by Science of Security, being a part of the program made this result possible. Our Industrial Advisory Board continues to help identify commercialization partners and identifying initial markets. Our NSA project champion, Peter Loscocco, continues to support us in this effort.

KU's PIs participated in an effort to establish a cybersecurity best practices research study for the Federal Avionics Administration. KU partnered with Oregon State University and Drexel University on a proposal to define the need for cybersecurity oversight and risk management associated with un-piloted avionics systems. This effort was funded and will begin in Q1 of 2022.

Perry Alexander and Xenofon Koutsoukos (Vanderbilt Lablet PI) hosted a conversation with Professor Chris Hankin (Imperial College) to discuss interactions between The Research Institute in Trustworthy Interconnected Cyber-physical Systems (RITICS) and our Lablets. The potential for synergist activities between RITICS and the KU and Vanderbilt Lablets is substantial.


KU initiated an internal Software Assurance Meetup for researchers interested in high-assurance and secure systems but are new to the area. Our goal is to involve researchers outside our Lablet in our

future work. Initial meetings include faculty from EECS, mathematics, sociology and philosophy. In support of this effort, Dr. Emily Witt from the KU Mathematics Department received a Keeler Professorship to work with Lablet researchers specifically to develop collaborative research projects.

KU PIs are restarting the GenCyber program for high-school teachers which brings high school teachers to the KU campus for a week-long intensive introduction to cyber security. GenCyber is supported by NSA and NSF.

KU has hired Dr. Jennifer Lohofener as a research assistant professor to focus on development of educational technologies and outreach. This position is funded partially by our Lablet with additional support from The KU Office of Research and KU School of Engineering. For the Lablet, Dr. Lohofener will focus on development of course materials for integrating security and formal methods. Working with Dr. Lohofener, KU PIs began an effort to develop materials for a new course that teaches formal analysis techniques with laboratory examples from cyber-security. The objective of this effort is to develop a course that introduces formal techniques to undergraduate and graduate students in the context of verifying secure and trusted systems.

Formal Approaches to the Ontology and Epistemology of Resilience

PI:	John Symons
HARD PROBLEMS:	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

Accomplishments are reflected in the Education and Outreach and Publication entries below.

IMPACT ON HARD PROBLEMS

Security Science has focused on network-based measures of resilience. This is a valuable formal approach, but its range of application is narrower than the general problem requires. In order to make progress

on these questions, a broader theoretical approach is required, and we will need to call on a range of other formal and informal methods.


EDUCATION AND OUTREACH

- John Symons completed a review article for the British Journal of Philosophy of Science on Primiero's Foundations of Computing book
- John Symons will organize a conference on miscomputation, malfunction, and error during the spring semester. He and his colleagues will produce a journal special issue based on accepted papers in late 2022.

PUBLICATIONS

- John Symons and Jack Horner, "What Have Google's Random Quantum Circuit Simulation Experiments Demonstrated About Quantum Supremacy?" In Arabnia H.R., Deligiannidis L., Tinetti F.G., Tran QN. (eds) Advances in Software Engineering. *Transactions on Computational Science and Computational Intelligence*. Springer, Cham. https://doi.org/10.1007/978-3-030-70873-3_29.
- John Symons, Stephanie Harvard, Eric Winsberg, and Amin Adibi, "Value judgments in a COVID-19 vaccination model: a case study in the need for public involvement in health-oriented modelling," *Social Science & Medicine*, 2021.
- John Symons and Marco Camacho, "Millikan's biological explanation of mental representation—a critical introduction," *Saudi Journal of Philosophical Studies* 1 (1), 195-201, 2021.
- John Symons and Hassan Alsharif, "Open-mindedness as a Corrective Virtue," *Philosophy* 96, 1 73-97 (2021)

Scalable Trust Semantics and Infrastructure

PI:	Perry Alexander
HARD PROBLEMS:	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. Our project focuses on attestation among heterogeneous collections of systems and systems-of-systems. We focus on developing precise semantics informed by prototype implementations. Our goal is providing engineers with both the intellectual and programming tools to field effective remote attestation systems.

ACCOMPLISHMENTS

Lablet researchers continue their work on model development and prototype implementation. This year we published the first verified semantics for an attestation manager and published results from verification of an end-to-end layered attestation and appraisal using that semantics. With our colleagues at MITRE, JHUAPL and NSA we published a collection of patterns for flexible deployment of layered attestations, and continued development of a model for attestation protocol negotiation using dependent types for policy enforcement. Finally, we developed a model and prototype implementation of a general appraisal algorithm.

Our team helped maintain the base Copland language definition for attestation protocols and maintained the publicly available Copland repository including CakeML, Coq and Haskell attestation manager implementations. Throughout the year we released updates to our

attestation manager implementations. Finally, Lablet investigators contributed to the initial development and verification of a layered attestation implementation and trusted boot. We are in the initial stages of standing up an attestation testbed to support experimentation with our tools. The testbed will support collaboration and demonstration of our tools and technique.

Lablet investigators developed a health record for certificate/passport style attestation and began work with the TPM 2.0. The health record captures attestation results and is used to preserve those results over time. We currently use a blockchain to store results, but other structures are equally useful. We began integrating the TPM 2.0 into the Copland attestation framework in a manner that will allow using TPM operations in attestation protocols and use TPM keys for strong identity, evidence integrity, and confidentiality.

IMPACT ON HARD PROBLEMS

Scalability and Composability

A foundational contribution of our research is a verified semantics of trust embodied in scalable prototype implementations. We provide definitions of trust and metrics for soundness of evaluation and appraisal that include soundness and sufficiency of evidence, semantic mechanisms for identity and attestation, formal definitions of evidence, and meta-evidence appraisal. This work contributes directly to implementing and scaling trust infrastructure in the form of hierarchical frameworks for trust infrastructure including virtualized TPM implementations, trust aggregation and trust as a service.

Policy-Governed Secure Collaboration

The Copland attestation protocol representation and semantics provides a means for evaluating policy. Our work with negotiation implements selection and privacy policy in a way that ensures parties involved in an attestation understand the evidence they share, mechanisms used, and policies that enforce privacy. Using formal, executable representations for attestation protocols ensures that policies in play during protocol development and analysis are also honored at run-time.

EDUCATION AND OUTREACH

- We continued our support for the High Confidence Software and Systems symposium (HCSS) with Perry Alexander serving a member of the planning committee. Ian Kretz presented a paper on confining adversaries at HCSS'21 that included joint work between KU Lablet researchers, MITRE, and JHUAPL.
- Adam Petz presented a paper at The NASA Formal Methods conference on verification of an attestation manager.
- Perry Alexander and Raj Pal (NSA), along with representatives of T-Mobile and Arm began a dialog around attestation as a service for authentication in 5G networking and continue to meet biweekly. This group has significant visibility in several standards organizations including TCG and IETF providing us input into industrial standards and practice.
- We continue our joint work with MITRE, JHUAPL, and NSA to explore new layered attestation approaches. Joint work from this effort is available at <https://www.copland-lang.org> including the Copland Collection of utilities and tools, Copland formal semantics, and attestation manager implementations.
- Poster presentation: Anna Fritz, and Perry Alexander, "A Dependently Typed Attestation Protocol Language," Hot Topics in the Science of Security (HoTSoS 2021), April 13-15, 2021.


PUBLICATIONS

- Adam Petz and Perry Alexander, "An Infrastructure for Faithful Execution of Remote Attestation Protocols," *NASA Formal Methods Symposium (NFM'21)*, May 24-28, Norfolk, VA.
- Sarah Helble, Ian Kretz, Peter Loscocco, John Ramsdell, Paul Rowe, and Perry Alexander, "Flexible Mechanisms for Remote Attestation," *ACM Transactions on Privacy and Security*, 24(4), pp 1-23.
- Adam Petz, Grant Jurgensen, and Perry Alexander, "Design and Formal Verification of a Copland-based Attestation Protocol," *ACM-IEEE International Conference on Formal Methods and Models for System Design (MEMOCODE'21)*, Virtual, Nov 20-22, 2021.

Software Releases

- <https://www.copland-lang.org/>

Secure Native Binary Execution

PI:	Prasad Kulkarni
SUB-LABEL:	University of Tennessee
HARD PROBLEMS:	Security Metrics and Models, Scalability and Composability 

GOAL

This goal of this research is to build tools and techniques that will allow users to determine 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

Our team focused on two projects: (a) build an effective and efficient hybrid (combined static-dynamic) binary program analysis and instrumentation framework to harden software binaries on the client-side to protect them from different categories of memory-related vulnerabilities; and (b) build an infrastructure for client-side security assessment for binaries. We continued our work to develop tools and techniques to evaluate the client-side security properties of binary software, including (a) the programming language that was used to build the software, (b) the automated detection of any compiler inserted security checks in the binary, and (c) the security-relevant coding conventions used by the developers. We also explored new methods to evaluate the effectiveness of implementing source-level (compiler-based) security techniques at the binary-level and built

a new technique to protect binary software from memory-related attacks. Our contributions include developing new techniques, collecting diverse benchmarks, and designing detailed experimental frameworks to support this research. We built a rules-based technique to automatically detect the compiler-based security checks in the binary without relying on known code signatures; developed new ML models to infer the programming language used when given only the stripped binary executable; explored the challenges and evaluated the effectiveness of binary-level techniques to detect and prevent memory errors; started a new study to assess the effectiveness and efficiency of conducting Control-Flow Integrity (CFI) on binary code as compared to performing CFI on source code; and continued work on a high-performance static analysis and binary rewriting tool that inserts checks for spatial memory vulnerabilities during program startup.

IMPACT ON HARD PROBLEMS


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

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

PUBLICATIONS

- Raturaj Vaidya, Prasad Kulkarni, and Michael Jantz, "Explore Capabilities and Effectiveness of Reverse Engineering Tools to Provide Memory Safety for Binary Programs," *International Conference on Information Security Practice and Experience*, 2021.

Side-Channel Attack Resilience

PI:	Heechul Yun
HARD PROBLEMS:	Resilient Architectures 

GOAL

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

ABSTRACT

Cyber-Physical Systems (CPS)--cars, airplanes, power plants, etc.--are increasingly dependent on powerful and complex hardware for higher intelligence and functionalities. However, this complex hardware may also introduce new attack vectors--hardware side-channels--which can be exploited by attackers to steal sensitive information, to disrupt timing of time-critical functions that interact with the physical plants, or to break memory protection mechanisms in modern computers. Because these attacks target hardware, even logically safe and secure software such as a formally verified OS, could

still be vulnerable. Given the safety-critical nature of CPS, hardware side-channels should be thoroughly analyzed and prevented in CPS. This project focuses on micro-architectural side channels in embedded multicore computing hardware and aims to develop fundamental OS and architecture designs that minimize or 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 speculative execution-based attacks. As part of this effort, we validated our previously developed SpectreRewind covert channel, which targets non-pipelined floating point division unit, on several recent Intel architectures and found that our attacks work well after minor adaptation. We released SpectreRewind PoCs, which include a C version, a JavaScript version, and a modified end-to-end Meltdown attack using SpectreRewind, as open source on a public github repository.

We have continued to develop Denial-of-Service (DoS) attack and prevention techniques. As part of this effort, we continued to evaluate Intel RDT technologies and their effectiveness in preventing DoS

attacks on recent intel architectures. In addition, we have demonstrated memory-aware (DRAM bank-aware) variant of DoS attack techniques are substantially (by up to 5X) more effective in causing execution time increases to cross-core victims than memory-unaware state-of-the-art cache DoS attacks on embedded ARM platforms. More recently, we also successfully demonstrated a similar degree of improved effectiveness on Intel platforms.

We have continued to develop operating system-level solutions to mitigate micro-architectural attacks. As part of this effort, we proposed new interference-aware group (virtual gang) scheduling methods that minimize utilization loss of our previously proposed gang scheduling approach while still preventing cache DoS attacks

IMPACT ON HARD PROBLEMS

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

- Dr. Yun was invited to serve as TPC Chair of IEEE RTAS'22.

PUBLICATIONS

- 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
- Michael Bechtel and Heechul Yun, "Memory-Aware Denial-of-Service Attacks on Shared Cache in Multicore Real-Time Systems," *IEEE Transactions on Computers*, 2021

Software Releases


<https://github.com/CSL-KU/SpectreRewind-POC>

Vanderbilt University

The Vanderbilt University (VU) Science of Security (SoS) Lablet team is led by Principal Investigator (PI) Xenofon Koutsoukos. VU Sub-Lablets are Massachusetts Institute of Technology and University of California, Berkeley. The Lablets and Sub-Lablets' faculty, postdoctoral and PhD student researchers 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.


While educational robotics and makerspaces are useful to modern STEM education, they introduce both physical and economic barriers to entry. In 2021 we created a "Cyber Makerspace," a simulated, networked environment where we can facilitate instruction on cyber-physical systems and their security and related topics while reducing cost and complexity. The approach will facilitate reaching audiences from traditionally underrepresented groups. It also supports remote learning, an especially important feature due to the current pandemic. See Section 3 for more details on Cyber Makerspace.



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

PI:	Nazli Choucri
SUB-LABEL:	Massachusetts Institute of Technology
HARD PROBLEMS:	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

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

When NIST issued a fifth revision (Rev. 5) of its document 800:53, it impacted the formal connections, or interfaces, between multiple sources of “raw” data for this project. It also coupled very closely the controls and control families for security and privacy. Further, this revision raised serious questions about the implications of this new version of 800:53 regarding current NIST perspectives and priorities pertaining to security.

For this project, especially important is the fact that we faced a necessary “re-do” of research steps, and a review of results, with respect to: Data Linkage Process; information pertaining to security controls and control families; an unexpected entanglement between the Security and the Privacy controls, and control families, thereby creating new ambiguities; data-based signals that, in the security domain, “everything is related to everything else and to privacy as well”; and serious impediments to the “reversing the arrows” test and

a validation strategy for our research design and results. We devoted ourselves to the “re-do” as well as to delineating and understanding the implications of the “entanglements” of security and privacy controls for policy implementation of the Cybersecurity Framework (CSF). The validated “re-do” shows an unexpected result, namely that privacy and security controls are heavily dependent on each other, thus creating “noise” in a focused analysis of security controls.

We have identified practical uses of research and results so far:

Data Linkages: The full value of the CSF is difficult to capture given the set of intervening tasks required and the distributed nature of the database. CSF points to what has to be done and why, but not how. It is up to the user to work through the process outlined by CSF. In this case, the practical use is created by providing a method to streamline access to, and use of, essential data required to implement the security-related actions required by CSF. Because CSF points to a

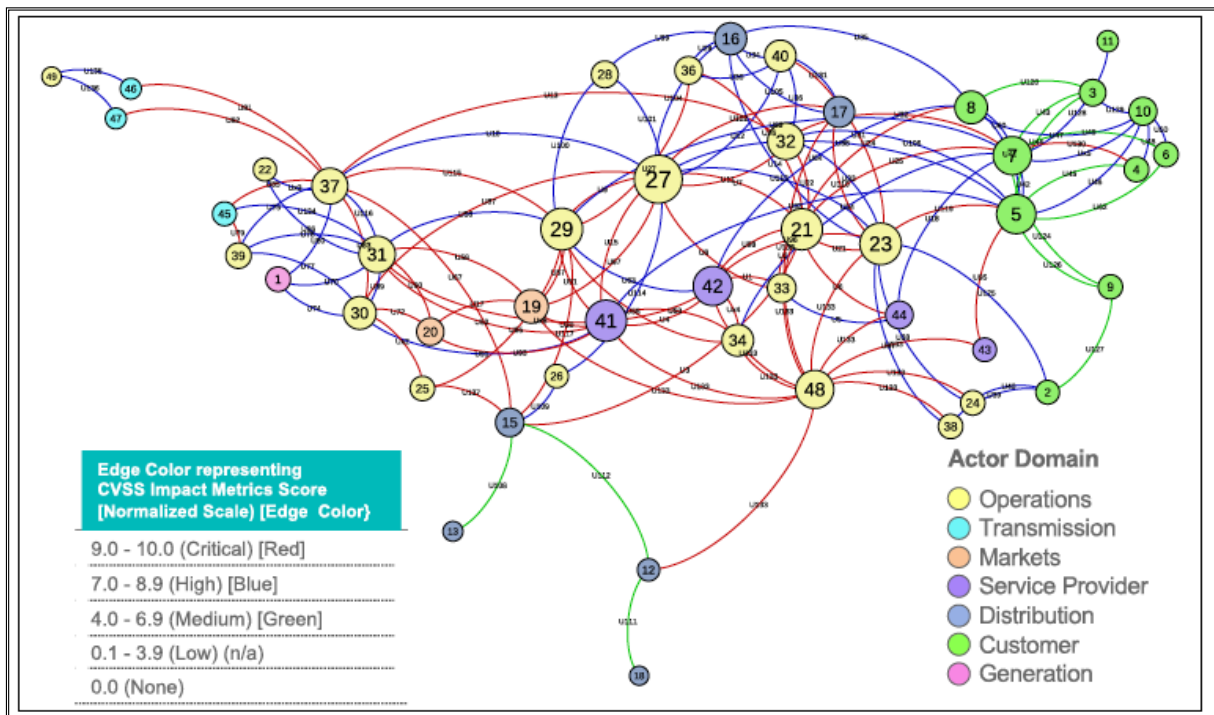
number of individual documents hosting different directives, the users' task is to identify and make connections among them as needed. Moreover, modifications and updates by NIST on the content of key intervening documents require users, in turn, to identify the updates and determine requirements for change.

Metrics & Measures: Given that policy documents and directives are conveyed in text form, in linear and sequential order, it is common practice to retain information in that form. We developed a method to transform text into metrics to deal with numerals, not letters. The practical use is compelling: metrics and measures enable more precision, with more flexibility in scale and scope of analysis, than can ever be done with the text form. This in itself takes away much of the built-in ambiguity of policy documents. Since the method is portable, it can be applied to all forms of policy texts - irrespective of issue area or domain.

One part of the research design focuses on analytics for the cyber-physical system itself; the other part is on analytics for

cybersecurity policies and directives. Both "parts" share a common process that must be applied to each side separately because the data are separate. We simplify the process in terms of: Text to Data; Data to Metrics; and Metrics to Model. Part I is completed with the smart grid network model. This generates and identifies the nodes and the logical interfaces. Earlier we focused on task testing for Part II and identified empirically the logical interfaces for the system "as is" that connect to CSF directives.

We present the results for the computed base network model of the reference system for the NIST Smart Grid as a cyber-physical system. NIST provides very detailed information on vulnerability impacts for violation of each of the security objectives (confidentiality, integrity, and availability). In order to facilitate action, a net assessment for vulnerability impacts on the system across three dimensions and three levels of intensity is needed and can be performed using the Common Vulnerability Scoring System (CVSS). The figure below shows the consolidation of analysis and results in one integrated system view.



NIST Smart Grid reference network with edge weighted by the impact level based on CVSS 3.0.

IMPACT ON HARD PROBLEMS

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.

PUBLICATIONS

- Nazli Choucri and Guarav Agarwal, "Complexity of International Law for Cyber Operations," 20th IEEE Symposium on Technologies for Homeland Security (HST21), November, 2021.
- Nazli Choucri, Stuart Madnick, Keman Huang, and Fang Zhang, "A Systematic Framework to Understand Transnational Governance for Cybersecurity Risks from Digital Trade," *Global Policy* (Durham University and John Wiley & Sons Ltd.): pp: 1-14.

Foundations of CPS Resilience

PI:	Xenofon Koutsoukos
HARD PROBLEMS:	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

Modeling and simulating attacks in power systems

Due to the increased deployment of novel communication, control and protection functions, the grid has become vulnerable to a variety of attacks. Designing robust machine learning based attack detection and mitigation algorithms require large amounts of data that rely heavily on a representative environment, where different attacks can be simulated. We have developed a comprehensive tool-chain for modeling and simulating attacks in power systems. First, we present a probabilistic domain specific language to define multiple attack scenarios and simulation configuration parameters. Secondly, we extend the PyPower-dynamics simulator with protection system components to simulate cyber-attacks in control and protection layers of power system. We demonstrate multiple attack scenarios with a case study based on IEEE 39 bus system.

Byzantine Resilient Aggregation in Distributed Reinforcement Learning

Recent distributed reinforcement learning techniques utilize networked agents to accelerate exploration and speed up learning. However, such techniques are not resilient in the presence of Byzantine agents which can disturb convergence. In this work, we present a Byzantine resilient aggregation rule for distributed reinforcement learning with networked agents that incorporates the idea of optimizing the objective function in designing the aggregation rules. We evaluate

our approach using multiple reinforcement learning environments for both value-based and policy-based methods with homogeneous and heterogeneous agents. The results show that cooperation using the proposed approach exhibits better learning performance than the non-cooperative case and is resilient in the presence of an arbitrary number of Byzantine agents.

Assurance Monitoring of Learning-enabled CPS

Machine learning components such as deep neural networks are used extensively in CPS but such components may introduce new types of hazards that can have disastrous consequences and need to be addressed for engineering trustworthy systems. Although deep neural networks offer advanced capabilities, they must be complemented by engineering methods and practices that allow effective integration in CPS. We proposed an approach for assurance monitoring of learning-enabled CPS based on the conformal prediction framework. In order to allow real-time assurance monitoring, the approach employs distance learning to transform high-dimensional inputs into lower size embedding representations. By leveraging conformal prediction, the approach provides well-calibrated confidence and ensures a bounded small error rate while limiting the number of inputs for which an accurate prediction cannot be made. The experimental results demonstrate that the error rates are well-calibrated while the number of alarms is very small. The method is computationally efficient and

allows real-time monitoring. Current and future work includes using the approach for detection and classification of attacks in CPS/IoT.

Reliable Probability Intervals for Classification

Deep neural networks are frequently used by autonomous systems for their ability to learn complex, non-linear data patterns and make accurate predictions in dynamic environments. However, their use as black boxes introduces risks as the confidence in each prediction is unknown. Different frameworks have been proposed to compute accurate confidence measures along with the predictions but at the same time limitations such as execution time overhead or inability to be used with high-dimensional data. In this research we use Inductive Venn Predictors for computing probability intervals regarding the correctness of each prediction in real-time. We propose taxonomies based on distance metric learning to compute informative probability intervals in applications involving high-dimensional inputs. Empirical evaluation on botnet attacks detection in Internet-of-Things (IoT)

applications demonstrates improved accuracy and calibration. The proposed method is computationally efficient, and therefore, can be used in real-time.

Inductive Conformal Out-of-distribution Detection

Machine learning components are used extensively to cope with various complex tasks in highly uncertain environments. However, Out-Of-Distribution (OOD) data may lead to predictions with large errors and degrade performance considerably. This research first introduces different types of OOD data and then presents an approach for OOD detection for classification problems efficiently. Our approach utilizes an Adversarial Autoencoder (AAE) for representing the training distribution and Inductive Conformal Anomaly Detection (ICAD) for online detecting OOD high-dimensional data. Experimental results using several datasets demonstrate that the approach can detect various types of OOD data with a small number of false alarms. Moreover, the execution time is very short, allowing for online detection.

IMPACT ON HARD PROBLEMS

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 attacks.

EDUCATION AND OUTREACH

- We held a meeting with Chris Hankin and Perry Alexander to discuss the synergies between the Lablets and Research Institute in Trustworthy Inter-connected Cyber-Physical Systems (RITICS).

PUBLICATIONS

- Feiyang Cai, Ali Ozdagli, and Xenofon Koutsoukos, "Detection of Dataset Shifts in Learning-Enabled Cyber-Physical Systems using Variational Autoencoder for Regression," *IEEE International Conference on Industrial Cyber-Physical Systems (ICPS 2021)*, May 10-12, 2021.
- Ajay Chhokra, Carlos Barreto, Abhishek Dubey, Gabor Karsai, and Xenofon Koutsoukos, "Power-Attack: A comprehensive tool-chain for modeling and simulating attacks in power systems," *9th Workshop on Modeling and Simulation of Cyber-Physical Energy Systems (MSCPES 2021)*, May 18, 2021.
- Dimitrios Boursinos and Xenofon Koutsoukos, "Assurance Monitoring of Learning Enabled Cyber-Physical Systems using Inductive Conformal Prediction based on Distance Learning," *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 35(2), 251-264, May 31, 2021.
- Waseem Abbas, Mudassir Shabbir, Yasin Yazicioğlu, and Xenofon Koutsoukos, "Edge Augmentation with Controllability Constraints in Directed Laplacian Networks," *IEEE Control Systems Letters*, 60th IEEE Conference on Decision and Control, 2021.
- Dimitrios Boursinos and Xenofon Koutsoukos, "Reliable Probability Intervals for Classification Using Inductive Venn Predictors Based on Distance Learning," *IEEE International Conference on Omni-layer Intelligent systems (COINS 2021)*, August 23-25, 2021.
- Feiyang Cai, Ali Ozdagli, Nicholas Potteiger, and Xenofon Koutsoukos, "Inductive Conformal Out-of-distribution Detection based on Adversarial Autoencoders," *IEEE International Conference on Omni-layer Intelligent systems (COINS 2021)*, August 23-25, 2021.
- Himanshu Neema, Leqiang Wang, Xenofon Koutsoukos, CheeYee Tang, and Keith Stouffer, "Model-Based Risk Analysis Approach for Network Vulnerability and Security of the Critical Railway Infrastructure," *The 16th International Conference on Critical Information Infrastructures Security (CRITIS 2021)*, September 27-29, 2021.
- Bradley Potteiger, Feiyang Cai, Zhenkai Zhang, and Xenofon Koutsoukos, "Data Space Randomization for Securing Cyber-Physical Systems," *International Journal on Information Security*, 2021.
- Jiani Li, Feiyang Cai, and Xenofon Koutsoukos, "Byzantine Resilient Aggregation in Distributed Reinforcement Learning," *18th International Conference on Distributed Computing and Artificial Intelligence (DCAI'21)*. Lecture Notes in Networks and Systems, vol 327, pp. 55-66, Springer, Cham.
- Yi Li, Xenofon Koutsoukos, and Yevgeniy Vorobeychik. "Adversarial Gaussian Process Regression in Sensor Networks," *Game Theory and Machine Learning for Cyber Security*, 149-159, 2021.

Mixed Initiative and Collaborative Learning in Adversarial Environments

PI:	Claire Tomlin CO-PI: Shankar Sastry
SUB-LABELT:	University of California, Berkeley
HARD PROBLEMS:	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

Resilience and the Science of Security for multi-agent systems is certainly the hot area in a number of fields right now: in areas straddling AI/ML, control theory and robotics. The work on aggregative games, model predictive learning applied to multi-player games seems to provide the best set of underpinning tools for this purpose. Our focus has been on trying to analyze and design attacks on learning in multiplayer games. The research led to the publications noted below. Adversarial attacks in the context of Recurrent Neural Network (RNN) based sequence classifiers involve carefully crafting very small perturbations to a nominal input, in order to fool the network into misclassifying the input signal. In this project, the main goal is to find analytical conditions under which RNNs are susceptible to such adversarially constructed input perturbations. By taking a dynamical system theoretic approach, the problem of finding adversarial perturbations is reframed as a control synthesis problem, with the disturbances viewed as control inputs. Stability and robustness of RNNs can then be studied using the well-developed tools from control theory. Concretely, we obtain sufficient conditions in terms of the trainable parameters of the RNN under which adversarial perturbations exist, and show how such perturbations can be dynamically constructed as a feedback control, by leveraging the sequential nature of RNNs. Our algorithm scales gracefully with the length of the nominal input sequence, and can be deployed for real time attacks. However, given our controls framework, perturbations may also be designed using off-the-shelf tools from optimal control. This work has been tested on a wide variety of sequence classification examples, including tasks like sentiment classification from the NLP domain.

Convolutional and RNNs have been widely employed to achieve state-of-the-art performance on classification tasks. However, it has also been noted that these networks can be manipulated adversarially

with relative ease, by carefully crafted additive perturbations to the input. Though several experimentally established prior works exist on crafting and defending against attacks, it is also desirable to have rigorous theoretical analyses to illuminate conditions under which such adversarial inputs exist. Our work in using a control structure to dynamically compute adversarial perturbations for RNNs provides both the theory and supporting experiments for real-time attacks. The focus is specifically on recurrent architectures and inspiration is drawn from dynamical systems theory to naturally cast this as a control problem, allowing dynamic computation of adversarial perturbations at each timestep of the input sequence, thus resembling a feedback controller. We have worked on a number of examples in Human Activity Recognition (HAR) and sentiment analysis in text.

A canonical problem for unmanned vehicles (ground, air and water borne) is to provide provably safe algorithms for not running into moving obstacles, including other agents. A huge number of probabilistic complete algorithms called RRT has been proposed for this purpose. But they are basically unusable in practice for such elementary applications as driving in cluttered environments because they are centralized and kinematic (that is that they are based on computational geometry rather than dynamics of the agents). This glaring omission has resulted in problems for air space management (sometimes called UAS for UAVs), and also for the certification of driving cars. The resilience aspect of these algorithms is that they need to be designed to be robust to adversarial attack by rogue vehicles. The most common approach here has been a game theoretic approach with inverse reinforcement learning to determine adversarial intent. The practical problems here are these solutions are usually so conservative as to be useless in real world scenarios. A possible way around this is the use of Model Predictive and Learning Games.

IMPACT ON HARD PROBLEMS

We have been developing a framework for incorporating **human behavior** into resilient robot motion planning. We have been using reachability analysis to develop **scalable**, online safety updates of

these motion plans. We have been developing **scalable**, analyzable methods for learning unknown dynamics within the framework of feedback linearization.

EDUCATION AND OUTREACH

- Shankar Sastry launched a new Institute entitled the C3 Digital Transformation Institute (<https://c3dti.ai>) a partnership of Berkeley, UIUC (co-leads) with U Chicago, CMU, MIT, Princeton, Stanford, Royal Institute of Technology 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 have been furthering the agenda of SoS in the workshops that this institute has run in the Spring see <https://c3dti.ai/events/workshops> Two workshops, one in March 2021 and the other in May 2021, focused on new vulnerabilities which are introduced through the introduction of learning algorithms in multi-player systems.
- PI Claire Tomlin and Co-PI Shankar Sastry have taken the lead in revamping large amounts of the undergraduate and graduate curriculum to feature the recent confluence of AI/ML, robotics, and control. In the Fall, Sastry taught his


course on Introduction to Robotics (a mezzanine course for undergrads and Masters students) for about 150 students (see <https://ucb-ee106.github.io/106a-fa20site/>) for the resources associated with this class). In a partnership with OSD's National Security Innovation Network (NSIN) we planned to place the top students from this class at various labs in the DoD in Summer 2021 (this was piloted successfully in Summer 2020, see <https://blumcenter.berkeley.edu/news-posts/national-security-innovation-with-uc-berkeley/>)

- In Spring 2021 Sastry taught the second term version of this class EECS 106B/206B which will include the research results from both of the projects discussed above. He, along with PI Tomlin and others collaborated in teaching a new graduate class combining vision, deep learning and control EECS 290-5. These three courses are all part of the revamping of the curriculum mentioned above.

PUBLICATIONS

- Shankar Deka, Dušan Stipanović, and Claire Tomlin, "Feedback-Control Based Adversarial Attacks on Recurrent Neural Network," In *59th IEEE Conference on Decision and Control (CDC)*, pages 4677–4682, 2020. <https://doi.org/10.1109/CDC42340.2020.9303949>.
- Shankar Deka, Dušan Stipanović, and Claire Tomlin, "Dynamically Computing Adversarial Perturbations for Recurrent Neural Networks," *IEEE Transactions on Control Systems Technology* <https://arxiv.org/abs/2009.02874>.
- Victoria Tuck, Yash Pant and Shankar Sastry "Decentralized Path Planning for Moving Obstacles and Multi-Agent Systems," *2021 IEEE Conference on Control Theory and Applications, CCTA*, San Diego, August 2021.

Multi-Model Test Bed for the Simulation-Based Evaluation of Resilience

PI:	Peter Volgyesi	CO-PI: Himanshu Neema
HARD PROBLEMS:	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

Deep Learning Testbed Infrastructure and Graph Neural Networks on AWS

We made significant improvements to DeepForge, our web-based collaborative design and experimentation platform for deep neural network-oriented research, including developing a graph neural network support for DeepForge to create a more accessible design and evaluation environment in this domain. Our current research in developing novel graph descriptor representations is supported by an Amazon Web Services (AWS)-based scalable deployment. This work borrows some ideas from the controllability of Laplacian dynamics and obtains more expressive representations of the graph structure (graph embedding) based on how some phenomenon spreads/propagates/evolves in the structure. Network training and evaluation requires significant computational power, thus we rely on customized on-demand AWS instances to support this effort.

Threat Modeling and Risk Analysis in Industrial Control Systems

In this effort, we are working on developing a modeling and analysis framework for threats and cybersecurity risks in Industrial Control Systems (ICS). Identification of system vulnerabilities and implementation of appropriate risk mitigation strategies are crucial for ensuring the cybersecurity of ICS. These system vulnerabilities must be evaluated depending on their exploitability, impact, mitigation status, and target platform and environments. Therefore, in order to assess

system vulnerabilities and risk mitigation strategies quantitatively, we are focusing on threat modeling and risk analysis methods for the cybersecurity of Railway Transportation Systems (RTS), which are real-world ICS and have become increasingly vulnerable to cyberattacks due to growing reliance on networked physical and computation components. An interesting aspect of RTS is that these systems have a continuously changing network topology due to moving locomotives. These systems, in general, are CPS with integral but non-stationary components. The key challenge posed by non-stationarity is the evolving nature of threats and vulnerability propagation owing to dynamic network connections that form and disappear as components move! Our framework dealing with this effort is called the Risk Analysis Framework (RAF). In our work, we have been successful in modeling the dynamic network connections and integrating it into dynamic vulnerability propagation algorithms. We extended the framework to incorporate cyber-gaming of exploits versus mitigations to plan for worst-case attacks; we also developed methods to deal with dynamic network connections where the vulnerabilities and their propagation via changing network connectivity continually changes.

Physics-guided Learning and Surrogate Modeling - Resilient CPS Applications

We continued our experimentation work for structural design and health monitoring for CPS applications, and developed two alternative

methods for auto-generating FAE static stress simulation results with relatively simple parametric CAD models (pressure vessel capsules). The generated datasets are used for developing physics-guided ML models and are also used for experimenting with a graph machine learning-based approach for FEM surrogate modeling and/or topology optimization. While the physics-guided learning approach has a broad use-case in CPS (e.g. buildings, transportation infrastructure) we successfully applied the results in the design process of Unmanned Underwater Vehicles (UUV) as part of the DARPA Symbiotic program.

Resilient Consensus using Centerpoint Algorithm and Hashgraph Based Communication

Unmanned Aerial Vehicles (UAVs) are used for a variety of tasks, such as inspection of dangerous environments, surveillance, and pursuit of a target. These systems use distributed machine learning algorithms to cooperate towards achieving an objective and are prone to Denial of Service (DoS) and integrity attacks. We investigate an approach for integrating a messaging mechanism and a coordination algorithm

based on Stochastic Gradient Descent (SGD) in a multi-agent network for target pursuit that is resilient against such attacks. The network consists of agents sending messages containing local data and state estimates and uses the SGD algorithm to optimize the global loss by aggregating state estimates from immediate neighbors. The network can suffer from a DoS attack to disrupt the ordering of messages or an integrity attack where one agent sends arbitrary estimates to neighbors to disrupt the convergence of normal agents towards an optimal state. The messaging mechanism uses Hashgraph, a distributed ledger technology, to guarantee a correct ordering of messages. The SGD algorithm uses centerpoint-based aggregation for converging to a target in the presence of compromised agents. We evaluate the approach using scenarios of target pursuit for multi-UAV systems using simulations in Microsoft AirSim with PX4 flight controllers. The evaluation results demonstrate cases for which the multi-agent system under attack is resilient and converges to the approximate optimal state.

IMPACT ON HARD PROBLEMS

As part of our framework's metrics-driven evaluation capability, we are developing a library of cyber-attacks (as DeepForge workflows/pipelines) and neural network models. The testbed provides a strictly versioned store for these models along with input datasets (not 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 continued our collaboration and technical exchange with the Cybersecurity Research Group at Fujitsu System Integration Laboratories Ltd. Discussion topics included the following: Threat Intelligence - Sharing Policy Enforcement (SPE); WebGME-based modeling tool development for SPE; and Integration with the MITRE ATT@CK framework
- Based on our collaboration with NIST on threat modeling and risk analysis in ICS, Vanderbilt and NIST presented our work on risk analysis and management to a group of top-level people from the Association of American Railroads, Railway Information Security Committee (RISC), and Railway Suppliers Committee. The meeting was well received and we have planned to continue our discussion with them in the near future.
- As part of our summer internship program, we were presenting our ongoing research projects for undergraduate students. PI Peter Volgyesi gave a presentation and led a discussion on AI-driven communication infrastructure CPS topics on May 26, 2021. The material covered important parts of the testbed (WebGME/DeepForge) with the goal of enticing prospective graduate students for the Lablet project.
- Presented the work-in-progress paper "A Model-Based Risk Analysis Approach for Network Vulnerability of Railway Infrastructure" at the *8th Symposium on Hot Topics in the Science of Security (HotSoS '21)*. The paper was later published (see below).

PUBLICATIONS

- Himanshu Neema, Leqiang Wang, Xenofon Koutsoukos, CheeYee Tang, and Keith Stouffer, "Model-Based Risk Analysis Approach for Network Vulnerability and Security of the Critical Railway Infrastructure," *16th International Conference on Critical Information Infrastructures Security (CRITIS' 2021)*, Sep. 27-29, 2021, Lausanne, Switzerland.
- Ali Ozdagli, and Xenofon Koutsoukos, "Model-based Damage Detection through Physics-guided Learning for Dynamic Systems," November 2021.

Science of Security

Quarterly Meetings

Winter 2021 Lablet

Quarterly meeting

The Winter 2021 Science of Security and Privacy (SoS) Quarterly Lablet meeting was hosted virtually by Vanderbilt University (VU) on January 12-13, 2021. The virtual attendees from the government and six SoS Lablets were welcomed by Xenefon Koutsoukos, the Principal Investigator (PI) at the VU Lablet, and Adam Tagert, the NSA SoS Technical Director. Adam reported on the winner of the [8th Annual Best Scientific Cybersecurity Paper Competition](#) (“Spectre Attacks: Exploiting Speculative Execution”) and encouraged attendees to submit a paper for the [9th Annual competition](#). He also presented details about HotSoS 2021, which will be hosted by NSA and held virtually from 13-15 April; HotSoS 2021 will encourage interaction among presenters and attendees and focus on Works-in-Progress (WiPs), posters, and student presentations.

NSA Champions Panel

The first presentation consisted of a panel discussion by NSA Champions Adam Tagert, Ahmad Ridley, and Mike Collins, from NSA’s Laboratory for Advanced Cybersecurity Research. Adam Tagert moderated the panel. The discussion session addressed such issues as communications with the projects for which they are the Champions, the genesis of the research projects, what it means to be a Champion, how they connect with other researchers about the projects, balancing short-term and long-term research contributions, and NSA’s grand challenges.

The Champions noted that it is important to expose the SoS researchers and graduate students to the mission problems or the mission use cases from a national security perspective since they are often different than what researchers might think about in a purely academic or industry case. It is also important to provide additional insight into use cases that drive the national security mission. That makes research more interesting for both the government sponsors and the Lablet researchers.

In addressing the question of how to ensure that others at NSA are made aware of Lablet research, a panelist said that it starts with understanding the challenges in an operational setting. Areas that are not working closely with the Lablets may not be ready for

the full research theory or practice being developed, but there are small milestones within the research efforts that could be relevant to challenges in other areas. Even if it is just reducing the amount of time needed to solve a problem, the contribution is worth it. Non-research organizations, for example, may need a way to transition or understand research-heavy concepts and ideas in a way that makes sense for their mission. The role of the Champions is to serve as a bridge. They suggested that if the Lablet research is broken into goals, certain milestones could be achieved such as a white paper, code or research level code that will generate some results.

The Champions agreed that by knowing where the academic and industry research is headed and the tough problems that NSA is trying to solve, Champions can break the research down and explain to mission customers and partners what type of environment they might need to create to be ready for that research when it becomes more widely available. The goal is to allow them to take advantage of research successes—even if it is just a small component that would help a mission partner, that short term goal is still important.

Invited Talk

Brad Martin, Formal Methods at Scale

The quarterly keynote presentation was by NSA’s Brad Martin. The topic was his work to increase the use of formal methods in security. Mr. Martin discussed the two workshops held in 2019 under the auspices of the National Science and Technology Council (NSTC) Special Cyber Operations Research and Engineering subcommittee (SCORE) that were designed to improve understanding of how the formal methods community, in partnership with sponsors and users, might achieve broader use of the technologies at increasing levels of scale. The workshops also sought to identify successes, barriers, opportunities, and challenges regarding the use of formal methods in cyber systems. The workshops explored tools, key enablers of success, emerging formal methods capabilities, and opportunities for key R&D investments. Workshop participants concluded that

tools, practices, and ecosystems are already facilitating commercial, government, and academic application of formal tools across many application domains and types of systems, but work remains to advance the scope, capability, and usability of the key formal methods technologies, tools, and practices. While the use of formal methods is increasing, the technologies are still at an early stage of development with respect to potential benefits to security, quality, and other kinds of assurance, as well as ancillary benefits to developing systems that are both readily adaptable and, on the basis of formal evidence, also readily recertified.

Lablet project presentations

The quarterly meeting included six technical talks, one from each Lablet. The talks were chosen on the quarterly meeting theme of empirical studies and simulation. The theme enables researchers and interested parties to gather with common interests to increase collaboration within the community.

Prasad Kulkarni (KU): Secure Native Binary Executions

Project: Secure Native Binary Executions

Professor Kulkarni described two projects that are under way to achieve the goal of developing a high-performance framework for client-side security assessment and enforcement for cots binary software. The overall project aims to provide greater control to the end-user to actively assess and secure the software they use. The first project assesses the security of software binaries, while the second project enhances the security of software binaries. The projects are designed to produce easy to use, reliable, and high-performance software tools for assessing and enforcing the security of binary software. One of the questions posed was how his research might help mitigate threats from supply chain attack which was recently highlighted by the SolarWinds compromise.

Kyle Crichton (CMU): How Do Home Computer Users Browse the Web?

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

Using data collected from the Security Behavior Observatory, a CMU CyLab initiative that has collected a wide array of system, network, and browser data from over 500 home Windows computer users over several years, researchers are seeking to better understand how users browse the web in order to identify where browsing breaks down and how to use this data to better mitigate or prevent some of the consequences. This study sought to identify how user behavior has changed compared to previous measurement studies, and researchers identified five notable changes in user behavior. These changes included, among others, an increase in the use of multiple browser tabs as well as an increase in the number of web pages visited each day. One point raised during the follow-up discussion was how this study could be used to educate users on how to avoid common pitfalls associated with web browsing.

Alisa Frick (ICSI): A Qualitative Model of Older Adults' Contextual Decision-Making About Information Sharing

Project: Operationalizing Contextual Integrity

The overall project 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. The research goal of this particular study is to develop a comprehensive understanding of older adults' information-sharing decision-making process and what contextual factors affect it. Dr. Frick described the methods and participants used in the study, the model flow that the researchers developed, implications, and future work. During the discussion period, she noted that the focus of the interviews and model flow was to quantify the impact and the relative importance of different factors.

Xenofon Koutsoukos (VU): Resilient Distributed Optimization and Learning in Networked Cyber-Physical Systems

Project: Foundations of CPS Resilience

This research deals with the system science of secure and resilient CPS and aims to develop a systematic body of knowledge with strong theoretical and empirical underpinnings to inform the engineering of secure and resilient cps that can resist not only known but also unanticipated attacks. This presentation focused on how normal agents can aggregate their neighbors' information to achieve the global objective even if some of the neighbors are adversarial. Professor Koutsoukos addressed vector consensus, distributed diffusion using centerpoint, and distributed multi-task learning in multi-agent environments, as well as simulations and empirical studies. The technical discussion following the presentation focused on the centerpoint method, its comparison to other methods, and possibly combining it with optimization algorithms.

Andy Meneely (RIT) NCSU: Deriving Vulnerability Discoverability Insights from a Penetration Testing Competition

Project: Predicting the Difficulty of Compromise through How Attackers Discover Vulnerabilities

The overall goal of this project is to provide actionable feedback on the discoverability of a vulnerability. The researchers' approach combines the attack surface metaphor and attacker behavior to estimate how attackers will approach discovering a vulnerability. In this presentation, Professor Meneely began by noting that the lifetime of a vulnerability is usually many years, and because vulnerabilities are typically small, they are easy to miss. He pointed out that discoverability aids risk assessment and behavior drives discoverability. Using over 14 TB of attacker behavior data gathered from National Collegiate Penetration Testing Competitions (CPTC), the RIT and NCSU researchers sought to answer multiple research questions. Professor Meneely presented some of their findings. In responding to a question about whether these findings can be used to improve the competition, he noted that the competition organizers have been able to see what people tend to do, and they are now able to use actual data to inform competition development.

Hussein Subai (UIUC): Accelerating Autonomous System Verification Using Symmetry

Project: Resilient Control of Cyber-Physical Systems with Distributed Learning

The project aims to bring together techniques from Machine Learning (ML) and formal verification to improve resiliency and risk-reduction in autonomous systems and CPS. This presentation addressed recent successes in CPS verification by the researchers at UIUC and the University of Texas at Austin. The presenter noted that safety analysis of autonomous systems is challenging, and one of the key challenges in analyzing autonomous systems is scalability. He addressed how symmetry can be used to enhance the scalability of autonomous systems verification and provided several examples of how that had been applied in the research. He summarized that they had designed caching reachsets and abstraction-refinement-based algorithms to accelerate autonomous systems verification and developed safety verification tools for autonomous systems that utilize symmetry, and that experimental results are showing orders of magnitude savings in verification time.



Summer 2021 Lablet

Quarterly meeting

The Summer 2021 Science of Security and Privacy (SoS) Quarterly Lablet meeting was hosted virtually by Carnegie Mellon University (CMU) on July 13-14, 2021. The virtual attendees from the government and the six SoS Lablets were welcomed by Jonathan Aldrich, the Principal Investigator (PI) at the CMU Lablet, and Heather Lucas, the SoS team lead in NSA's Laboratory for Advanced Cybersecurity Research. The theme of the Summer Quarterly was Artificial Intelligence (AI) and Machine Learning (ML).

SoS Hard Problems Panel

The SoS initiative is working on setting the direction for the future of the Science of Security. The Hard Problems, developed nearly a decade ago, are integrated into all Lablet research, but need to be reviewed in order to direct research going forward. Because AI and ML weren't part of the original Hard Problems and have developed significantly in the past ten years, the SoS community needs to consider how AI and ML can be applied to cybersecurity challenges. The Summer Quarterly kicked off with a panel discussion on the Hard Problems, one of the purposes of which was to elicit input from the SoS community. Moderator Adam Tagert and panel members Ahmad Ridley, Jonathan Aldrich, and Carl Landwehr encouraged attendees to share their research ideas and suggest areas for further exploration.

In order to frame the subject and prompt discussion, Dr. Ahmad Ridley opened the session with a presentation on the National Security Commission on Artificial Intelligence (NSCAI) Final Report (available [here](#)).

Some of the questions raised during the follow-on discussions included:

- What are the national questions for AI and cybersecurity?
- How do we ensure the security of ML and AI and how to develop trust in such systems?
- How do we address the ethical issues surrounding ML and AI?
- How do individuals protect themselves from AI?
- If ML and AI are the next revolution, how do we use history to predict possible changes?

Technical Presentations

Events and Stories: NLP toward Secure Software Engineering

Hui Guo, NCSU

The research interests described in this presentation are in natural language processing in events and stories in natural language text; The researchers looked at breach reports targeting text related to software development and applied the research questions: How can we effectively extract targeted events from text?; How can we effectively extract targeted event pairs from text?; and How can we effectively extract targeted stories from text? Professor Guo presented the research methodology, representative examples, and findings, and then proposed future work, including more reliable extraction from

low-quality text, pre-defined event types, deeper understanding of event relations, and how story understanding helps.

Trustworthy, Robust, and Explainable ML: Joined at the Hip

Matt Fredrikson, CMU

Under this Lablet proposal, researchers aim to analyze and quantify the extent to which attacks can be mounted using implementation-level explanations. Professor Fredrikson introduced attribution methods and showed how they can be used to explain and quantify attacks. He also discussed security pitfalls of explanations and introduced robustness, a practical remedy to these pitfalls. He summarized by noting that faithful attribution methods are useful for gaining insight into attacks and measuring vulnerability. In certain cases, attackers can manipulate attributions with feature perturbations. Lipschitz-continuous models mitigate this threat and are often more explainable to begin with. Globally robust nets are an efficient, effective way to achieve strong Lipschitz continuity.

Beyond L_p balls: Attacks on real-world uses of machine learning

Michael Reiter, Duke

This presentation dealt with attacks on practical uses of ML in a setting where attacks aren't constrained by nearness, noting that traditional attack approaches and traditional defenses may not apply. The attack scenario he described was fooling face recognition, and he provided several approaches and methodologies. His attack scenarios were different from previous ones in that the adversarial examples were, among other differences, far in L_p space from source inputs. The takeaways from his presentation were that real-world applications of ML are vulnerable; small L_p distance attacks and defenses may not apply; and defenses have a long way to go. In response to a question, he pointed out that most of the attacks he described were white-box attacks--if the classifier doesn't tell us anything other than that it was malware, for example, it is a significant challenge. This presentation relates to the CMU project Securing Safety-Critical Machine Learning Algorithms.

Flexible Mechanisms for Remote Attestation

Perry Alexander, KU

The goals of this research deal with formal semantics of trust—a definition of trust sufficient for evaluating systems; a verified remote attestation infrastructure—verified components for assembling trusted systems; enterprise attestation and appraisal; and sufficiency and soundness of measurement. Professor Alexander provided a remote attestation example that included three attestation managers and seL4 implementation infrastructure. This presentation builds upon the KU project entitled Scalable Trust Semantics and Infrastructure.

A First Look at Soft Attestation in Android Apps

Abbas Razaghpanah, ICSI

Dr. Razaghpanah began by defining remote attestation as a way for a remote party to ensure that the client software and the platform it runs on are trustworthy, and noted that this research provides evidence that shows apps are abusing remote attestation to hide bad data collection practices and to evade app store policy enforcement. The research questions include why mobile apps use soft attestation; how it works on mobile platforms and the heuristics; what apps are protecting when using attestation; and how differently apps would behave when not “watched.” He proposed answers to these questions during the presentation and noted that the researchers used these insights to improve their testing apparatus. He said that proper attestation is very hard with the android platform, and unless Google can compel all their OEMs to put a TPM on the device, strong universal attestation won't happen on Android. This research is under the ICSI project Scalable Privacy Analysis.

Analytics of Cybersecurity Policy: Value for Artificial Intelligence?


Nazli Choucri, MIT

To provide context for her presentation, Professor Choucri briefly described the ongoing VU project, Policy Analytics for Cybersecurity of Cyber-Physical Systems. She noted that cybersecurity policies and guidelines are in text form, and the project aims to introduce analytics for cybersecurity policy of cyber-physical systems in order to produce tools for analytics of cybersecurity policy; the project uses the NIST Cyber Security Framework applied to a smart grid as a testbed. She discussed the AI global ecology and pointed out that analytics of cyber security policy for cyber-physical systems is extremely difficult for any system operation to apply the cyber policies that are given to them because there are impediments to just using words. AI applies in that everything the researchers have done by hand must be automated.

Emulating Cybersecurity Simulation Models Using Metamodels for Faster Sensitivity Analysis and Uncertainty Quantification

Michael Rausch, UIUC

Dr. Rausch noted that because cybersecurity models are often large, have long execution times, and many uncertain input variables, this research approach uses metamodels. Traditional approaches to uncertain input variables include Sensitivity Analysis (SA) and Uncertainty Quantification (UQ). A metamodel is a model designed to emulate the behavior of a model referred to as the base model. Given the same input, the metamodel should produce the same or similar output as the base model. A metamodel trades some accuracy for faster runtimes. Once constructed, SA and UQ can be applied to the metamodel. He discussed how to construct metamodels and collect training and testing data, and then provided test cases to evaluate the metamodels and the results of the modeling. He emphasized that metamodeling is a powerful approach for analyzing cybersecurity models that are slow running and have many uncertain variables. He also concluded that metamodels approximate the base model behavior relatively accurately, while being much faster than the base model. Use of metamodels was explored in the UIUC project, Monitoring, Fusion, and Response for Cyber Resilience.



Fall 2021 Lablet

Quarterly meeting

The Fall 2021 Science of Security and Privacy (SoS) Quarterly Lablet meeting was hosted virtually by the University of Kansas (KU) on November 15-16, 2021. The virtual attendees from the government and six SoS Lablets were welcomed by Perry Alexander, the Principal Investigator (PI) at the KU Lablet, and Heather Lucas, the National Security Agency (NSA) SoS Initiative lead. The Quarterly included two invited talks, six Lablet project talks, and a presentation on Challenge Problem sponsored by the National Science Foundation (NSF).

Invited Talks

Dr. Robert Runser, NSA Research Directorate Technical Director Unclassified Challenges Facing Cybersecurity

Dr. Runser opened his presentation by noting that cybersecurity is embedded in every pillar of research that NSA is currently doing in the Research Directorate (RD), which demonstrates how important and foundational cybersecurity is to national security and to all the future systems. He described the NSA missions and how RD contributes to advancing those missions by conducting world-class scientific research to develop new technologies and innovative techniques. RD's five technical focus areas, all of which are interlinked, are 1) Future computing systems; 2) Crypto Mathematics; 3) Science of analysis; 4) SIGINT collection research; and 5) Cybersecurity. He noted that multidisciplinary research is critical to cybersecurity, and since human factors make up our cybersecurity defenses, that often involves bringing in operational psychologists, people that understand how humans operate and can instrument in various trials, human experiments, and red team activities to better understand how cyber defenses can react to human elements and how networks can be more responsive to human errors. He also said that data science is a critical discipline to cybersecurity with the amount of network logs, net flow data, and other types of information, and that human-machine teaming and data science have become critical foundational aspects of the type of cybersecurity research NSA conducts. Within cybersecurity research, he cited large-scale graph analytics, malware and software reverse engineering analytics, and secure operating systems as areas of focus. He commented that the biggest challenge faced today is sifting through large-scale graphs. While graphs used to be constructed simply to look for adversaries, graph analytics are now used to understand the infrastructure that attackers are using to establish the command and control networks before an attack occurs. If we can spot their entry and exit points before they penetrate the system, we can block their attacks and potentially trace them back to their source and perform attribution. He believes that this will dominate much of the research NSA will do in the coming years, trying to optimize graph analytics, adapt it to the cyber domain, and understand how to ensure the right observables go into the high-performance computing environment. High-performance computing systems working in tandem with graphic analytics will drive much research. Within Crypto Mathematics, he cited quantum-resistant cryptography, lattices/modules, coding theory, multivariate quadratics, hash functions, zero-knowledge proofs, and elliptic curve isogenies as areas of interest. He concluded by describing both the academic and career opportunities that are available through the Research Directorate.

The follow-up discussion dealt with such areas as 5G network challenges, implementation of crypto protocols that have been developed, and supply chain issues.

Dr. Natarajan Shankar, SRI International Composing High-Assurance Software with Evidential Tool Bus

Dr. Shankar began his talk by discussing the software stack, describing it as one of mankind's greatest engineering achievements, but whose power comes with a price- a large attack surface where bugs can have serious consequences leading to software errors and cybercrime. He said that, unlike other engineering artifacts, software supports greater flexibility, resilience, and versatility in the design and maintenance of a system, but we lack a mature engineering discipline of principled software construction, and attacks can wreak havoc on a global scale. To counter this, we need to invest in a discipline that provides composable assurance. He addressed what can go wrong in software design and what can lead to vulnerabilities, and pointed out that while formal modeling and analysis is practical and even necessary, it is not a panacea. He called for software to be designed hand-in-hand with assurance artifacts that are verifiable by clients (or trusted third parties) and that software designs be centered around software architectures (models of computation and interaction) that deliver efficient arguments for isolation and composition. He addressed Evidence-Based Assurance, saying that an assurance case is a formal method for demonstrating the validity of a claim by providing a convincing argument together with supporting evidence. It is a way to structure an argument to help ensure that top-level claims are credible and supported. The Evidential Tool Bus (ETB2) is a distributed tool integration framework for constructing and maintaining claims supported by arguments based on evidence generated by static analyzers, dynamic analyzers, satisfiability solvers, model checkers, and theorem provers. The key ideas associated with ETB2 are: Data as a metalanguage; Denotational and operational semantics; Interpreted predicate for tool invocation, and uninterpreted predicates for scripts; Datalog inference trees as proofs; Git as a medium for file identity and version control; and Cyberlogic, a logic of attestations, to authenticate the claims and authorize the services. He provided technical details on the ETB2, and described goals, challenges, and approaches as well as key issues in its application. He concluded by describing a Software Proof of Virtues (SPOV), noting that software failures and cyber-attacks weaken trust and the current strategy of applying larger and larger band-aids is only fueling an arms race. He believes that we have the tools and insights to build the infrastructure of trust in software from the ground up to include: software development lifecycle workflows that continuously maintain both process and outcome-based assurance evidence; tools and models that support designs annotated with traceable ontic information that are founded on efficient arguments; verified platforms and services whose integrity is certified by audit logs and audits; and composable assurance cases validating intent, correctness, and innocuity.

Lablet project presentations

KU Project Talk

Heechul Yun (KU): Micro-Architectural Attacks and Defenses

Project: Side-Channel Attack Resilience

In addressing micro-architectural attacks, Professor Yun said that they are software attacks on hardware that can have multiple adverse effects. His talk focused on three elements: a new contention-based covert channel; a new DoS attack; and a hardware defense mechanism for DoS attacks. He described Spectre Rewind, a novel contention-based covert channel that transmits secret speculative instructions to past instructions through non-pipelined functional units on a single hardware thread, and bypasses all existing defenses against cache or SMT-based covert channels, and summarized both its benefits and limitations. In describing a new DoS attack, he noted that DoS attacks are more effective when an attacker's memory requests are processed slowly. He and his team developed memory-aware DoS attacks that target a subset of DRAM banks, and evaluation results show significantly improved attack efficiency on the tested embedded computing platforms. Finally, he addressed the Bandwidth Regulation Unit (BRU), which was created in response to the fact that DoS attacks are possible because of unregulated access to the shared resources and the resultant need for a simple low overhead mechanism to regulate access to shared resources. BRU is a synthesizable hardware IP that regulates memory traffic at the source core that demonstrates the feasibility of fast and predictable processors.

UIUC Project Talk

Matt Caesar (UIUC), Kevin Jin (University of Arkansas) and Gabriella Xue (UIUC)

An Automated Synthesis Framework for Network Security and Resilience

Project of the same name

This project was described as building a rigorous methodology for the science of security and addressing challenges in applying science to security. The specific outcome of the project is a resilient network architecture with a specific focus on network data flow. Their research approach leverages network synthesis to automate experiments and apply results; enables practical uses, including deriving patches and automating configuration; and builds upon mathematics. There are three task plans: network control synthesis; network software analysis and modeling; and resilient and self-healing network applications. The talk focused on specific work the team is doing on self-driving service provider infrastructures, resilient power systems, and supporting teaching and research with virtualized IoT systems.

VU Project Talk

Lillian Ratliff (University of Washington), Eric Mazdumar (Caltech), S. Shankar Sastry (UC Berkeley)

Digital Transformation of Social Systems: A New Hope: Hackers Strike Back

Project: Mixed Initiative and Collaborative Learning in Adversarial Environments

This presentation addressed the digital transformation of societal systems (IoT, AI, the Cloud, Big Data) and raised the question of how AI and edge computing fit into intelligent transportation systems. Intelligent systems require rethinking ML, since classical ML assumes

the past is representative of the future, an unintended consequence of which is feedback reinforced bias. An emerging new domain is learning-enabled intelligence. Their research 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 and trying to analyze and design attacks on learning in multiplayer games. They discussed how they have applied a model to their work and their findings.

NCSU Project Talk

Jeffrey Carver and Matthew Armstrong, University of Alabama
Guidelines for Reporting Scientifically Rigorous and Valid Cyber Security Research

Project: Development of Methodology Guidelines for Security Research

Because cybersecurity is a complex, maturing field and there is no consistent, community accepted standard for reporting cyber security research, this project was initiated to facilitate the development of the science of cyber security by developing guidelines that bring validity and rigor to cyber security reporting. The researchers described their methodology and their accomplishments thus far, including the development of Version 1.0 of the Guidelines for Scientific Reporting in Cyber Security. As part of the presentation, researchers organized an interactive session to gather early feedback on the guidelines.

CMU Project Talk

David Garlan (CMU)
Model-Based Explanation for Automated Decision Making

Project: Model-Based Explanation for Human-in-the-Loop Security

Professor Garlan provided context for the problem, noting that while autonomy is increasingly important for modern systems, many systems require a combination of automated and human involvement to handle security attacks. The problem, therefore is how to create effective coordination by deciding which tasks to allocate to the system vs the human, ensuring that humans have confidence in automated actions, enabling correction of errors, improving automation by learning from humans, and understanding what the system does. He described the current research approach and progress, and discussed specific initiatives underway.

ICSI Lablet Project Talk

Julia Bernd (ICSI)
Perspectives of Stakeholders in Data Governance

Project: Governance for Big Data

This presentation began by noting that while stakeholders are dealing with the same data, they have different goals for the use of the data, different technological capabilities and resources, and different approaches to data management and governance. For the purposes of this research, data stakeholders included technology developers, technology and data platforms, data users, tech/data regulators, and data subjects. This talk presented research findings on ad developers how they approached user privacy as well as a case study on stakeholders in health app privacy. Also presented was the role of privacy champions on software teams and how to support them.

Special Presentation

NSF-funded C3E Challenge Problem Opportunity

Dan Wolf and Don Goff, Cyber Pack Ventures, Inc.

This presentation described a National Science Foundation-funded grant and offered attendees an opportunity to contribute to a challenge problem. The Special Cyber Operations Research and Engineering (SCORE) committee sponsors a yearly Computational Cybersecurity in Compromised Environments (C3E) Workshop, which includes follow-up challenge problems, each of which goes along with the theme of the workshops. NSF has provided funding for honoraria for research on topics related to the workshop challenge problem. All of the submissions go through a peer-review process. The theme of the 2021 C3E workshop was real issues in securing the supply chain, with emphasis on software security, and the 2022 challenge problem options shown below all relate to that topic:

- Supply chain software static analysis coverage.
- Artificial intelligence applied to supply chain cybersecurity
- Computational victimology for developing risk models for supply chain cybersecurity





Promoting Rigorous Scientific Principles

In 2021 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 9th Annual Best Scientific Cybersecurity Paper Competition.

The National Security Agency's Research Directorate selected "On One-way Functions and Kolmogorov Complexity," written by Yanyi Liu, Cornell University and Rafael Pass, Cornell Tech. as the winner of its 9th

Annual Best Cybersecurity Research Paper competition. The paper was published at the 2020 IEEE Symposium on Foundations of Computer Science.

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.


9th Best Paper Competition



Winning Paper



Yanyi Liu



Raphael Pass


The National Security Agency and Science of Security selected “On One-way Functions and Kolmogorov Complexity,” written by Yanyi Liu, Cornell University and Rafael Pass, Cornell Tech. The paper was published at the 2020 IEEE Symposium on Foundations of Computer Science. One-way functions (OWF) are a key underpinning in many modern cryptography systems, and were first proposed in 1976 by Whitfield Diffie and Martin Hellman. These functions can be efficiently computed but are difficult to reverse, as determining the input based on the output is computationally expensive. OWFs are vital components of modern symmetric encryptions, digital signatures, authentic schemes and more. Until now, it has been assumed that OWF functions exist even though research shows that they are both necessary and sufficient for much of the security provided by cryptography.

Honorable Mention


Receiving honorable mention was the paper “Retrofitting Fine Grain Isolation in the Firefox Renderer” written by Shравan Narayan, Craig Disselhoe, Tal Garfinkel, Nathan Froyd, Sorin Lerner, Hovav Shacham and Deian Stefan. This paper was originally published at the USENIX Security Conference 2020 and provides a security solution for use in the Firefox web browser while also demonstrating that that technology can be utilized for other situations. The solution, RLBox, is a culmination of many advances that enable software to securely use software components, such as libraries, which have not been verified as trustworthy. RLBox has been incorporated into Firefox 95.




Shравan Narayan
UC, San Diego




Craig Disselhoe
UC, San Diego



Tal Garfinkel
Stanford University




Nathan Froyd
Mozilla



Sorin Lerner
UT, Austin



Hovav Shacham
UC, San Diego



Deian Stefan
UC, San Diego

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 collectively reviewed 34 nominated papers. After review and ranking, the Distinguished Experts forwarded their recommendations to NSA Director of Research Gilbert Herrera for final selection.

The Distinguished Experts were:

- Dr. Whitfield Diffie, Cybersecurity Advisor
- Prof Kathleen Fisher, Tufts University
- Dr. Dan Geer, In-Q-Tel
- Dr. Eric Grosse, Cybersecurity Advisor
- 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, Cybersecurity Advisor
- Ms. Mary Ellen Zurko, MIT Lincoln Laboratory

The 10th Annual Best Scientific Cybersecurity Paper Competition is open for nominations for papers published during calendar year 2021 in peer-reviewed journals, magazines, or technical conferences that show an outstanding contribution to cybersecurity science.



Growing the Science of Security Community

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 2021, despite the continuing pandemic, the SoS continued to grow the Science of Security community of interest.

For the eighth year, SoS sponsored the Hot Topics in the Science of Security Symposium (HotSoS) which was held virtually and hosted by the NSA from 13-15 April 2021. Over 1200 individuals registered for HotSoS '21, and more than 625 participated over the three days. The participants, a mix of government, academia, and industry, came from 36 countries. HotSoS 2021 was designed to encourage interaction among presenters and attendees and focus on WiP, posters, and student presentations.

The SoS-Virtual Organization (SoS-VO), a longstanding initiative designed to grow the Science of Security community, reached over 2000 members in 2021. 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 nearly 1800 subscribers. There also continued to be over 300 Facebook postings to more than 100 members.

The Vanderbilt University Labet created a new program to grow the SoS community in 2021 with its Cyber Makerspace initiative. Cyber Makerspace is a simulated, networked environment that facilitates instruction on cyber-physical systems, their security and related topics while reducing cost and complexity. The approach will facilitate reaching audiences from traditionally underrepresented groups.

Details on HotSoS, Cyber Makerspace, and outreach initiatives are found in the following pages.



Hot SoS

The National Security Agency (NSA) virtually hosted the 8th Annual Symposium on the Science of Security (HoTSoS), from 13-15 April 2021. The General Chair was Adam Tagert (NSA) and Program Co-Chairs were Özgür Kafali (University of Kent) and Ahmad Ridley (NSA). HotSoS brings together researchers from diverse disciplines to promote the advancement of work related to the Science of Security and Privacy initiative (SoS) and features a mix of invited keynotes, Works-in-Progress (WiP) discussions, presentations of already published work, and student presentations. The format of HotSoS was revised this year to focus on discussions of ongoing research. This year's virtual event provided the opportunity for those who might otherwise be unable to attend to engage with other SoS researchers. Over 1200 individuals registered for HotSoS '21, and more than 625 participated over the three days. The participants, a mix of government, academia, and industry, came from 36

countries. HotSoS 2021 was designed to encourage interaction among presenters and attendees and focus on WiP, posters, and student presentations. In addition to 4 keynote presentations, HotSoS 2021 included presentations of 12 published papers, 8 WiP manuscripts, 5 student presentations, and 15 posters which, in total, represented the work of 116 authors from 37 universities. In addition to the keynotes and posters, there were 7 topical sessions which included paper presentations and WiP discussions, and a student presentation session. In keeping with the goal of collaborative community engagement, HotSoS 2021 again featured WiPs which provide an opportunity for authors to get early feedback on a research direction, technology, or idea before it has been fully evaluated, or to discuss systems in an early, pre-prototyping phase; submissions were restricted to session attendees. The agenda also included a special session on Science of Security Hard Problems.



Keynote Presentations

1. "Securing Data in Clouds: Making the Most of Trusted Hardware,"

Nick Felts, NSA

In this talk, Mr. Felts focused on how trusted hardware can be incorporated into the cloud infrastructure to protect both keys and data in use. He provided examples of trusted hardware, addressed assumptions for enclaves, and discussed data protection at scale. He noted that meaningful encryption, at scale, is difficult for a number of reasons, and suggested that trusted hardware can help protect data at scale. He concluded that trusted hardware has the potential to protect data-in-use within cloud computing and that protecting keys-in-use paves the way for data-in-use, both in capabilities and knowledge, and becomes more practical as secure enclaves.

2. "Spectre Attacks: Exploiting Speculative Execution - and why the heck is the computer speculating anyway?"

Werner Haas, Cyberus Technology

Werner Haas was one of the authors of the paper that won the SoS 8th Annual Best Scientific Cybersecurity Paper competition and a related paper that received an Honorable Mention at the 7th Annual Paper Competition. His presentation addressed the earlier research in the context of today, including what has changed or what is different in the Spectre class of vulnerabilities. He began by describing the iron law of processor performance (performance associated with the time required to perform a task) and then updated the equation to include Cycles Per Instruction (CPI). He continued his presentation with a discussion of why Spectre wasn't discovered earlier and included ongoing research.

3. "Why rigorous underpinnings for cyber security education and training matter? Experiences from CyBOK: the Cyber Security Body of Knowledge"

Awais Rashid, University of Bristol,

This presentation focused on a science of security approach to cyber education and training and the development of such an approach in the form of CyBOK. Dr. Rashid noted that as a topic and as a discipline cybersecurity does not have the common foundation that other disciplines have to build education and training programs. CyBOK is designed to codify foundational and generally recognized knowledge in cyber security following broad community engagement nationally and internationally. It also provides a guide to the body of knowledge, the focus of which is on the established foundation on the subject rather than on everything that has ever been written or on still-emerging, nascent, topics. The guiding principles of the project are that it is 1) an international effort; 2) for the community by the community; 3) open and freely accessible; and 4) transparency. Dr. Rashid identified the major knowledge areas, CyBOK use cases, and the mapping framework and discussed potential applications for CyBOK.

4. "Working with Academia at the UK National Cyber Security Centre,"

Paul Waller, UK's Government Communications Headquarters (GCHQ)

Mr. Waller provided background on the NCSC mission, products and services, and programs, and then discussed research areas and academic partnerships, noting that NCSC is investing for the long term. He described the following academic partnership programs:

- RISCs—Research Institute for Sociotechnical Cyber Security
- RIVeTSS—Research Institute for Verified Trustworthy Software Systems
- RITICS – Research Institute in Trustworthy Interconnected Cyber-physical Systems
- RISE – Research Institute in Secure Hardware and Embedded Systems
- ATI – Alan Turing Institute

Topical Sessions

There were seven topical sessions each organized around a theme that included a mix of published papers and facilitated WiP discussions. The WiP sessions were introduced in HotSoS 2019 and have been a part of both HotSoS 2020 and 2021. 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. Since the manuscripts were in the early stages, access to the submissions remains confidential and is restricted to session attendees.

Theme: Cloud Security

Papers

"Formal Foundations for Intel SGX Data Center Attestation Primitives"

Muhammed Usama Sardar, Rasha Faqeh, and Christof Fetzer, TU Dresden

Hardware-based Trusted Execution Environments (TEEs) such as Intel SGX offer potential solutions for protecting data in use, and Intel has recently offered third-party attestation services, called Data Center Attestation Primitives (DCAP), for a data center to create its own attestation infrastructure. This paper deals with the work the team is doing with respect to formal proof for Intel's DCAP. They propose an automated and rigorous approach to specify and verify the remote attestation under the assumption that there are no side-channel attacks and no vulnerabilities inside the enclave. The team's findings include the discovery of various discrepancies in the literature. The presenter also discussed potential future work such as removing simplifying assumptions and considering side-channels.

"A Secure and Formally Verified Linux KVM Hypervisor"

Shih-Wei Li, Xupeng Li, Ronghui Gu, Jason Nieh, and John Zhuang Hui, Columbia University

This paper introduces microverification, an approach that decomposes a commodity hypervisor into a small core and a set of untrusted services in order to prove security properties of the entire hypervisor by verifying the core alone. Using microverification, the researchers retrofitted the Linux KVM hypervisor with only modest modifications to its codebase. Using Coq, they proved that the hypervisor protects the confidentiality and integrity of VM data, while retaining KVM's functionality and performance. This work is the first machine checked security proof for a commodity multiprocessor hypervisor.

Work-In-Progress Discussion Session

"JavaScript Attacks"

Ross Copeland and Drew Davidson, University of Kansas

Theme: To Err is to be Human

Work-In-Progress Discussion Sessions

"Election Security"

Natalie Scala and Josh Dehlinger, Towson University

Paul L. Goethals, United States Military Academy

"Phishing"

Dennis Roellke, Salvatore Stolfo, and George Litvinov, Columbia University

Shlomo Herschkopp, Allure Security

Mark Seiden, Internet Archive

Theme: Flanking the Defense

Papers

"Leveraging EM Side-Channel Information to Detect Rowhammer Attacks"

Zhenkai Zhang, Texas Tech University

Zihao Zhan, Texas Tech University, Vanderbilt University

Daniel Balasubramanian, Peter Volgyesi, and Xenofon Koutsoukos, Vanderbilt, University

Bo Li, University of Illinois at Urbana-Champaign

This paper proposes a novel approach named RADAR (Rowhammer Attack Detection via A Radio) that leverages certain electromagnetic (EM) signals to detect rowhammer attacks. The researchers found that there are recognizable hammering-correlated sideband patterns in the spectrum of the DRAM clock signal, and because such patterns are inevitable physical side effects of hammering the DRAM, they can "expose" any potential rowhammer attacks. The patterns of interest may become unapparent due to the common use of Spread-Spectrum Clocking (SSC) in computer systems, but the research team used a de-spreading method that can reassemble the hammering-correlated sideband patterns scattered by SSC. There has been little prior work that uses physical side-channel information to perform rowhammer defenses, and this appears to be the first investigation on leveraging EM side-channel information for this purpose.

"Counting Broken Links: A Quant's View of Software Supply Chain Security"

Dan Geer, Bentz Tozer, and John Speed Meyers, In-Q-Tel

The number of supply chain attacks and the count of reports about these attacks continue to increase—open-source libraries, source

code, compiled code, and released software are all vulnerable. This paper presents the results of the researchers' review of news releases, the internet, and white papers to assess the extent of software supply chain. One of the driving questions behind the research was trying to quantify the risks associated with the attacks and then getting into costs; future work will look at quantifying severity as well. They concluded that it is essential to increase transparency, automate assessments, monitor all activities, and illuminate risks.

Work-In-Progress Discussion Session

"Device Profiling"

Tushar Jois and Aviel D. Rubin, Johns Hopkins University

Claudia Moncaliano, Johns Hopkins University Applied Physics Lab

Khair Henderson, Morgan State University

Theme: Humans aren't only Users

Papers

"Can Advanced Type Systems be Usable? An Empirical Study of Ownership, Assets, and Tystate in Obsidian"

Michael Coblenz, University of Maryland

Jonathan Aldrich, Brad Myers, and Joshua Sunshine, Carnegie Mellon University

Obsidian is a new tystate-oriented programming language that uses a strong type system to rule out some of the vulnerabilities associated with some blockchain programs. Ownership, tystate, and assets, which Obsidian uses to provide safety guarantees, have not seen broad adoption together in popular languages and result in significant usability challenges. This paper presents the results of an empirical study comparing Obsidian to Solidity, which is the language most commonly used for writing smart contracts today. The researchers observed that Obsidian participants were able to successfully complete more of the programming tasks than the Solidity participants. They also found that the Solidity participants commonly inserted asset-related bugs, which Obsidian detects at compile time.

"Adversarial Thinking - Teaching Students to Think Like a Hacker"

Frank Katz, Georgia Southern University

Training students in cybersecurity requires teaching technical tools such as firewalls, VPNs, IDS/IPS, and packet sniffers, but the premise of the presentation is that it is not enough to teach technology. It is also important to understand hacker motivations and learn how hackers are thinking and develop the student's abilities to anticipate the strategic actions of cyber adversaries, including where, when, and how they might attack, and their tactics for avoiding detection. This paper describes the content and implementation of a 6 hour 15 minute (5 class sessions) module in Adversarial Thinking in a Network Security course, the students' perceptions of the value and importance of the module as a result of their anonymous responses to a survey on the module, and the statistical results of a Data Breach Pretest-Posttest Assessment to measure how well they understood the concepts involved in Adversarial Thinking as part of learning cybersecurity.

Work-In-Progress Discussion Session

"Practices in Software Development"

Laysan Nurgalieva, and Gavin Doherty, Trinity College Dublin, Ireland,

Alisa Frik, University of California, Berkeley

Theme: Saving the Physical World from Cyber

Papers

“Verified Hardware/Software Co-Assurance: Enhancing Safety and Security for Critical Systems”

David Hardin, Collins

For security-critical architecture and design, identification of the attack surface has emerged as a primary analysis technique. This paper describes the identification of and mitigation against attacks along that surface, using mathematically-based tools. The motivation behind these efforts is emerging application areas, such as assured autonomy, that feature a high degree of network connectivity, require sophisticated algorithms and data structures, are subject to stringent accreditation/certification, and encourage hardware/software co-design approaches. The researchers focus first on software implementation, but are extending to hardware as well as hardware/software co-designs. The paper describes how they utilized RAC (Restricted Algorithm C) to translate their JSON filter to ACL2, presents proofs of correctness for its associated data types, and describes validation and performance results obtained through the use of concrete test vectors.

“überSpark: Practical, Provable, End-to-End Guarantees on Commodity Heterogeneous Interconnected Computing Platforms”

Amit Vasudevan, Anton Dimov Hristozov, and Bruce Krogh, Software Engineering Institute, Carnegie Mellon University

Petros Maniatis, Google Research

Ruben Martins, Computer Science Department, Carnegie Mellon University

Raffaele Romagnoli, Department of Electrical and Computer Engineering, Carnegie Mellon University

Today’s computing ecosystem, comprising Commodity Heterogeneous Interconnected Computing (CHIC) platforms, is increasingly being employed for critical applications, consequently demanding fairly strong end-to-end assurances. However, the generality and system complexity of today’s CHIC stack seem to outpace existing tools and methodologies towards provable end-to-end guarantees. This paper describes the team’s on-going research, and presents überSpark (<https://uberspark.org>), a system architecture that argues for structuring the CHIC stack around Universal Object Abstractions (üobjects), a fundamental system abstraction and building block towards practical and provable end-to-end guarantees. The paper addresses the CHIC stack challenges, illustrates design decisions, describes the überSpark architecture, presents the foundational steps, and outlines on-going and future research activities. The researchers expect überSpark to retrofit and unlock a wide range of unprecedented end-to-end provable guarantees on today’s continuously evolving CHIC stack.

Work in Progress Discussion Session

“Railway Infrastructure”

Himanshu Neema and Xenofon Koutsoukos, Vanderbilt University
Leqiang Wang, CheeYee Tang and Keith Stouffer, National Institute of Standards and Technology

(This paper was subsequently published. See Vanderbilt University’s project Multi-Model Test Bed for the Simulation-Based Evaluation of Resilience in Section 1.)

Theme: Go Where I Send Thee

Papers

“ZeRØ: Zero-Overhead Resilient Operation Under Pointer Integrity Attacks”

Mohamed Tarek, Miguel Arroyo, and Evgeny Manzhosov, Columbia University

Simha Sethumadhavan, Columbia University, Chip Scan

Low performance overhead and convenience are key to widespread adoption of security techniques. This paper describes ZeRØ, a hardware primitive for resilient operation when pointers are targeted with zero overhead. ZeRØ enforces code and data pointer integrity with minimal metadata. ZeRØ incurs 0% performance degradation compared to 14% for the state-of-the-art ARM PAC when applied to its full extent. ZeRØ matches or offers better security guarantees than ARM’s PAC and Intel’s CET. Moreover, hardware measurements show that ZeRØ can be implemented with minimal latency/area/power overheads.

“Insights for Systems Security Engineering from Multilayer Network Models”

Adam Williams, Gabriel C. Birch, Susan Caskey, Elizabeth Fleming, Thushara Gunda, Thomas Adams and Jamie Wingo, Sandia National Laboratories

Next-generation systems security engineering approaches for High Consequence Facilities (HCF) need to address challenges stemming from complex risk environments, innovative adversaries, and disruptive technologies. This paper focuses on key insights from complexity, systems, and network theories—and support from Subject Matter Expert (SME) elicited empirical data—that support using a multilayer network model for HCF security to address these challenges. Early results of modeling HCF security system performance in terms of multilayer network characteristics seemed to meet the needs expressed in the SME data, provide a suite of mathematically tractable metrics to describe more complex security behaviors, and align with characteristics related to incorporating systems security engineering into the future of systems engineering.

Work-In-Progress Discussion Session

“Studying App Reviews”

Vaibhav Garg, Hui Guo, and Munindar Singh, North Carolina State University

Nirav Ajmeri, University of Bristol

Theme:

Potpourri for 1000

Papers

“AI-Powered Ransomware Detection Framework”

Subash Poudyal and Dipankar Dasgupta, University of Memphis

Various approaches have been proposed to combat ransomware, and while there are commercial tools available in the market for ransomware analysis and detection, their performance is questionable. This paper aims at proposing an AI-based ransomware detection framework and designing a detection tool (AIRaD) using a combination of both static and dynamic malware analysis techniques. Dynamic binary instrumentation is done using PIN tool, function call trace is analyzed leveraging Cuckoo sandbox and Ghidra. Features extracted at DLL, function call, and assembly level are processed with NLP, association rule mining techniques and fed to different machine learning classifiers. Support vector machine and Adaboost with J48 algorithms achieved the highest accuracy of 99.54% with 0.005 false-positive rates for a multi-level combined term frequency approach.

“On Managing Vulnerabilities in AI/ML Systems”

Jonathan Spring, April Galyardt, and Allen Householder, Community Emergency Response Team/Coordination Center; Software Engineering Institute, Carnegie Mellon University
Nathan M. VanHoudnos, Software Engineering Institute, Carnegie Mellon University

This paper explores how the current paradigm of vulnerability management might adapt to include machine learning systems through a thought experiment: what if flaws in machine learning (ML) were assigned Common Vulnerabilities and Exposures (CVE) identifiers (CVE-IDs)? The researchers consider both ML algorithms and model objects. The hypothetical scenario is structured around exploring the changes to the six areas of vulnerability management: discovery, report intake, analysis, coordination, disclosure, and response. While algorithm flaws are well-known in academic research community, there is no apparent clear line of communication between this research community and the operational communities that deploy and manage systems that use ML. The thought experiments identify some ways in which CVE-IDs may establish some useful lines of communication between these two communities. In particular, it would start to introduce the research community to operational security concepts, which appears to be a gap left by existing efforts.

Work-In-Progress Discussion Session

“Attestation and Game Theory”

Shanto Roy, Salah Uddin Kadir, and Aron Laszka, University of Houston
Yevgeniy Vorobeychik, Washington University in St. Louis

Student Presentations:

1. Analysis of the Impact of Varying Statistical RF Fingerprint Features on IoT Device Classification

Asia Mason, Morgan State University

Ms. Morgan noted that the growth of the Internet of Things (IoT) has also led to a growth in the exploitation of vulnerabilities of interconnected devices and created a need for security mechanisms appropriate for the IoT domain. She suggested RF Fingerprinting as a promising security technique that uses the physical characteristics of a wireless signal to identify a device. These characteristics, known as features, are combined to create an RF fingerprint and are then run through a machine learning classification algorithm for device identification. She

presented her work analyzing three sets of fingerprints that were tested with 25 classification models, noting that the results thus far show that using feature combinations specific to a device has an impact on the RF fingerprint performance.

2. Maze: A Secure Cloud Storage Service Using Moving Target Defense and Secure Shell Protocol (SSH) Tunneling

Vasco Xu and Sherif Khattab, University of Pittsburgh

The authors presented MAZE, a secure cloud storage system built from simpler security primitives (virtual machines, Secure Shell Protocol (SSH) tunnels, system randomization, and proactive secret sharing) that uses Moving Target Defense (MTD) techniques for cloud security. They provided an overview of the MAZE design and described the experiments they performed that demonstrated the potential of an MTD-based cloud storage system to protect against attackers while providing reasonable response time.

3. Performance Improvement of Anomaly Detection on IoT Network

Latha Suryavanshi Karakos and Jumoke Ladeji-Osias, Morgan State University

In addressing the challenges of securing the Internet of Things, this research deals with statistical and machine learning based countermeasures for Botnet attacks on IoT devices using IoTID20, one of the latest botnet datasets available for evaluation. This work focuses on the decision tree algorithm, which is fast, efficient and most suitable for deployment on resource constrained IoT devices. The author described the experiments that were performed and presented the results using datasets of varying sizes.

4. Uighurs and Facial Recognition Technology

Camille Catania, University of Kent

The author posited that Chinese facial surveillance companies are using ethnic clashes in the province of Xinjian to generate profits, and she linked the disappearance of members of the Uighur minority to the use of Automated Facial Recognition Technology (AFRT). By looking at the way AFRT companies develop, use, and market this technology, this research aims to demonstrate how by targeting members of the Uighur minority, these companies are able to perfect their technology through what is effectively a human test trial. She presented the theory behind her research, the results of her literature review, and the rationale behind her hypothesis.

5. Vulnerability Evaluation and Prioritization for Cyber Resilient Systems

Omer Keskin, Nick Gannon, Brian Lopez, and Unal Tatar, University at Albany

The presenter addressed the challenges of vulnerability management to include the fact that it is resource-intensive and, because patching all vulnerabilities is impractical, vulnerabilities need to be prioritized according to their criticality. Mr. Keskin described NIST’s Common Vulnerability Scoring System (CVSS) and then presented the authors’ research method to prioritize the criticality of vulnerabilities by integrating environmental factors in two steps: connecting the network topology by analyzing asset-based dependencies from an adversarial perspective; and analyzing the depended of the organizations business process of vulnerable assets. His presentation included the research methodology and the system architecture. He concluded that their method results in more accurate vulnerability rankings than the CVSS base metrics provided in the National Vulnerability Database.

Special Session on Hard Problems

HotSoS 2021 included a special breakout discussion session centered around what should constitute the Science of Security Hard Problems. SoS Lablet Principal Investigators originally identified five Hard Problems when the SoS program was initiated. SoS community influencers are revisiting the SoS Hard Problems and their definitions in preparation for a second decade of the NSA SoS Program. The Hard Problems session consisted of small discussion groups followed by a joint session with summaries from the discussion group moderators. The topics were as follows:

- **AI Trustworthiness**
- **Adversaries**
- **Human Behavior**
- **Human Weakness**
- **Time**
- **Rethinking Security Measures**
- **Systems**
- **Adoption of Tech**
- **Data Provenance**

Poster Session

Fifteen posters were presented at HoTSoS representing the work of over 20 authors from multiple universities and institutions.

The HotSoS 2021 Best Poster Award “Managing the Security Risk of Open-Source Dependencies: Current Tools and Challenges” was given to Nasif Imtiaz and Laurie Williams from North Carolina State University.

Best Poster

1. Managing the Security Risk of Open-source Dependencies: Current Tools and Challenges

Nasif Imtiaz, Laurie Williams
North Carolina State University

Software Composition Analysis (SCA) tools can detect open-source dependencies and report known vulnerabilities in them. A key strength of an SCA tool is the accuracy, up-to-dateness, and completeness of its vulnerability database. Future research opportunity lies in 1) understanding how developers assess security risk and make fix decisions for dependency vulnerabilities; and 2) automation techniques to ensure continuous monitoring of vulnerability data in the open source ecosystem.

2. A Raspberry Pi Mesh Network to Monitor Biodiversity

Timothy Park, Jacob Scriffiny, Noah Smith, and LTC Thomas Babbitt
United States Military Academy

This is the fourth consecutive HotSoS that USMA personnel have presented a poster on the Raspberry Pi Sensor Network. Military bases contain significant biodiversity including many endangered species and this network allows for an inexpensive and dynamic method to monitor wildlife on military bases.

3. Automatic Security Patching for Containerized Java Applications

Olufogorehan Tunde-Onadele, Xiaohui Gu North Carolina University

The motivation behind this research is the fact that containerized applications pose a set of new security challenges to distributed computing environments. The authors seek to determine whether code analysis can be improved by finding patterns connected to core Java library code and whether static analysis is practical for security patching containerized applications.

4. A Hybrid Approach to Security Attack Detection in Containerized Applications

Yuhang Lin and Xiaohui Gu
North Carolina State University

Previous work on examining security vulnerabilities of containerized applications uses supervised and unsupervised approaches. This work suggests a hybrid approach combining both supervised and unsupervised learning that reduced false positive rate by 52.6%

5. Privacy-Preserving framework for Anomaly Detection Models in Smart Homes

Sai Sree Laya Chukkapalli, Anupam Joshi
University of Maryland, Baltimore County

The unexpected security risks in the IoT systems deployed in smart homes has led to privacy concerns of the smart home device users. The authors collected and integrated behavioral data from IoT devices and applied an anomaly detection model.

6. Cyber KG + RL: Guiding Reinforcement Learning Algorithms for malware mutation and detection with knowledge graphs

Aritran Piplai, Anupam Joshi
University of Maryland, Baltimore County

Cyber Knowledge Graphs (CKG) provide important information to a Deep Neural Network-based Reinforcement Learning (RL) algorithm such that the length of the action sequence can be decreased while retaining the same performance on defeating a malware classifier.

7. Disambiguation of Policy Documents for Improved Access Control Decision Making

Anantaa Kotal, Anupam Joshi
University of Maryland, Baltimore County

Knowledge extraction from policy documents is affected by ambiguity in policy language and the authors developed a method to objectively identify ambiguity in policy document. They propose a model that uses knowledge from this work to disambiguate policy text for improved access control decision making.

8. Capturing and analyzing windows kernel events for anomaly detection

Swapnil Bhosale, Anupam Joshi, Jeff Avery
University of Maryland, Baltimore County

The authors created a kernel resident system to tap into file, registry and network events in Windows OS. In related work, they have developed an unsupervised neural network LSTM based sequence anomaly detection model that uses this data.

9. Generating Security Requirements from Threat Reports

Rayhanur Rahman, Laurie Williams
North Carolina State University

The challenges in addressing this issue include obtaining and labelling the dataset, changes in attack patterns and the lack of benchmarks.

10. From Collaboration Characteristics to Code Quality

Amanul Haque, Nirav Ajmeri, Ruijie Xi, Laurie Williams, Munindar Singh
North Carolina State University

Quantifying collaboration characteristics of an Open-Source Software (OSS) team and investigating its correlation with the quality of the software artifact produced by the team can lead to predictive models for early detection of potentially poor-quality code in software development projects. To understand how OSS team collaboration relates to the quality of the software produced, we applied text and social analytics to a dataset of 168 OSS projects. Our findings highlight the significance of effective communication in team collaboration and identifies some key attributes that can potentially be used in predictive models for early bug identification in software projects.

11. A Dependently Typed Attestation Protocol Language

Anna Fritz and Perry Alexander
University of Kansas

Remote attestation is the act of making trust decisions about a communicating party. During this process, an appraiser asks a target to execute an attestation protocol that generates and returns evidence. The appraiser then evaluates the evidence to make claims about the target. We use the language Copland to formally specify attestation protocols and introduce Copland centered negotiation as prerequisite to attestation. The goal of negotiation is to generate a protocol that meets the target's needs for constrained disclosure and the appraiser's desire for comprehensive information.

12. DNA Data Storage and Cryptography

Anna Vinnedge
United States Military Academy

The typical storage systems that currently hold the worlds data in the form of 0's and 1's will not be able to keep up. For this reason, finding ways to store data efficiently has become an increasingly relevant problem. This poster provides an overview of DNA cryptography, a new area of research in data storage and security.

13. AIRMAIL: Scaling Mobile Vulnerabilities through the AI Supply Chain

Bradley Potteiger and Rachel Cohen
Johns Hopkins Applied Physics Laboratory

AI and ML libraries are being increasingly incorporated within mobile applications. Supply chain distribution channel and homogeneous software structure means that one software bug can lead to an exploit scaling to millions of devices around the world. AI dependencies for mobile applications provide a new attack vector beyond traditional adversary machine learning approaches to covertly obtain and maintain a foothold into adversary systems. Developing an autonomous reverse engineering and exploitation framework will allow designers to more rapidly identify vulnerabilities in critical applications.

14. WIP: Guidelines for Improving Scientific Reporting in Cyber Security

Matthew Armstrong, Jeffrey Carver
University of Alabama

In order to build the Science of Security, cyber security research must be reported in a scientifically rigorous and valid manner. This poster presents the results of interviews with cyber security experts that were structured to further the goal of developing a set of guidelines for reporting scientifically rigorous cyber security research.

15. A Pilot Study on Website Fingerprinting Vulnerability of Tor Onion Services and General Websites

Loc Ho1, Young-Ho Kim, Won-gyum Kim, Donghoon Kim, Doosung Hwang
Arkansas State University, Dankook University, AiDeep

Website fingerprinting attacks have exposed a vulnerability in Tor network. Tor browsers can access both onion services and general websites. The onion (hidden) services are services that can only be accessed over Tor while the general websites (non-hidden services) that can also be accessed with regular web browsers. However, there has not been much research conducted on how secure the onion services are in the website fingerprinting attack compared to the general websites. To find out the vulnerability of Tor onion services and general websites, we implemented a framework that can collect network traffic. The framework can filter the collected traffic to consist only of Tor-related traffic.

The agenda and selected presentations are available [here](#).

For non-members, information about the SoS VO community and the process for requesting membership is available [here](#).

In 2022, the 9th Annual HotSoS will be hosted by the [University of Illinois at Urbana-Champaign](#), with Sayan Mitra as the General Chair.



Outreach

Again in 2021 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. Whether within their own organizations, in the classroom, or at local and international symposia and conferences, SoS community influencers ensured that the science of security and privacy was included in discussions, presentations, and curricula.

The major means of community 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 2000 members in 2021, 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 2021 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.

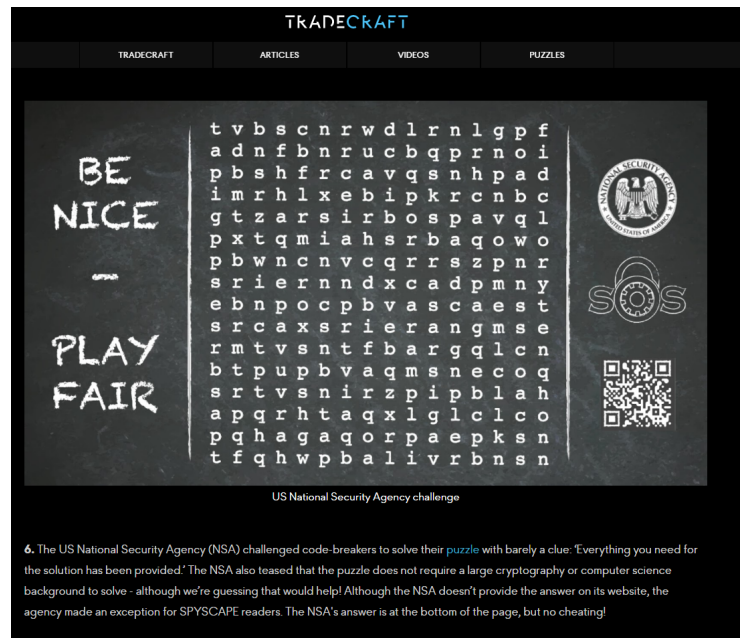
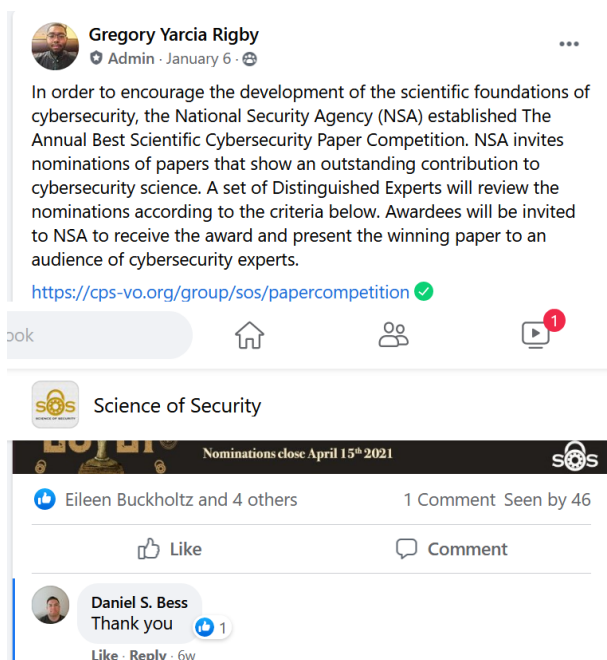
New members are encouraged and can join by signing up via the SoS VO website at www.sos-vo.org

The SoS Reviews and Outreach (SoS R&O) newsletter published 12 editions in 2020 to almost 1800 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 2300 Pub Crawl items published in 2021 covering over 330 topics and representing the curated work of over 7500 authors.
- In the News: A consolidated list of selected articles from recent SoS-VO postings that are focused on SoS-related research, advancements, and discoveries, and are published daily on the SoS-VO. In 2021 approximately 1200 news items were included in the R&O.
- Upcoming Events: Information on SoS-related conferences, symposia, and workshops. There were close to 200 events identified and publicized in 2021.
- 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 and 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

The Science of Security also maintains a Facebook presence, and there were over 300 Facebook postings in 2020.

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.



Cyber Makerspace

While educational robotics and makerspaces are useful to modern STEM education, they introduce both physical and economic barriers to entry. In 2021 Vanderbilt University created a “Cyber Makerspace,” a simulated, networked environment that facilitates instruction on cyber-physical systems, their security and related topics while reducing cost and complexity. The approach, which is led by Principal Investigator Ákos Lédeczi, will facilitate reaching audiences from traditionally underrepresented groups. It also supports remote learning, an especially important feature due to the pandemic.

The virtual robotics environment is connected to a block-based programming language, NetsBlox, to allow students to engage with the curriculum regardless of programming experience. The networked simulation and collaborative programming environment combine to become especially effective for distance learning. Three summer camps for high school students and a week-long Professional Development (PD) workshop for 4 high school teachers were conducted. The first camp specifically targeted students from Vanderbilt’s partner, the MLK STEM Magnet School in Nashville; 11 students participated in the week-long virtual camp. The curriculum was the same as in previous years in-person cybersecurity-focused camps utilizing physical robots, but this time, using the virtual environment. Student feedback was very positive. This was followed by the teacher PD workshop. Finally, two two-week long camps were conducted (17 students total) taught by the teachers who had gone through the PD workshop. The curriculum included the virtual robotics environment but was also Internet of Things-focused, utilizing the PhonoIoT mobile app. The final project used PhonoIoT to make the phone a remote controller for the virtual robots. Again, students remained engaged throughout the camp and provided positive feedback. The teachers were equally excited and are planning on using the material in their regular classrooms next year.

Students will be able to move from the block-based user interface of NetsBlox to a new, Python-based environment with a gentle learning curve: the aim is to preserve many of the block-based features that students are used to: sprites, the stage, events, the concurrency model, etc., but they can implement their scripts in Python. To enable Python in the environment, Vanderbilt has started the development of a Python-based environment that enables students to write Python programs that access the services and features provided by the NetsBlox server. Hence, they will be able to use the physical and virtual robots with all the cybersecurity features as well as PhonoIoT from Python. The program sponsors have also worked on a revised version of the virtual robotics environment; one that will work in the browser and hence, will not require installing a separate executable. Another significant advantage is the seamless integration with the NetsBlox programming environment.

Vanderbilt has launched an outreach program Vanderbilt Digital Nights (VDN) focusing on local high school students. VDN is a series of online workshops held monthly using NetsBlox to introduce advanced computing concepts to young learners. The program has reached about 50 high school students mainly from the middle Tennessee area, but since all workshops were virtual, some of the participants came from out of state from a variety of places. So far, approximately

40% of participants have been female. The VDN series concluded in April with a workshop on climate change. Students created a program that displayed atmospheric CO2 concentrations using Antarctic ice core data going back 800,000 years. The series will continue in 2022 utilizing the virtual robotics environment and the PhonoIoT app as well.

Vanderbilt has prepared an application to the State of Tennessee for a new high school computer science course that will include all the tools and curriculum developed under this project.

The Cyber Makerspace technology and accomplishments have been shared in the following ways:

Community Engagements

- Technology Demonstration, “Beyond CS Principles: Bringing the Frontiers of Computing to K12,” by Ákos Lédeczi, Shuchi Grover, Veronica Catete, and Brian Broll at the 52nd ACM Technical Symposium on Computer Science Education (SIGCSE 21), March 2021
- Teacher PD workshop at the annual CSTA conference for about 30 computer science teachers
- Two presentations at the annual Snap!Shot workshop attended by many high school teachers (<https://www.snapcon.org/conferences/snapshot2021/schedule/events>): PyBlox: A Snappy Python Environment by Devin Jean and Virtual Robots in the Browser by Gordon Stein
- Technology Demonstration, “Using Mobile Devices and Visual Programming to Introduce IoT in the Classroom,” by Devin Jean and Akos Ledeczi, IEEE VL/HCC Conference, October 2021, Best Showcase Award

Publications

- Devin Jean, Gordon Stein, and Ákos Lédeczi, “Demo Abstract: Hands-On IoT Education with Mobile Devices,” *ACM Information Processing in Sensor Networks*, 2021
- Gordon Stein, Devin Jean, and Ákos Lédeczi, “Demo Abstract: Distributed Virtual CPS Environment for K12,” *ACM Information Processing in Sensor Networks*, 2021
- Brian Broll and Devin Jean, “Making CS More Engaging in an Interconnected World,” *CSTA Workshop/Professional Development for K12 Teachers*, 2021
- Devin Jean, Brian Broll, Gordon Stein, and Ákos Lédeczi, “Your Phone as a Sensor: Making IoT Accessible for Novice Programmers,” *Frontiers of Education Conference*, October, 2021
- Brian Broll, Akos Ledeczi, Gordon Stein, Devin Jean, Corey Brady, Shuchi Grover, Veronica Catete, and Tiffany Barnes, “Removing the Walls Around Visual Educational Programming Environments,” *IEEE VL/HCC Conference*, October, 2021
- Gordon Stein and Akos Ledeczi, “Enabling Collaborative Distance Robotics Education for Novice Programmers,” *IEEE VL/HCC Conference*, October, 2021



produced by cyber pack ventures, inc.



sos-vo.org