SCIENCE OF SECURITY AND PRIVACY

2018 Annual Report

Science of Security

and

Privacy (S0S) Initiative

2018

Table of Contents

digital copies are available at https://cps-vo.org/sosannualreport

Executive Summary

The Science of Security and Privacy Initiative (SoS) continued to contribute to the advancement of cybersecurity science in 2018 through its support of research, commitment to scientific principles, and outreach aimed at growing the community. Under the sponsorship of the National Security Agency (NSA) Research Directorate (RD) whose mission is to secure the future by conducting ground-breaking research, the SoS initiative advanced the goal of protecting cyberspace.

The SoS initiative is focused on the establishment of a cybersecurity discipline producing scientifically supported cybersecurity advancements. By replacing ad hoc and common practice approaches to security with scientifically supported methods, SoS is developing strategic rather than tactical methods of approaching cybersecurity. These 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 strategic goals:

- Engage the academic community for Foundational Research
- Promote regiorous scientific principles
- Grow the SoS community

The SoS Initiative engaged the academic community for foundational research in 2018 through sponsorship of the third generation of SoS Lablets of which there are six. The Lablets are 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). These Lablets are focused 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 and Models
- Resilient Architectures
- Human Behavior

Privacy and Cyber-Physical Systems are cross cutting interest areas that span the five Hard Problems. While all the Lablets tackle these challenges, there is a strong focus on privacy at the ICSI Lablet. The third generation of Lablets continues the same goals as the prior two generations. A major new objective is to expand collaborative engagements by taking advantage of the synergy between the scientific and NSA operations perspectives. The SoS initiative began the year with twenty projects and added another six in the fall of 2018. During 2018, the Lablet projects published 34 papers.

The projects, organized by Hard Problems, are as follows:

Scalability and Composability

- Cloud-Assisted IoT Systems Privacy (KU)
- Side-Channel Attack Resistance (KU)
- Scalable Trust Semantics and Infrastructure (KU)

Policy-Governed Secure Collaboration

- Principles of Secure Bootstrapping for the Internet of Things (NCSU)
- Obsidian: A Language for Secure-By-Construction Blockchain Programs (CMU)
- Reasoning about Accidenal and Malicious Misuse via Formal Models of User Expectations and Software Systems (NCSU)
- Analytics for Cyber-Physical System Cybersecurity (VU)

Privacy Emphasis

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

Security Metrics and Models

- Predicting the Difficulty of Compromise through Modeling How Attackers Discover Vulnerabilities (NCSU)
- Securing Safety-Critical Machine Learning Algorithms (CMU)

- Multi-model Test Bed for the Simulation-based Evaluation of Resilience (VU)

Resilient Architectures

- An Automated Synthesis Framework for Network Security and Resilience (UIUC)
- Formal Approaches to the Ontology and Epistemology of Resilience (KU)
- A Monitoring, Fusion, and Response for Cyber Resilience (UIUC)
- Uncertainty in Security Analysis (UIUC)
- Coordinated Machine Learning-Based Vulnerability Discovery and Security Patching for Resilient Virtual Computing Infrastructures (NCSU)
- Model-Based Explanation for Human-in-the-Loop Security (CMU)
- Secure Native Binary Executions (KU)
- Foundations of CPS Resilience (VU)
- Resilient Control of Cyber-Physical Systems with Distributed Learning (UIUC)

Human Behavior

- Mixed Intitiative and Collaborative Learning in Adversarial Environments (VU)
- Characterizing User Behavior and Anticipating its Effects on Computer Security with a Security Behavior Observatory (CMU)

Details about the Hard Problems and 2018 research on the specific projects can be found in Section 1.

While SoS sponsorship of fundamental research also contributes to the achievement of the second goal of promoting rigorous scientific principles, there are several other activities undertaken by SoS that reinforce that effort. In 2018 the SoS initiative sponsored the 6th Annual Best Scientific Cybersecurity Paper Competition, sponsored awards at the Intel International Science and Engineering Fair (ISEF), and for the first time, sponsored two Best Paper Awards, one at the Hot Topics in the Science of Security: Symposium and Bootcamp (HotSoS), and the other at the Symposium on Usable Privacy and Security (SOUPS). There were 28 papers submitted in the 6th Annual Paper Competition bringing the total submissions to over 225 during the competition's six years. This year's winning paper was "How Shall We Play a Game? A Game-theoretical Model for Cyber-warfare Games" by researchers from Carnegie Mellon University (CMU) and University of California, Santa Barbara. This was the fourth consecutive year that SoS has sponsored prizes at ISEF, and this year there were seven winners. The winners of the HotSoS and SOUPS Best Paper competitions were selected by their respective program committees based on evaluation criteria that SoS personnel developed with them, and the winning papers were automatically submitted to the 7th Annual Best Scientific Cybersecurity Paper Competition.

Details on SoS activities to promote rigorous scientific principles in 2018 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. The 2018 Hot Topics in the Science of Security: Symposium and Bootcamp (HoTSoS) brought together over 120 researchers and practitioners from academia, government, and industry for thought-provoking presentations on the cybersecurity science. The SoS Virtual Organizatoin (www.sos-vo.org) grew to over 1500 members and continued as the centralized, online location for researchers and all interested parties to engage in discussion and access the most current research in cybersecurity. The Science of Security, both the NSA SoS-sponsored activities and the topic of Science of Security, regularly appeared in social media.

Details on what the SoS initiative did in 2018 to grow the Science of Security can be found in Section 3.

Privacy

In 2018 with the third generation of Lablets, SoS sponsored its first privacy-focused Lablet at the International Computer Science Institute (ICSI) in Berkeley, California. ICSI is an independent, non-profit research organization with an affiliation agreement with the University of California, Berkeley. ICSI's research focus is aimed at understanding the implications of data use for privacy. NSA privacy research includes prototyping in three areas: Chat Bot (Topic Modeling); the use of secure multi-party communication techniques to help prevent more intrusive measures; and a Common ML Studio. There are four privacy research areas that NSA believes need more attention: the impact of General Data Production Regulation (GDPR) on data sets used for training models; impact of GDPR on models that remember training data; whether ML can help disrupt online tracking; and whether digital canaries can be created to discover biasing attempts within social media or news.

Details on 2018 Privacy activities can be found in the following section.

In 2019 the SoS initiative will continue to sponsor foundational research at the Lablets and seek to increase the impact of Lablet research on cybersecurity operations at NSA. The initiative will sponsor the 7th Annual Best Scientific Cybersecurity Paper Competition, the Intel ISEF awards, and Distinguished Paper competitions at cybersecurity conferences. HotSoS 2019 will be held at Vanderbilt University in April and continue its focus on the advancement of scientific methods in approaching the Hard Problems in cybersecurity. SoS personnel will continue to grow the Science of Security through outreach efforts at all academic levels to raise awareness of the need for foundational cybersecurity science to ensure a mature and reliable cyberdefense.

Privacy

The mission of the NSA Research Directorate is to conduct world-class scientific research with the objective of developing new and innovative techniques (and applications) to support and enable, amongst others, the Information Assurance mission. To that end, the Research Directorate is directed to collaborate with industry and academia to address current research needs to anticipate technological advances that may be disruptive to the mission. In that respect, 2018 has been a remarkable year for the privacy research effort. However, any initial discussion must address the shiny tech coin that has engrossed the tech community: Machine Learning.

With Machine Learning, a machine learns to do better in the future based on what it has experienced in the past. Previously, engineers used code to program all the activities that they wanted a computer to complete. With Machine Learning, a computer is trained with large amounts of data to figure out how to do the task by itself; questions remain as to how much data is needed and what constitutes the right data (free of bias or undue influence). The Machine Learning paradigm can then be viewed as "programming by example." Although we often have a specific task in mind, we don't program the computer to solve the task directly in Machine Learning. Instead, the goal is to seek methods by which the computer will come up with its own "program" based on examples that we provide. The goal that is sought, and continues to be explored, is the development of models that are similar to the way the human brain thinks so that the computer can recognize complex patterns in the data and react accordingly.

What the privacy team has focused on in the latter end of 2018, and moving into 2019, is the development of learning algorithms that advance state of the art in reinforcement learning (e.g., Random Network Distillation) so that we can begin to develop Natural Language Processing (NLP) agents that interact with human users as humans would. This, in turn, builds a level of comfort with the user to trust the agent, from accepting its alerts/warnings to helpful analysis identification, and thus reduce potential privacy intrusions.

In 2018 the Science of Security and Privacy initiative (SoS) sponsored its first privacy-focused Lablet at the International Computer Science Institute (ICSI) in Berkeley California. The overarching goal of the Lablet, in 2018, was to begin to facilitate conducting and disseminating fundamental scientific research on privacy to understand the implications of data use better. In particular, they are concerned with systematically exploring several deeply-connected issues to address six privacy challenges:

- Defining privacy across varying contexts and conceptions, so that researchers and practitioners who approach privacy from varying disciplines can describe their work using a common lexicon;
- Providing transparency in data collection and usage, so that researchers and practitioners can better convey issues stemming from data usage and collection to stakeholders;
- Understanding privacy perceptions that surround the usage of personal data across varying contexts, so that decision-support systems and frameworks can account for human behavior;
- Assessing privacy risks using formal reasoning to account for data usage across varying backgrounds so that researchers and practitioners can utilize mathematical models to predict future privacy risks that are introduced by the composition and aggregation of collected data from varying heterogeneous sources;
- Designing and validating new methods for Big Data accountability that provide hard guarantees and are context-aware; and
- Exploring how current advances in privacy engineering can be applied to solve privacy as mentioned above challenges.

The remaining Lablets, in particular, University of Kansas (KU), Carnegie Mellon University (CMU) and North Carolina State Univiersity (NCSU) have begun to identify and transition previous efforts to conform, or instead satisfy privacy-related NLP tasks. For example, NCSU Professor William Enck's work noted in the upcoming paper "PolicyLint: Investigating Internal Privacy Contradictions on Google Play" has been selected for internal usage in identifying conflicting statements amongst the corpus of internal and external policies governing NSA's mission. KU Professor Fengjun Li's work, noted in paper "Privacy-preserving Classifications for IoT Applications" has been handed off to internal IoT research teams for additional follow-up. Finally, CMU Professor Nicolas Christin's work, "SBO Privacy Research" is showing promise in the detection of Personally Identifiable Information (PII) leakage during active online activity of a user. Expect prototype code representative of these works are to be available by the end of Fall 2019.

Academia is not alone in this privacy research growth. In 2018 NSA's Privacy Research Team tripled its team of researchers focused on privacy-related technologies, acquired a budget, and documented its strategy moving forward--all steps which signify leadership backing and support for this effort. With this growth externally and internally, the NSA Privacy Research team is channeling 2018's accomplishments and inroads towards the following 2019 goals:

- What is the impact of General Data Protection Regulation (GDPR) on datasets used as training for a particular model?
- What is the effect of GDPR on models which have 'memory' and can be de-aggregated?
- Can Machine Learning be used to disrupt online track-

ing for a given user?
- Can we create digital canaries that can detect bias or adversarial influence attacks on social media/news for a given geographical region?

In closing, 2018's privacy efforts can be succinctly wrapped up into two concepts: operational application and internal/external focus. As noted by the current Lablet efforts, papers are not enough – research ideas need to come with some form of necessary code representation to help engineers build upon their success. Finally, privacy enhanced technology doesn't have to be mission-centric. We can leverage the research ICSI completed in 2018 towards their Children's Online Privacy Protection Act (COPPA) effort to include it in applications that benefit society writ large – thus enhancing the privacy understanding and ecology of all.

Science of Privacy

Secure System Architecture and Analysis

Adversarial Machine Learning

Machine learning (ML) is proposed as a scalable defensive and offensive capability in cyber security. Individual proposals range from semi-automated decision support tools to fully-automated capabilities. However, ML models can be exploited—by *adversarial* machine learning—in at least four ways: (a) attackers can *poison* training data used for training ML algorithms to degrade prediction quality, or redirect predictions altogether; (b) attackers can *evade* by manipulating runtime data to ensure ML models misclassify malicious behavior as benign; (c) attackers can *infer* records in the training data; and (d) attackers can approximately reconstruct the ML model for further analysis and exploitation. When ML models of varying qualities are integrated into an ensemble, the attacker can exploit weaknesses in individual models to coordinate a malicious effect in the overall system.

Project outline:

- Understand Attacks and Vulnerabilities
- Develop Mitigations
- Design, Test, and Evaluate
- Modeling the Adversary and ML Decision Support
- Devise Influence Metrics
- Measure Adversarial Influence

Adversaries can Use ML to fool ML

Gu, Dolan-Gavitt, and Garg, "BadNets: Identifying Vulnerabilities in the Machine Learning Model Supply Chain", http://arxiv.org/abs/1708.06733.

About: The Science of Privacy is an innovative approach to the advancement of privacy protections built on mathematical foundations.

General Focus Areas

Abstract: The Information Assurance Research Group (IAR) is attempting to advance privacy research through four pathways: (1) as a *clearinghouse*; (2) as an *incubator*; (3) as *outreach*; and finally, (4) as a purveyor of *infrastructure*.

- As a *clearinghouse*, the program surveys state-of-the-art research and technology for transition opportunities to mission elements.
- As an *incubator*, the program invests time in developing prototypes to evaluate concepts discovered by government, industry and academic partners in order to empirically measure technical impacts on privacy risk. For example, we are developing data-driven techniques to automate legal, privacy, and compliance decision-making across machine learning, artificial intelligence, and formal methods.
- As a point of *outreach*, the program provides access to unclassified partnerships at other governmental, industry and academic institutions; this includes developing necessary relationships to recruit talent to support privacy-relevant mission areas.
- Finally, the program acts as a purveyor of *infrastructure* to support the design and development of privacy-enabling tools and datasets. Through this application of both theoretical and practical methodologies we intend to enhance privacy-oriented architecture(s) for mission. For example, we are driving research into practical cloud-based infrastructure services based on data modeling, legal automation, and data-driven privacy risk evaluations to conduct at-scale, data-driven operations on Big Data. Areas of study include: supporting Privacy as a Service (PRVaaS); automated analytic vetting for privacy data-leakage in support of Development Operations (DevOps); and the enabling of real-time privacy impact assessments of analytics and analytic workflows.

Constructively **Mitigating Adversarial Machine Learning (MAML)**

Pattern Recognition and Applications Lab, http://pralab.diee.unica.it/en/SecurityEvaluation

Knowledge Management

A project focusing on ensemble classification techniques, ranging from entity extraction to more advanced question-answering suites to:

- Engineer a framework in which legal and policy decisions can be encoded.
- Automatically synthesize business rules from new legal and policy documents.
- Identify the policies and laws from which a business rule was derived. Can NLP techniques advance our compliance with privacy policies/laws?
- Identify the policies and/or laws that were used in generating a response to a general compliance question.

Deep Learning

Deep learning is a subfield of machine learning where models inspired by how the human brain works are expressed mathematically, and the parameters defining those mathematical models—which can number in the few thousands to 100+ million—are learned automatically from the data. In this project we aim to use deep learning to predict the types of questions people are likely to ask about our corpus of legal documents, so we are better able to conduct pre-query analysis. As illustrated below, we are leveraging Unity's multi-agent AI concept as a potential visualization and controller for our MAML effort and data fusion projects.

Unity AI architecture being used for our Compliance Assistant Research

PII Data Leakage

A project to further our understanding of how various applications routinely leak private information without the user's, or possibly developers', knowledge—the leakage occurring through use of third-party SDKs. We will attempt to advance a user's autonomy through our existing mobile-monitoring techniques. The project uses observations of SDK and third-party advertising techniques as a means to mitigate the leaks. This includes working with Android simulators, iOS simulators, and dynamic program analysis to observe run-time leaks.

privacy.foundations.research@nsa.gov

Engaging the Academic Community for Foundational Research

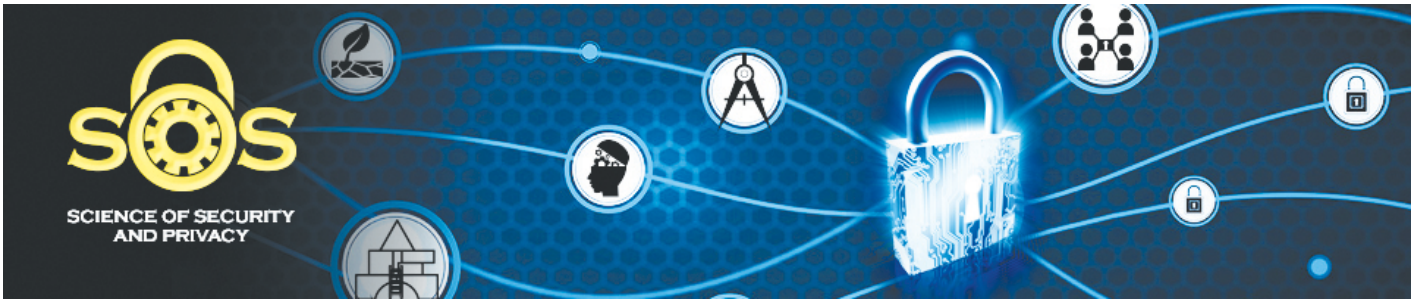

In 2018 the Science of Security and Privacy (SoS) initiative engaged the academic community for foundational research by awarding contracts to six Science of Security Lablets. The Lablets were competitively selected in 2017 from a Broad Agency Announcement (BAA) released by NSA. The specific research projects were selected by NSA to create a portfolio of projects that have technical excellence, NSA mission relevance, broad applicability beyond NSA, and in total would span the SoS five Hard Problems. The six Lablets and their thirteen Sub-Lablets performed research on 26 projects in 2018 and published 34 peer-reviewed papers bringing to almost 600 the number of papers published by Lablet researchers since the SoS initiative was established in 2012. The papers have addressed multiple aspects of the five Hard Problems, and have been presented at conferences, symposia, and workshops around the world. The foundational research embodied by the papers has contributed significantly to enhancing the scientific rigor of research into cybersecurity. The Principal Investigators (PIs) of the Science of Security Lablets, along with the NSA Research organization, developed five Hard Problems. These Hard Problems serve as a means of establishing challenging and critical research goals, 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. The Lablets also engage regularly in community outreach, participate in international conferences and workshops, and integrate Science of Security principles into their curricula. In addition to providing quarterly and annual reports on their activities, the six Lablets, along with Sub-lablets, collaborators, and NSA researchers, meet quarterly to present updates on their research projects and exchange information about issues related to Science of Security. **The Hard Problems, Lablet activity, and Lablet quarterly meetings are detailed in this section.**

Hard Problems

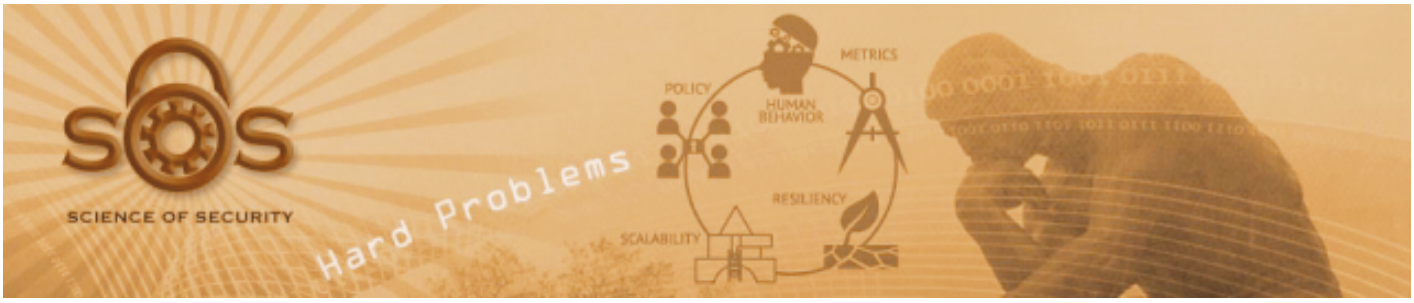

The Principal Investigators (PIs) of the Science of Security Lablets, in collaboration with NSA Research, developed the Hard Problems as a means of establishing challenging and critical research goals. The Hard Problems also serve as the beginnings of a common language and a way to assess progress. These problems were selected for their level of technical challenge, their potential operational significance, and the likelihood that these problems would benefit from emphasis on scientific research methods and improved measurement capabilities. The five Hard Problems are not intended to cover all cybersecurity research challenges, but rather five specific areas that need scientific progress. Fundamental research undertaken by the Lablets is tied to at least one Hard Problem. A discussion on the history and future of the five Hard Problems took place at the Fall Quarterly Lablet meeting held at Carnegie Mellon University. **See section 1, Lablet Quarterly Meetings, for details on the Hard Problem discussion.**

The five problems are addressed 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 of human actors in cybersecurity. The Hard Problem focuses on the behaviors and actions of malicious attackers, system users, and software/system developers. 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.

Carnegie Mellon University

The Carnegie Mellon University (CMU) Science of Security and Privacy (SoS) Lablet is currently focusing on projects in four Hard Problems areas: Human Behavior, Secure Collaboration, Security Metrics and Models, and Resilient Architectures.

The first project is Securing Safety-Critical Machine Learning Algorithms and deals with metrics and models to improve robustness in machine learning algorithms. This includes understanding both how classifiers can be spoofed, including in ways that are not apparent to human observers, and also how robustness of classifiers can be enhanced, including through explanations of model behavior, and also means to harden models against attacks.

The second project, Model-Based Explanation for Human-in-the-Loop Security, focuses on combining human and automated actions in response to security attacks. 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.

CMU's third project, Obsidian: A Language for Secure-By-Construction Blockchain Programs addresses models for secure collaboration and models for contracts. These models are applicable 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. This project, 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.

CMU's fourth project, Characterizing User Behavior and Anticipating its Effects on Computer Security with a Security Behavior Observatory continues the development of an observing infrastructure to better asses how users in home setting actually behave when faced with security threats.

Details on these projects can be found in Fundamental Research, below, and in the Fall Lablet Quarterly section.

The CMU Sub-Lablets are Chatham University, George Mason University, University of North Carolina at Chapel Hill, University of California, Berkeley, and Indiana University.

The CMU Lablet is led by Principal Investigators (PIs) Bill Scherlis and Jonathan Aldrich.

PI William Scherlis

PI Jonathan Aldrich

FUNDAMENTAL RESEARCH

Project: Securing Safety-Critical Machine Learning Algorithms
Lead PI: Lujo Bauer
Co-PI: Matt Fredrikson
Participating Sub-Lablet:
University of North Carolina at Chapel Hill
Hard Problems: Security Metrics and Models, Resilient Architectures

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 with 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 (say, for surveillance and authentication) 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. Armed with this understanding of the weaknesses of popular classifiers and their modes of use, we will develop explanations of model behavior to help identify the presence of a likely attack; and generalize these explanations to harden models against future attacks.

Improving explanations for neural networks: Significant progress was made implementing a framework for explaining the predictions made by deep neural networks, and incorporating it into a graphical tool for use by researchers and practitioners. Explanations may identify the network-internal factors that cause misclassifications, and leverage this capability to make progress on the hard problems above. We also believe that certain types of explanations can comprise a runtime defense, with a human in the loop, by exposing cases where predictions appear to be "made for the wrong reasons." Our approach to explanations allows analysts to parameterize queries of network behavior on the aspect being explained, the set of samples in question, and the portion of the network under study, and our tool gains flexibility by exposing these as options. While this tool is useful for our activities on the project, we plan to release it as an open-source project, as well as a more limited interactive web application, for other researchers as well.

We presented a paper at the CV-COPS workshop (The Bright and Dark Sides of Computer Vision: Challenges and Opportunities for Privacy and Security) discussing the suitability of different L_p-norms as metrics for determining whether inputs to machine-learning algorithms are "close" or "far" from each other. An input is held to be "adversarial" when it appears very similar to (or indistinguishable from) a benign input, but is classified differently. In the literature, similarity is typically measured using L_p-norms: two inputs are considered to be similar if their distance according to an L_p-norm is smaller than some application-specific threshold. This paper shows that such a notion of similarity is neither necessary nor sufficient for at least several datasets commonly used in research on adversarial machine learning. The paper's goal in pointing this out is to steer research away from using seemingly fragile metrics for defining adversarial inputs and what it means for a system to be robust to them.

Project: Model-Based Explanation for Human-in-the-Loop Security
Lead PI: David Garlan
Hard Problems: Security Metrics and Models, Resilient Architectures

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. 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).

Building on prior work that we have carried out in this area, 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. Critically, these models capture an explicit value system (in a multi-dimensional utility space) that forms the basis for determining courses of action. Because the value system is explicit, we believe that it will be possible to provide a rational explanation of the principles that led to a given system plan. Moreover, our approach will allow the user to make corrective actions to that value system (and hence, fu-

ture decisions) when it is misaligned. This will be done without a user needing to know the mathematical form of the revised utility reward function.

Resiliency with observability: The adversarial nature of the security domain, and APTs in particular, poses unresolved challenges to the design of self-adaptive systems, such as defending against multiple types of attackers with different goals and capabilities. In this interaction, the observability of each side is an important and under-investigated issue in the self-* domain. We have proposed a model of APT defense that elevates observability as a first-class concern. We evaluate this model by showing how an informed approach that uses observability improves the defender's utility compared to a uniform random strategy, as well as demonstrate how the approach can enable robust planning through sensitivity analysis, can inform observability related architectural design decisions, and can scale to realistically long time horizons. This work builds on techniques to learn strategies for online games. To support experimentation, we developed parsers and feature extractors to pull out the interesting information from the state and make it feasible to use Inverse Reinforcement Learning (IRL) to learn the action policies. We were then able to use, test, and analyze a Python IRL algorithm to demonstrate that it could indeed learn different strategies for different players. Several explanation algorithms were implemented on top of the policies/strategies to be able to summarize the differences in strategies automatically.

Project: Obsidian: A Language for Secure-By-Construction Blockchain Programs
Lead PI: Jonathan Aldrich
Co-PI: Brad Myers
Participating Sub-Lablet: George Mason University
Hard Problems: Scalability and Composability, Policy-Governed Secure Collaboration, Resilient Architectures, Human Behavior

This project considers models for secure collaboration and contracts in a decentralized environment among parties that have not established trust. A significant example of this is blockchain programming, with platforms such as Ethereum and HyperLedger. There are many documented defects in secure collaboration mechanisms, and some have been exploited to steal money. Our approach builds two kinds of models to address these defects: typestate models to mitigate re-entrancy-related vulnerabilities, and linear types to model and statically detect an important class of errors involving money and other transferrable resources.

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.

Safer blockchain transactions: Hackers have exploited security vulnerabilities in existing blockchain programs. To address this, we are designing a new language, Obsidian, using principles of user-centered design. Obsidian uses the technical approaches of typestate (expressing both the types of objects and their state in a way that supports static reasoning) and linearity (to avoid loss or duplication of tracked assets). These are intended to support stronger safety guarantees than current approaches for programming blockchain systems.

PUBLICATIONS

Project: Securing Safety-Critical Machine Learning Algorithms

- Mahmood Sharif, Lujo Bauer, and Michael K. Reiter, "On the suitability of L_p-norms for creating and preventing adversarial examples," in *Proceedings of The Bright and Dark Sides of Computer Vision: Challenges and Opportunities for Privacy and Security (in conjunction with the 2018 IEEE Conference on Computer Vision and Pattern Recognition)*, June 2018.

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

- Javier Camara, Wenxin Peng, David Garlan, and Bradley Schmerl, "Reasoning about Sensing Uncertainty and its Reduction in Decision-Making for Self-Adaptation," in *Science of Computer Programming,* 2018. Accepted for publication.
- Roykrong Sukkerd, Reid Simmons and David Garlan, "Towards Explainable Multi-Objective Probabilistic Planning," in *Proceedings of the 4th International Workshop on Software Engineering for Smart Cyber-Physical Systems (SEsCPS\'18)*, Gothenburg, Sweden, 27 May2018.

Project: Obsidian: A Language for Secure-By-Construction Blockchain Programs

- Michael Coblenz, Jonathan Aldrich, Brad Myers, and Joshua Sunshine, "Interdisciplinary Programming Language Design (Essay)" in *Proceedings of Onward!*, 2018. To appear.

EDUCATIONAL

The top-rated undergraduate program in computer science at Carnegie Mellon now has a Concentration in Security and Privacy. This concentration was developed over a period of several years, was recently approved, and is now accepting its initial cohort of students. Details regarding the program are on the web (http://isri.cmu.edu/education/undergrad/secpriv/index.html).

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

- With non-Lablet sponsorship, the CMU ISR in the School of Computer Science supported a Research Experience for Undergraduates with a focus on Software Engineering (REUSE) in Summer 2018. The REUSE program targets US undergraduate students, and includes several students with an interest in software security.
- Three undergraduates also supported this project.

COMMUNITY OUTREACH

Project: Securing Safety-Critical Machine Learning Algorithms

- Lujo Bauer presented work on adversarial machine learning at the Computational Cybersecurity in a Compromised Environment (C3E) workshop in Atlanta in September.
- Matt Fredrikson began co-organizing a workshop on Security in Machine Learning, to be held in conjunction with the 2018 Neural Information Processing Systems conference in December, a premier conference in Machine Learning.

Project: Obsidian: A Language for Secure-By-Construction Blockchain Programs

- Josh Sunshine, who co-leads the Obsidian project, also co-directs the CMU ISR Research Experience for Undergraduates (REU) program. Two undergraduate students participated in Obsidian research.
- In many developing countries, farming insurance markets have not developed. A severe weather event like a very late hard frost or a drought can therefore devastate these farmers. We are working with the World Bank to develop a parametric insurance platform on the Blockchain with Obsidian to address this need. The platform will serve as an evaluative case study of the expressiveness and effectiveness of the Obsidian language design.

Carnegie Mellon University

International Computer Science Institute

New Science of Security Lablet at ICSI

The International Computer Science Institute (ICSI) in Berkeley, CA is the home to one of six new NSA-funded lablets focused on security and privacy research over the next five years. ICSI's lablet is led by Dr. Serge Egelman, head of the Usable Security and Privacy Research group at ICSI, and includes collaborators at Cornell Tech and UC Berkeley.

Dr. Egelman said, "We're really excited that the NSA has committed to funding foundational research into online privacy. Under the auspices of the lablet, my group is performing research to better understand new threats to online privacy, people's privacy preferences and decision-making, as well as how to design usable and effective privacy controls for emergent technologies, such as IoT devices."

The International Computer Science Institute (ICSI) Science of Security and Privacy (SoS) 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. "The implications of data use," is concerned with systematically exploring several deeply-connected issues to address six privacy challenges:

- **Defining privacy** across varying contexts and conceptions, so that researchers and practitioners who approach privacy from varying disciplines can describe their work using a common lexicon;
- **Providing transparency** into data collection and usage, so that researchers and practitioners can better convey issues stemming from data usage and collection to stakeholders (e.g., data subjects, policymakers, and the general public);
- **Understanding privacy perceptions** that surround the usage of personal data across varying contexts, so that decision-support systems and frameworks can account for human behavior (e.g., concerns, preferences, expectations, and potential reactions);
- **Assessing privacy risks** using formal reasoning to account for data usage across varying contexts, so that researchers and practitioners can utilize mathematical models to predict future privacy risks that are introduced by the composition and aggregation of collected data from varying heterogeneous sources;
- Designing and validating new methods for **Big Data accountability** that provide hard guarantees and are context-aware; and

- Exploring how current advances in **privacy engineering** can be applied to solve the aforementioned privacy challenges.

The Lablet represents a multi-disciplinary and multi-institutional collaboration to address these six challenges, while framing privacy as a scientific pursuit. The Lablet defines science of privacy as research that is grounded in the three pillars of science: conceptual modeling, formal reasoning about precise models, and empiricism. Using methods from all three pillars, the Lablet researchers intend to rigorously perform foundational research that yields generalizable knowledge into how privacy can be better protected, managed, and reasoned about. Rather than simply engineering new systems, the aim is to formulate and empirically validate new frameworks and methods that can be readily used by others and are generalizable to a myriad of privacy use cases. In short, this work will enable others to build systems grounded in scientific principles that evaluate privacy risks, rather than incremental improvements that are designed using ad hoc methods.

ICSI initiated five projects, which together aim to make advances in defining privacy, privacy engineering, Big Data accountability, understanding privacy perceptions, and assessing privacy risks. The specific projects are as follows:

- **Designing for Privacy**: Conducting a series of workshops on Privacy by Design to examine, improve, and refine privacy engineering practices and outreach efforts.

- **Governance for Big Data**: Conducting a series of workshops with stakeholders from government, industry, and academia to examine issues pertaining to Big Data governance.

- **Operationalizing Contextual Integrity**: Performing empirical research towards applying the theory of Contextual Integrity (CI) to the design of future privacy controls.

- **Contextual Integrity for Computer Systems**: Using formal methods to improve CI theory by refining it to support a wider range of privacy contexts.

- **Scalable Privacy Analysis**: Developing a framework to compare policy against practice. In studying practice, advancing a privacy testbed to studying data flows within the Android operation systems.

Details on these projects can be found below in Fundamental Research and in the Fall Quarterly meeting section.

We have also presented published work related to these projects at multiple international conferences, and are planning several outreach events over the coming year. Our planned outreach events include international workshops and symposia, as well as meeting with stakeholders. We have begun working with educators at the K-12 level to integrate privacy and security education related materials into new curricula being piloted as part of the NSA's GenCyber efforts. In addition, we are developing a new graduate level course that integrates research results stemming from the Lablet (this course will be offered at U.C. Berkeley in 2019).

The ICSI Sub-Lablets are Cornell Tech and University of California, Berkeley. The ICSI Lablet is led by Principal Investigator (PI) Serge Egelman.

Pi Serge Egelman

FUNDAMENTAL RESEARCH

Project: Designing for Privacy
PI: Deirdre Mulligan
Hard Problems: Policy-Governed Secure Collaboration, Human Behavior

The project focuses on designing for privacy holistically: from "privacy by design" to "privacy with design", i.e., designing with privacy throughout whole life cycle. 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. 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.

Identifying potential participants for proposed workshops is ongoing. The pool includes those who have previously attended Computing Community Consortium Privacy by Design workshop series, as well as other privacy-focused scholars, researchers, and practitioners. These participants work in academic, industry, and civil society organizations, and span disciplinary fields including law, engineering, design, and social science. Planning has also begun for an initial series of workshop topics and activities, including discussions on privacy patterns, the state of current privacy practice, and design thinking exercises.

We have started exploring the problem space. First by mapping the existing literature, especially technical literature, on privacy by design, provable privacy, and values-sensitive engineering. Second, measuring how likely it is that data analysis will reach unacceptable models naturally to understand when frameworks must overcome problems vs. when known workflows will lead

to acceptable outcomes. This understanding helps focus the research on the need. Finally, we are comparing existing toolkits for measuring fairness in classification models that examines how these tools can be built into common data science workflows.

To understand current concepts, tools, and practices used in current programs, and to identify education, research, training that could advance privacy work, we are interviewing industry engineers and designers. The interviews are semi-structured on practices related to privacy, and are qualitative interviews with "privacy professionals" (individuals directly responsible for, or involved with, the privacy function at their firms) across nine different information-intensive companies that have or are building out privacy-specific programs. The qualitative work will inform a survey to develop quantitative measures.

We have piloted an interactive design thinking workshop activity to help think about privacy and design practice together with graduate students in a professional technical degree program, which we will refine to use in our practitioner and researcher workshops.

A paper mapping the relationships between privacy and HCI design practices was workshopped at the 2018 Privacy Law Scholars Conference and a revised version of that paper is slated for publication.

Project: Governance for Big Data
PI: Deirdre Mulligan
Hard Problems: Policy-Governed Secure Collaboration, Human Behavior

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. A new data governance approach focuses on accountability and relates more to accounting and auditing. This project aims to synthesize computer science abstractions with governance goals. 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.

This project has begun examining current approaches. We have undertaken to understand and compare approaches currently pursued by NIST (the Privacy Framework), the NTIA (an open RfC on Developing the Administration's Approach to Consumer Privacy), and the consulting firm Nymity (the Privacy Management Accountability Framework).

The project is considering toolkits. By comparing existing toolkits for measuring fairness in classification models, we'll understand how these tools can be built into common data science workflows.

We have also conducted a preliminary analysis of risks that privacy-enhancing technologies provide privacy in name only, examining the extent to which deployments of differential privacy by industry actors in fact protect the privacy of individuals' data, producing a public position paper on the topic. In addition, we are examining provable privacy protection technologies as representations which reduce the dimensionality of the input data to limit inference capacity.

Project: Operationalizing Contextual Integrity
Lead PI: Serge Egelman
Co-PI: Helen Nissenbaum
Participating Sub-Lablet: Cornell Tech
Hard Problems: Scalability and Composability, Policy-Governed Secure Collaboration, Security Metrics and Models, Human Behavior

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; contextual information norms are modeled using data subjection, 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 work includes studies done on permission requests when a phone was inactive (a training exercise), the use of ML to detect when context has changed from expected data use to unexpected, and then a second experiment done in real-time that confirmed earlier findings. The next steps to determine what parameters are actually important to users are as follows:

- 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)

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. This aspect focuses on the Hard Problem of Scalability and Composability. Another 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, a goal that relates to the Hard Problem of Policy-Governed Secure Collaboration. With respect to the Metrics Hard Problem, 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. Finally, to address the Hard Problem of 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.

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.

We submitted the final version of our paper on Children's Online Privacy Protection Act (COPPA) compliance at scale to the Privacy Enhancing Technologies Symposium (PETS), and then presented the work in June. This was one of only two papers that were accepted for publication without mandatory revisions. This work documents the implementation of our dynamic analysis platform, which allows us to examine the privacy behaviors of Android apps under realistic conditions. As a proof-of-concept, we applied our infrastructure to detecting violations of COPPA, finding that a majority of Android apps in the Google Play Store directed at children appear to be violating federal law. In May, we presented a version of this work for feedback at the Privacy Law Scholars Conference (PLSC). We are in the process of adapting and improving this infrastructure to support future research activities, such as:

- Examining compliance with other privacy regulations, including the General Data Protection Regulation (GDPR) in the EU, as well as various state laws (e.g., CalOPPA in California).
- Using natural language processing to compare observed app privacy behaviors with stated practices in privacy policies.
- Creating an API to allow others to use our infrastructure, including existing regulatory agencies, who we are already collaborating with.

We are currently planning future studies in the domain of in-home IoT devices, to explore users' current privacy needs, the capabilities of current devices, and the design of future privacy controls. 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. Toward this end, we have designed several studies surrounding both current commercially-available in-home personal assistants (e.g., Google Home, Amazon Echo, etc.) and prototypes of future devices that we expect to see. As part of this, we submitted three papers to the Symposium on Applications of Contextual Integrity, hosted by Co-PI Helen Nissenbaum in September at Princeton. In one study, we're examining contextual norms around in-home audio monitoring, which is likely to proliferate. Following IRB approval, we're performing a study that involves users of either the Google Home or Amazon Echo answering questions about previously-recorded audio from their devices. These studies are designed to examine users' expectations surrounding data capture and transmission. Both manufacturers make audio recordings accessible to device owners through a web portal, and so our study involves using a browser extension to randomly present these clips to users, and then have them answer questions about the circumstances surrounding the recordings. We're interested in whether they were aware that the recordings were made, how sensitive the content was, as well as participants' preferences for various data retention and sharing policies.

In another set of studies, we're examining existing audio corpora, and then using crowdworkers to identify sensitive conversations, that we can then label and use to train a classifier. The goal is to design devices that can predict when they should not be recording or sharing data.

Finally, in the mobile space, we're looking at disclosure of data-sharing practices in Android app privacy policies. Since GDPR and the soon-to-be-enacted CA privacy law require disclosing data recipients (or categories of data recipients), we want to examine compliance and whether we can detect violations. Using our existing testbed and data, we know the third parties who receive data (ground truth). The question is, are these practices adequately disclosed? To examine this, we've designed a crowdsourcing task to label policies at scale. Using a test corpus of 100 policies, we've found very high inter-rater reliability, and so this method appears to be promising. Once we have labeled policies, we're also going to explore the idea of using this to train a classifier, to determine whether we can automatically extract named entities from policies to compare with our observations of data flows. Through this process, we have identified several security vulnerabilities on Android that we have reported to Google and are working on a paper on covert channels that apps are using to collect user data without consent.

Project: Contextual Integrity for Computer Systems
Lead PI: Michael Tschantz
Co-PI: Helen Nissenbaum
Participating Sub-Lablet: Cornell Tech
Hard Problems: Scalability and Composability, Policy-Governed Secure Collaboration

The overall goal of the research is converting the philosophical theory of contextual integrity into terms computer scientists can use. There is no agreement on what a context is: philosophers and computer scientists have different understandings, with philosophers focusing on abstract spheres of life, and computer scientists focusing on the concrete. The goal is to develop models of context and contextual integrity that meet computer scientists on their own truth. Relevant research questions include accounting for privacy in the design of multi-use computer systems that cut across contexts; modeling the adaptation of contexts to changes in technologies; and determining how contextual integrity relates to differential privacy. The current organizing hypothesis is that contexts are defined by a purpose. The privacy norms of a context promote the purpose, and purpose restrictions are ubiquitous. There are several possible models including game models, Markov decision process models, partially observable Markov decision process models, and multi-agent influence diagrams. Some of the challenges are that contexts don't exist in a vacuum, contexts might be in competition, privacy is multifaceted, and people often disagree. Potential outcomes are progress on defining privacy, further accountability for big data systems that cut across contexts, and enabling policy governed privacy with respect to collaboration.

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

PUBLICATIONS

Project: Operationalizing Contextual Integrity

- Irwin Reyes, Primal Wijesekera, Joel Reardon, Amit Elazari Bar On, Abbas Razaghpanah, Narseo Vallina-Rodriguez, and Serge Egelman, "Won't Somebody Think of the Children?" Examining COPPA Compliance at Scale. *Proceedings on Privacy Enhancing Technologies (PoPETS)*, 2018(3):63-83.
- Primal Wijesekera, Joel Reardon, Irwin Reyes, Lynn Tsai, Jung-Wei Chen, Nathan Good, David Wagner, Konstantin Beznosov, and Serge Egelman, "Contextual Permission Models for Better Privacy Protection," *Symposium on Applications of Contextual Integrity*, 2018.
- Julia Bernd, Serge Egelman, Maritza Johnson, Nathan Malkin, Franziska Roesner, Madiha Tabassum, and Primal Wijesekera, "Studying User Expectations about Data Collection and Use by In-Home Smart Devices," *Symposium on Applications of Contextual Integrity*, 2018.
- Nathan Malkin, Primal Wijesekera, Serge Egelman, and David Wagner, "Use Case: Passively Listening Personal Assistants," *Symposium on Applications of Contextual Integrity*, 2018.

EDUCATIONAL

Project: Operationalizing Contextual Integrity

- We have been developing a privacy and security curriculum suitable for K-12 classrooms, and expect to pilot it in 2019.

COMMUNITY OUTREACH

Project: Designing for Privacy

- We've had several lengthy meetings with teams of privacy practitioners and privacy researchers (some together some separate) at three large companies to understand the tools and approaches they are taking to design privacy into systems and processes, as well as research they are pursuing to address pain points.
- In November, we participated in a privacy workshop at the ACM Computer Supported Cooperative Work (CSCW) conference, presenting a position paper on ways to connect privacy theory to privacy design practice.

Project: Governance for Big Data

- We are planning workshops for stakeholders, which we expect to conduct over the coming year.
- We attended a NIST workshop on their privacy framework, as part of their effort to meet with stakeholders to develop guidance on "governance for big data.

Project: Operationalizing Contextual Integrity

- We presented our research at the Privacy Enhancing Technologies (PETS) symposium this month, and had three papers accepted for presentation in at the Symposium on Contextual Integrity in September. We also presented results at the Privacy Law Scholars Conference (PLSC) in May.
- We reported several security vulnerabilities to Google based on our mobile app analysis findings. Google is providing us with a bounty for one of the vulnerabilities. This vulnerability is actively being exploited by multiple ad SDKs, which we've reported to the FTC, and expect to follow up with them.
- Serge Egelman has had several meetings with regulators to discuss findings on mobile privacy. This includes ongoing consultations with FTC staff, guidance on a lawsuit being brought against platforms and app developers by the NM AG, as well as speaking to international regulators earlier this month at an event organized by the OECD.
- Our PETS paper has generated interest from regulators. We reported several apps that appeared to be violating COPPA to Google, which decided to ignore our report. As a result, they're now being sued by a state AG (alongside the app developers), and have backpeddled and claim to now be taking action based on our reports. So far, this has resulted in the removal of hundreds of child-directed apps from the Play Store.

Project: Contextual Integrity for Computer Systems

- Helen Nissenbaum hosted a symposium on Contextual Integrity at Princeton.

North Carolina State University

June 8, 2018 | Brent Lancaster

North Carolina State University has again been awarded a Science of Security Lablet by the National Security Agency (NSA) to continue its work in developing the cybersecurity and privacy breakthroughs needed to safe guard cyberspace.

The Science of Security Lablet at NC State was established in 2012. NSA this spring announced that NC State would again host a Lablet for an additional five years under a new contract.

Science of Security Lablets are small multi-disciplinary labs at leading U.S. research institutions that are part of NSA's Science of Security and Privacy (SoS) Initiative. Launched in 2012, SoS promotes security and privacy science as a recognized field of research and encourages rigorous research methodologies.

The North Carolina State University (NCSU) Science of Security and Privacy (SoS) Lablet has embraced and helped build a foundation for NSA's vision of a Science of Security and an SoS community. We have emphasized data-driven discovery and analytics to formulate, validate, evolve, and solidify the theory and practice of security. Our research has yielded significant findings, thus providing a deeper understanding of users' susceptibility to deception, developers' adoption of security tools, and how trust between people relates to their commitments. Motivated by NSA's overarching vision for SoS and building on our experience and accomplishments, we are both developing a science-based foundation for the five Hard Problems that we previously helped formulate, and fostering an SoS community with high standards for reproducible research. Our approach involves a comprehensive, rigorous perspective on SoS, focused on an integrated treatment of technical artifacts, humans (both stakeholders and adversaries), and relationships and processes relevant to the Hard Problems.

The SoS Lablet at NCSU has been contributing to conceiving and advancing the Science of Security (SoS) since 2011. We have advanced a variety of focus areas in SoS, including research, scientific methods, and community development. First, we have performed scientific research on the five Hard Problems, which we helped establish. What has unified our research is a basis in practical challenges in cybersecurity and evaluations in terms of their potential or actual impact on practice. Second, we have investigated and developed research methods customized for specific challenges, including theoretical models based on methodologies ranging from mathematics to grounded theory, along with best practices and research guidelines. Third, we have pursued community development by engaging with stakeholders in academia, government, and industry. Such engagement informs our research and enhances our ability to influence practice. Fourth, our Lablet incorporates comprehensive evaluation by specialists in the Science of Science, focused on assessing the

extent to which we have been successful in developing foundational knowledge for SoS, applying rigorous scientific methods, and developing a SoS community of practice.

The NCSU Lablet performed research under five projects in 2018. The first project, Coordinated Machine Learning-Based Vulnerability Discovery and Security Patching for Resilient Virtual Computing Infrastructure has begun applying unsupervised machine learning to container vulnerability discovery. The second, Predicting the Difficulty of Compromise through Modeling how Attackers Discover Vulnerabilities, has published a survey of research on attack surfaces that helps understand how attack surfaces are discussed in the literature. Principles of Secure Bootstrapping for IoT identified a vulnerability in the paging protocol used by intermittently connected devices (such as on the 4G and 5G cellular networks) by which an attacker could identify whether a device is present in a physical region. Reasoning about Accidental and Malicious Misuse via Formal Methods of User Expectations and Software Systems is designed to aid security analysts in identifying and protecting against accidental and malicious actions through automated reasoning on unified representations and software implementation. Development of Methodology Guidelines for Security Research builds upon earlier work to aid the security research community in conducting and reporting methodologically sound science through implementation of community guidelines. Details on these projects can be found below in Fundamental Research and in the Fall Quarterly meeting.

The NCSU Lablet includes researchers from Purdue University, Rochester Institute of Technology, University of Alabama, University of North Carolina at Chapel Hill, and University of North Carolina at Charlotte. The NCSU Lablet is led by Principal Investigators Laurie Williams and Munindar Singh.

PI Laurie Williams

PI Munindar Singh

FUNDAMENTAL RESEARCH

Project: Coordinated Machine Learning-Based Vulnerability Discovery and Security Patching for Resilient Virtual Computing Infrastructure
PI: Helen Gu
Hard Problem: Resilient Architectures

Our research aims at aiding administrators of virtualized computing infrastructures in making services more resilient to security attacks 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. The existing approach is static security analysis and scheduled patching. In the researchers' experiments, this approach fails to detect 90% of vulnerabilities, displays high false alarms, and shows memory inflation caused by unnecessary security patching. This proposal 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.

Containers have become increasingly popular for deploying applications in cloud computing infrastructures. However, our previous study has shown that containers are prone to various security attacks. We conducted an empirical study on the effectiveness of various container vulnerability detection schemes to understand the challenges in real world container vulnerability discovery. In order to understand the tradeoffs between different anomaly detection techniques and their effectiveness on detecting real world container vulnerabilities, we implemented and adapted a set of commonly used unsupervised machine learning techniques and compare their anomaly detection accuracies.

Project: Predicting the Difficulty of Compromise through Modeling how Attackers Discover Vulnerabilities
PI:Andrew Meneely, Rochester Institute of Technology
Co-PI: Laurie Williams, NCSU
Participating Sub-Lablet: Rocester Institute of Technology
Hard Problems: Security Metrics and Modeling

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. Future activities include incorporating risky systems calls, architectural decisions, risky developer activity and human-in-the-loop. One of the goals of the project is how to turn the attack surface into a number to be able to provide actionable feedback. The researchers want to develop metrics that are useful and improve the metric formulation based on qualitative and quantitative feedback.

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. We have collected the intrusion detection data of over 4 billion events from the Regional Collegiate Penetration Testing Competition (https://nationalcptc.org/), and will collect data from the national competition as well. This data will provide 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.

To date, approaches for predicting which code artifacts are vulnerable have utilized a binary classification of code as vulnerable or not vulnerable. To better understand the strengths and weaknesses of vulnerability prediction approaches, vulnerability datasets with classification and severity data are needed. In this work, we use crash dump stack traces to approximate the attack surface of Mozilla Firefox. We then generate a dataset of 271 vulnerable files in Firefox, classified using the Common Weakness Enumeration (CWE) system. We use these files as an oracle for the evaluation of the attack surface generated using crash data. In the Firefox vulnerability dataset, 14 different classifications of vulnerabilities appeared at least once. In our study, 85.3% of vulnerable files were on the attack surface generated using crash data. We found no difference between the severity of vulnerabilities found on the attack surface generated using crash data and vulnerabilities not occurring on the attack surface.

Our systematic literature review on attack surface definitions examines the current body of literature to determine the various definitions of the "attack surface" metaphor and determines clusters of those definitions. The phrase "attack surface" can mean many things to many people, and this study helps clarify what is intended when using the metaphor. Our systematic literature review was approved for publication by the Information and Software Technology.

Project: Principles of Secure Bootstrapping for IoT
PI: Ninghui Li, Purdue University
Hard Problems: Policy-Governed Secure Collaboration

This research builds upon work begun several years ago, motivated by the fact that IoT devices need trust and secure communication—trust between devices and trust between device and users. Constraints, however, limit options, and deployment scenarios determine resource availability, including power supply, computing resources, and serviceability. The research goal is to develop a lexicon and principles to model the different IoT security bootstrapping scenarios and tools to help developers. There is a five-step research plan:

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

Metrics include the number and importance of protocols classified by the framework, the number of vulnerabilities, and the percentage of failed protocols. The success criteria include being able to see the developed lexicon and develop the most important IoT bootstrapping tool.

One area of research examines the need for trust in IoT systems. Since some IoT devices also use cellular networks, we also looked at privacy problems in 4G and 5G cellular networks. To conserve energy, a cellular device needs to stay mostly in an idle, low-power state when not used. To enable a device to respond to incoming calls and various types of messages, a paging (broadcast) protocol is used, such that a cellular device wakes up periodically. We identified inherent vulnerabilities in the paging protocols that enable an adversary who knows the phone number (or other software ids such as Twitter handle) to identify whether a cellular device is present in a physical region. We also discovered attacks that enable the recovery of persistent identity (such as IMSI) of cellular devices. Our findings have been shared with the GSM Association.

PUBLICATIONS

Project: Predicting the Difficulty of Compromise through Modeling how Attackers Discover Vulnerabilities

- Chris Theisen, Hyunwoo Sohn, Dawson Tripp, and Laurie Williams, "BP: Profiling Vulnerabilities on the Attack Surface," *IEEE SecDev Builiding Security*, Cambridge, MA, 2018.
- Chris Theisen, Nuthan Munaiah, Mahran Al-Zyoud, Jeffrey Carver, Andy Meneely, and Laurie Williams, "Attack surface definitions: A systematic literature review," *Information and Software Technology*, Available online, 27 July 2018.

EDUCATIONAL

NCSU faculty are continuing to include Science of Security materials in courses on attack surfaces and policy and social computing. NCSU worked with a Research Experience for Undergraduates (REU) student (supported from internal funds) for development of a cybersecurity game for inclusion in upcoming course offerings. We also interacted with a high-school student in West Lafayette to involve the student in our research.

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

- The PI added a Science of Security related module in the CSC 724 class she regularly teaches.

Project: Predicting the Difficulty of Compromise through Modeling how Attackers Discover Vulnerabilities

- Andy Meneely revised his presentation of attack surfaces and how they apply to risk management in the SWEN 331 Engineering Secure Software course, based

on the research in this Lablet. This course sees 60-80 students per academic year, and is required for all software engineering majors at Rochester Institute of Technology (RIT).

COMMUNITY OUTREACH

Throughout the course of the year, NCSU Lablet researchers discussed Science of Security with colleagues at various conferences that they attended and included information on the Science of Security in conference presentations. NCSU is continuing international collaborations on Science of Security and have added a collaboration with a former postdoc, who is now a faculty member in United Kingdom.

Project: Predicting the Difficulty of Compromise through Modeling how Attackers Discover Vulnerabilities

- Andy Meneely discussed the NSA SoS Lablet with members of the Rochester community at the ImagineRIT festival in May 2018. He taught and demonstrated the basics of cybersecurity to the festival attendees.
- Laurie Williams discussed the NSA Science of Security Lablet in her keynote address at the XP2018 conference in Porto, Portugal.
- Laurie Williams presented a paper on Theisen et al.'s crash dump/attack surface work at the SecDev conference in Cambridge, MA in October 2018.
- Andy Meneely presented this work to the Cybercorps Scholarship for Service program at RIT, getting feedback on the work.

Project: Principles of Secure Bootstrapping for IoT

- We have notified the GSM Association about our findings and are in communication with them.
- We have involved a high school student in the research on IoT security.

NC STATE UNIVERSITY

University of Illinois at Urbana-Champaign

ILLINOIS RECEIVES FUNDING FROM NSA TO DESIGN TRUSTED SYSTEMS

Jun 27, 2018

RESEARCH EDUCATION ABOUT NEWS EVENTS PEOPLE CONTACT US

The University of Illinois at Urbana-Champaign has received a contract from the National Security Agency to lead an effort that will deepen the scientific understanding of the design of trusted systems. Illinois is one of six research institutions to receive funding for a "Lablet" that will conduct foundational security and privacy work over the next five years.

Led by ECE Illinois Head William H. Sanders, the Illinois Lablet will focus on the "Science of Security for Resilient Systems (SSRS)." The multi-university initiative, which includes researchers from 13 institutions and builds on the work of a previous NSA Lablet at Illinois, will consider how to develop security at a system level, examining how the properties of different components contribute to a system's end-to-end security.

The University of Illinois at Urbana-Champaign (UIUC) Science of Security Lablet is contributing broadly to the development of security science while leveraging UIUC expertise in resiliency, which in this context means a system's demonstrable ability to maintain security properties even during ongoing cyber attacks. The Lablet's work draws on several fundamental areas of computing research. Some ideas from fault-tolerant computing can be adapted to the context of security. Strategies from control theory are being extended to account for the high variation and uncertainty that may be present in systems when they are under attack. Game theory and decision theory principles are being used to explore the interplay between attack and defense. Formal methods are being applied to develop formal notions of resiliency. End-to-end system analysis is being employed to investigate resiliency of large systems against cyber attack. The Lablet's work also draws upon ideas from other areas of mathematics and engineering as well.

The UIUC Lablet had four research projects in 2018. The first project, An Automated Synthesis Framework for Network Security and Resilience, is developing an automated synthesis framework, which will derive network state and repairs from a set of specified correctness requirements and security policies. It is also transferring its technology to industry through interactions with Veriflow, a startup company commercializing verification technology that came out of this project's SoS Lablet funding. The second project, A Monitoring, Fusion, and Response for Cyber Resilience is developing a framework to build resilient systems that can detect and respond to malicious activities. The third project, Uncertainty in Security Analysis, is developing uncertainty models associated with an attacker's ability to move through a network and inhibit functionality of the network when the state of the defended network is not fully known. The fourth project, Resilient Control of Cyber-Physical Systems with Distributed Learning, aims to develop methods to protect systems that are using predictive AI models focusing on control systems. Details on these projects can be found below in Fundamental Research and in the Fall Quarterly meeting section.

UIUC Sub-Lablets are Illinois Institute of Technology and University of Texas at Austin. The Lablet is led by Principal Investigator (PI), Sayan Mitra and Co-PIs, David Nicol and William Sanders.

PI Sayan Mitra

PI David Nicol

PI William Sanders

FUNDAMENTAL RESEARCH

Project: An Automated Synthesis Framework for Network Security and Resilience
PI: Matt Caesar
Co-PI: Dong Jin
Participating Sub-Lablet: Illinois Institute of Technology
Hard Problems: Scalability and Composability, Policy-Governed Secure Collaboration, Security Metrics and Models, Resilient Architectures

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. 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.

We continue to investigate automated synthesis of network control to preserve desired security policies and network invariants. Specific invariants include: 1) reduction of reaction time to fix problems; 2) avoidance if introduction of errors in the repair process; and 3) prevention of vulnerabilities. We are also exploring how to synthesize patches to automatically fix critical invariants that were violated by the network controller application. The candidate approach under consideration models both the forwarding behavior of data through the network, control operations conducted on the network, as well as operations between the two. We have formulated a simplified solution for Open Shortest Path First (OSPF) based network using Satisfiability Modulo Theories (SMT) constraints and executed it using a Z3 solver.

We continued the exploration of self-healing network management to address the resilient architecture Hard Problem and application of the methods to applications in cyber-physical energy systems. We developed an enhanced version of the self-healing algorithm that considers both power system observability and communication network characteristics. The current version also assigns weights to the end-devices (i.e., Power Management Unit, PMU) according to certain selected power system metrics before performing the optimization. We implemented a proof-of-concept system in Mininet and are conducting system performance evaluation using the IEEE 30-bus system.

We studied the network intrusion detection system with deep learning models in order to enhance network security and resilience. We designed and implemented a TensorFlow-based deep learning library, called NetLearner. We made the software code of NetLearner publicly available at https://github.com/littlepretty/NetLearner

We continued transfer of our technology to industry through interactions with Veriflow. This startup company now employs over thirty people in the United States and has conducted multiple pilots and deployments across several industry sectors including within the US Department of Defense. Current collaborations target deployment of our verification technology to distributed cloud environments.

Project: A Monitoring, Fusion, and Response Framework for Cyber Resilience
PI: William Sanders
Hard Problems: Policy-Governed Secure Collaboration, Resilient Architectures

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

Our Response and Recovery Engine (RRE) work incorporates modules to monitor current state of a system, detect intrusions, and respond to achieve resilience-specific goals. Intrusion detection in large-scale distributed systems, which is a necessary precondition for intrusion tolerance and resilience, is highly susceptible to malicious manipulation of system data used for detection (e.g., using rootkits and log tampering), which we term "monitor compromise". Existing literature attempts to counteract the problem using reputation systems, which weight the trustworthiness of monitor data based on past trustworthiness of the data, but such systems are themselves subject to "betrayal attacks" and "sleeper attacks". We instead propose the use of data-driven methods for detecting potential monitor compromise. We leverage the insight that systems usually contain multiple monitors that provide redundant information about system activity, so we can use discrepancies between observations of system activity across different monitors to identify potential monitor compromise. For monitor compromise detection, we have developed a data-driven ensemble method for detecting potential monitor compromise using evidential reasoning and data mining. To construct the model for our approach, we have devised a method to mine meaningful correlations between system activity (i.e., events) and the discrete data points produced by monitors (i.e., alerts) and between alerts of different types from heterogeneous historical system data. We have applied our mining method to real data from an enterprise system with meaningful results. We implemented our monitor compromise detection approach using Storm, a real-time stream processing framework, such that it runs in real-time on online monitor data and ran experiments on enterprise network and host data from the National Center for Supercomputing Applications (NCSA) with different, injected compromise scenarios.

We have improved upon the evidential reasoning-based online monitor compromise detection approach that we presented at the 2018 Hot Topics in the Science of Security: Symposium and Bootcamp (HotSoS). We have devised an approach to identify likely monitor compromise using association rule mining that is complementary to our existing evidential reasoning-based approach. We have defined an ensemble method to detect likely monitor compromise that uses the two approaches we have devised to improve overall efficacy. We are currently evaluating the efficacy of the overall approach and preparing a paper submission based on the results.

We constructed a framework to analyze the safety of a system under threat by various attacker models. This work has been accepted into the IEEE Pacific Rim International Symposium on Dependable Computing (PRDC 2018). We developed generic parameterizable state automaton templates that model the effects of an attack. For a given attacker model and system, we can then generate the full state automaton that models the normal system operations under the threat of the specified attacker model. We model the system using network of timed automaton which is suitable for modeling concurrent processes, the progression of time, and physical processes. We consider attacks on network protocols and device commands. More precisely, we assume the attacker has the capabilities of a Dolev-Yao attacker in that he or she can delay, inject, modify, and remove network packets. We develop state automaton templates for each of those capabilities that can be executed either probabilistically or in a deterministic fashion. These templates can be composed and combined together based on the given scenarios. We can then generate a full state automaton that models the normal system operation under a particular set of attacks. We apply our approach to a railway system use case to analyze the safety of the signaling system given a variety of attacker models. Our threat model considers both insiders and outsiders. Outsiders had the capability of delaying and jamming communication to and from the trainborne system and trackside equipment to the system servers. Insiders had the capability of manipulating packets within the system networks. We also considered several safety countermeasures that can potentially deter such threat vectors. We used statistical model checking to verify the safety of the system and our results show that while less skilled outsiders are unable to affect system safety, outsiders who can target vulnerabilities in the network protocol are able to bring the system to an unsafe state even with current modern security protection mechanisms. Insiders are also able to easily affect system state. The safety countermeasures we introduce are able to deter some or all of those attacks although at the added cost of maintenance.

Project: Uncertainty in Security Analysis
PI: David Nicol
Hard Problems: Security Metrics and Models, Resilient Architectures

Cyber-Physical System (CPS) security lapses may lead to catastrophic failure. We are interested in the scientific basis for discovering unique CPS security vulnerabilities to stepping-stone attacks that penetrate through network of intermediate hosts to the ultimate targets, the compromise of which leads to instability, unsafe behaviors, and ultimately diminished availability. Our project advances this scientific basis through design and evaluation of CPS, driven by uncertainty-aware formalization of system models, adversary classes, and security metrics. We propose to define metrics, develop and study analysis algorithms that provide formal guarantees on them with respect to different adversary classes and different defense mechanisms.

Our research focuses on understanding the network security risk and the uncertainty associated with the estimate when security-related properties of the network components are not exactly known. In previous study, we used Boolean random variables to model the existence of a link between two immediate hosts in the network, which indicates the possibility of a lateral movement. Our current investigation generalized this model by modeling the uncertainty in the link existence using beta distribution, a more versatile class of distributions that takes one of many different shapes depending on its two parameters.

Computing the existence of a pathway between two specifically chosen hosts (i.e. reachability analysis) in the generalized model reduces to identifying the reachability distribution, in the form of a multivariate reliability polynomial of beta distributed random variables. This is a difficult problem. However, our initial results highly suggest that in many cases, the reachability distribution can be well-approximated by another beta distribution. This observation aligns with several results from previous studies regarding approximating functions of beta distributed random variables. Our finding, however, applies to a much more general setup. The implication of this result is that under conditions in which the approximation is sufficiently good, reachability analysis on the generalized model can be significantly simplified.

We have completed the design of simulation-based experiments to support the observations made above. The simulations make use of a parameterized model to approximate the reachability polynomial, where the approximation relies on bivariate copula functions and cubic Bezier fitting curves. We use several methods in the literature (method of moments, Maximum Likelihood Estimation (MLE), etc.) to estimate the two parameters of the approximating Beta distribution, then compare it with the actual reachability distribution using several goodness-of-fit metrics (Kolmogorov-Smirnov, Cramer-von Mises, etc.). We proposed a better way of representing uncertainty in a security model using beta distributions. Our observation suggests that in many cases, the reachability distribution of the new model is approximately beta distributed. By knowing the class of the distribution beforehand, the complexity of reachability analysis (in particular) and security analysis (in general) can be greatly reduced.

PUBLICATIONS

Project: An Automated Synthesis Framework for Network Security and Resilience

- Yanfeng Qu, Xin Liu, Dong Jin, Yuan Hong, and Chen Chen, "Enabling a Resilient and Self-healing PMU Infrastructure Using Centralized Network Control," *ACM International Workshop on Security in Software Defined Networks & Network Function Virtualization (SDN-NFV Security 2018)*, Tempe, AZ, March 21, 2018.
- Jiaqi Yan, Dong Jin, and Cheol Won Lee, "A Comparative Study of Off-Line Deep Learning Based Network Intrusion Detection," *10th International Conference on Ubiquitous and Future Networks (ICUFN)*, Prague, Czech Republic, July 3-6, 2018.
- Santhosh Prabhu, Gohar Irfan Chaudhry, Brighten Godfrey and Matthew Caesar, "High Coverage Testing of Softwarized Networks". *ACM SIGCOMM 2018 Workshop on Security in Softwarized Networks: Prospects and Challenges (SecSoN 2018)*, Budapest, Hungary, August 24, 2018.
- Christopher Hannon, Jiaqi Yan, Dong Jin, Chen Chen, and Jianhui Wang. "Combining Simulation and Emulation Systems for Smart Grid Planning and Evaluation," *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, August 2018.
- Santhosh Prabhu, Gohar Chaudhry, Brighten Godfrey, and Matthew Caesar, "High Coverage Testing of Softwarized Networks," ACM SIGCOMM 2018 Workshop on Security in Softwarized Networks: Prospects and Challenges, Budapest, Hungary, August 24, 2018.

Project: A Monitoring, Fusion, and Response for Cyber Resilience

- Carmen Cheh, Ken Keefe, Brett Feddersen, Binbin Chen, William Temple, and William Sanders, "Developing Models for Physical Attacks in Cyber-Physical Systems," *ACM Workshop on Cyber-Physical Systems Security and Privacy*, Dallas, TX, November 3, 2017.
- Benjamin Ujcich, Adam Bates, and William Sanders, "A Provenance Model for the European Union General Data Protection Regulation," *7th International Provenance and Annotation Workshop (IPAW '18)*, London, UK, July 9-13, 2018.
- Benjamin Ujcich, Adam Bates, and William Sanders. "Data Provenance for Accountability Mechanisms and Properties." *First Workshop on Supporting Algorithm Accountability Using Provenance: Opportunities and Challenges*, London, UK, July 12, 2018.
- Benjamin Ujcich, Samuel Jero, Anne Edmundson, Qi Wang, Richard Skowyra, James Landry, Adam Bates, William Sanders, Christina Nita-Rotaru, and Hamed Okhravi, "Cross-App Poisoning in Software-Defined Networking," *2018 ACM Conference on Computer and Communications Security (CCS '18)*, Toronto, Canada, October 15-19, 2018.
- Carmen Cheh, Ahmed Fawaz, Mohammad A. Noureddine, Binbin Chen, William Temple, and William Sanders, "Determining the Tolerable Attack Surface that Preserves Safety of Cyber-Physical Systems," *IEEE Pacific Rim International Symposium on Dependable Computing*, Taipei, Taiwan, December 4-7, 2018.

EDUCATIONAL

Project: An Automated Synthesis Framework for Network Security and Resilience

- Kevin Jin served as the Ph.D. Colloquium Chair for the ACM SIGSIM Conference on Principles of Advanced and Distributed Simulation, May 2018
- Kevin Jin supervised two undergraduate student research projects in Fall 2017 and Spring 2018. Matthew Caesar supervised five undergraduate student research projects during this same timeframe. Matthew also participated in the 2018 University of Illinois Undergraduate Research Symposium.
- We presented a research poster "Distributed Virtual Time System for Embedded Linux Devices" at the 7th Greater Chicago Area System Research Workshop (GCASR) on April 2018.
- Matthew Caesar taught a Networking Laboratory class in the fall semester and developed a new Cybersecurity module for his class. This module gives students the opportunity to set up and configure security features of routers and switches in a virtualized environment. Students configure ACLs and VLANs to ensure desirable security properties such as segmentation and access control. The lab is structured to give students direct hands-on experience with these techniques, making them confident to use these techniques in the field.

COMMUNITY OUTREACH

Project: An Automated Synthesis Framework for Network Security and Resilience

- Kevin Jin served as a technical program committee member of the 2018 International Conference on Information and Communications Security (ICICS)
- Kevin Jin gave a technical talk at the CODES summer workshop at Argonne National Lab in July. The topic of the talk was "Scalable Simulation and Modeling Framework for Evaluation of Software-Defined Networking Design and Applications."
- Xin Liu presented our SDN-NFV Security'18 paper, in conjunction with the ACM Conference on Data and Application Security and Privacy (CODASPY), in Arizona on March 2018.
- Matthew Caesar continues to serve as Chief Science Officer of Veriflow, a company commercializing technology spun out of our Science of Security Lablet work.
- Kevin Jin was invited to the NSF Visioning Workshop on Programmable Security in a Software Defined World, August 2018.

ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

University of Kansas

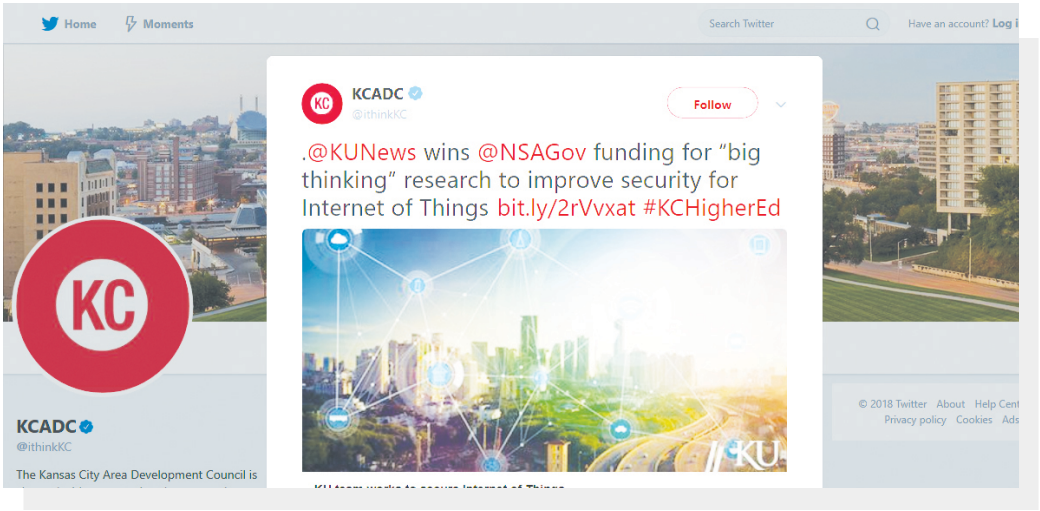

The University of Kansas (KU) Science of Security and Privacy Lablet is making interdisciplinary contributions to security science synthesizing knowledge and innovation from computer science, electrical engineering, psychology, sociology, and philosophy. The Lablet's work focuses on the foundational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

The KU Lablet initiated five projects in 2018. Our first project, Cloud-Assisted IOT Systems Privacy, is developing a method to enable cloud-assisted, privacy-preserving machine learning classification over encrypted data for IoT devices. Under the Formal Approaches to the Ontology and Epistemology of Resilience project we are developing an epistemology and ontology for framing resilience. Our third project, Side-Channel Attack Resilience, focuses on reducing micro-architectural side-channels by introducing new OS abstractions while minimally modifying micro-architecture and OS. The Scalable Trust Semantics and Infrastructure project is developing a formal basis for remote attestation that supports scalability to distributed, heterogeneous systems and demonstrating the sufficiency and soundness of remote attestation processes. The final project, Secure Native Binary Executions, aims to design a mechanism to both analyze and quantify the level of security provided and performance penalty imposed by different solutions to various security risks affecting native binaries.

Details on these projects can be found in Fundamental Research, below, and in the Fall Lablet Quarterly section.

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

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

- The Department of Psychology
- The Department of Sociology

The KU Lablet is led by Dr. Perry Alexander who serves as the Principal Investigator. He is assisted by Bob Rummer and Patricia Bergman who provide research and administrative support. Lead project PIs include Bo Luo, Fengjun Li, John Symons, Heechul Yun, Garrett Morris, and Prasad Kulkarni. The KU Sub-Lablet, the University of Tennessee, will be added in 2019. Approximately 10 graduate students and 7 faculty are supported by the program assisted by several research staff and undergraduate students. Our Lablet team meets regularly with an industrial advisory board that provides input on research directions, facilitates outreach and helps establish a regional, midwestern cyber security community.

PI Perry Alexander

FUNDAMENTAL RESEARCH

Project: Cloud-Assisted IOT Systems Privacy
Lead PI: Fengjun Li
Co-PI: Bo Luo
Hard Problems: Scalability and Composability, Security Metrics and Models, Resilient Architectures

The key to realizing the smart functionalities envisioned through the Internet of Things (IoT) is to securely and efficiently communicate, store, and make sense of the tremendous data generated by IoT devices. Integrating IoT with cloud platforms for their computing and big data analysis capabilities becomes increasingly important, since IoT devices are computational units with strict performance and energy constraints. However, when data are transferred among interconnected devices or to the cloud, new security and privacy issues arise. In this project, we investigate the privacy threats in the cloud-assisted IoT systems, in which heterogeneous and distributed data are collected, integrated and analyzed by different IoT applications. The goal of the project is to develop a privacy threat analysis framework to provide a systematic methodology for modeling privacy threats in the cloud-assisted IoT systems.

Recent security incidents exploited vulnerabilities embedded in IoT devices due to poor security design and problematic implementation. However, with a large number of heterogeneous IoT devices online, it is difficult to develop a universal security solution that addresses the vulnerability of each individual device. To tackle this problem, we investigate the attack surface through which IoT devices are exposed and develop an isolation-based approach to separate the private network where IoT devices are deployed from the public network via a newly designed IoT gateway.

Machine Learning plays an important role in making sense of the tremendous data generated by the IoT devices. An emerging machine intelligence platform, known as Machine Learning as a Service (MLaaS), makes it easier for users to analyze IoT data faster and deliver more accurate insights at smaller costs. However, it raises several concerns regarding the security and privacy of IoT data as well as the proprietary Machine Learning models. To address this problem, we developed a method to enable cloud-assisted, privacy-preserving Machine Learning classification over encrypted data for IoT devices that allows a cloud server to interact with Machine Learning service providers on behalf of the resource-constrained IoT devices in a privacy-preserving manner and shift the load of computation-intensive classification operations off the devices.

Project: Formal Approaches to the Ontology and Epistemology of Resilience
PI: John Symons
Hard Problem: Resilient Architectures

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.

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. Consider a communications network where an initial evaluation of resilience might involve deciding whether or not the system continues to function. Specifically, is it possible to send and receive messages reliably through the communications network? This is a functional account of the individuation of the system. 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.

Resilience is certainly tied to function in important ways. The value of a communications network is its functional properties, and it is likely more resilient if it continues to perform its functions reliably. While pragmatic considerations are important, conditions for persistence or individuation are not properly understood in terms of pragmatic preferences with respect to the functional properties of systems. The fact that it is important

that the network functions in accordance with requirements is distinct from the question of what it is that makes the network resilient. We might have, for example, an invulnerably resilient network with less than ideal functionality. As decisions on trade-offs are made in the context of security, it is necessary to understand distinctions of this kind.

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. Philosophers have tackled the problem of determining the correct approach to ontological questions (questions about the nature of the kinds of things that exist) and can shed light on many of the questions concerning resilience. Not only are many philosophers familiar with the graph theoretic foundations of network theory, but they are also used to dealing with questions concerning persistence using techniques from modal logic and category theory. More importantly, philosophers are used to recognizing distinctions in these domains that others often miss. It is the contention of this group, for example, that excessive attention to abstract functional level descriptions can potentially distract us from other aspects of systems that contribute to resilience and are important to defend.

Project: Side-Channel Attack Resilience
PI: Heechul Yun
Hard Problems: Resilient Architectures

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

Micro-architectural side-channel attacks such as Meltdown and Spectre have received great interest from the research community and general public alike, and new attacks are being discovered. 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. Our initial work provides a major technical foundation – new memory abstractions and supporting micro-architecture designs – that the continuing project intends to utilize to reduce micro-architectural side-channel attack surfaces. We have made progress on Spectre attack resistant architecture design. Specifically, we developed an initial proof-of-concept prototype, which extends the Gem5 simulator's O3CPU (out-of-order) model, that is able to mitigate the Spectre attack efficiently.

Project: Scalable Trust Semantics and Infrastructure
Lead PI: Perry Alexander
Co-PI: Garrett Morris
Hard Problems: Scalability and Composability, Policy-Governed Secure Collaboration

Remote attestation provides a run-time capability for appraising system behavior and establishing trust. Using remote attestation, an appraiser requests evidence describing a target. The target responds by performing measurement to gather evidence, then adds cryptographic signatures to assure integrity and authenticity. The appraiser takes the evidence and assesses the target's behavior to determine if the target is who and what it claims to be.

Remote attestation has enormous potential for establishing trust in highly distributed IoT and Cyber-Physical Systems. However, significant work remains to build an overarching science of remote attestation without which scaling trust to large, complex systems will prove difficult. To develop basis for remote attestation the Lablet is developing a semantics of measurement, attestation and appraisal. This semantics includes metrics for soundness and sufficiency of evidence, semantic mechanisms for identity and attestation, formal definitions of evidence, and meta-evidence appraisal. We are exploring systematic mechanisms for establishing and evaluating roots of trust in heterogeneous systems that include legacy components. Finally, we are working on developing formal, executable representations for attestation protocols and tools for static analysis.

Along with developing a semantics for trust and remote attestation, an important goal is implementing and scaling a prototype trust infrastructure. Our approach includes hierarchical frameworks for building trusted systems including reusable attestation managers and attestation protocols in support of layered attestation in multi-platform, heterogeneous systems. Sharing the same semantics for trust across distributed systems supports appraisal across enterprise systems. Furthermore, it reduces the burden of developing new infrastructure as systems grow and add new functionality. Of particular interest are systems that employ Trusted Platform Modules (TPMs) and Virtualized Trusted Platform Modules (vTPMs), trust aggregation, trust in legacy systems and trust as a service.

PUBLICATIONS

Project: Cloud-Assisted IOT Systems Privacy

- Lei Yang and Fengjun Li, "Cloud-Assisted Privacy-Preserving Classification for IoT Applications," *IEEE Conference on Communications and Network Security* (CNS), 2018.

- Lei Yang, Chris Seasholtz, Fengjun Li and Bo Luo, "Hide Your Hackable Smart Home from Remote Attacks: An Extra Network-Level Safeguard," *European Symposium on Research in Computer Security (ESORICS)*, 2018.

Project: Formal Approaches to the Ontology and Epistemology of Resilience

- John Symons, "Metaphysical and scientific accounts of emergence: Varieties of fundamentality and theoretical completeness," *Emergent Behavior in Complex Systems Engineering* (pp. 3-20), Saurabh Mittal, Saikou Diallo, Andreas Tolk (eds.), March 2018 DOI: 10.1002/9781119378952.ch1

- John Symons, "Brute facts about emergence," *Brute Facts*, Elly Vintiadis (ed), Oxford University Press. (pp.177-196).

- John Symons and Rasmus Rendsvig, "Epistemic Logic," *Stanford Encyclopedia of Philosophy*, (Revised Edition)

Project: Side-Channel Attack Resilience

- Farzad Farshchi, Prathap Kumar Valsan, Renato Mancuso, Heechul Yun, "Deterministic Memory Abstraction and Supporting Multicore System Architecture," *Euromicro Conference on Real-Time Systems (ECRTS)*, July 2018.

- Michael Garrett Bechtel, Elise McEllhiney, Minje Kim, and Heechul Yun, "DeepPicar, A Low-cost Deep Neural Network-based Autonomous Car," *IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA)*, 2018

EDUCATIONAL

Lablet funding directly supports training of PhD, MS and BS students working towards their degrees. Additionally, the Lablet provides synergistic opportunities to students involved in non-Lablet projects. Several students in the KU Scholarship for Service program participate in Lablet research. Numerous undergraduate students participate in Lablet activities as a part of senior projects and undergraduate research projects.

Working with the KU Electrical Engineering and Computer Science Department Lablet, researchers have developed an Undergraduate Certificate in Cyber Security to complement the existing MSIT in Cyber Security. A unique aspect of the certificate is receiving credit for active participation in the KU competitive cyber security team, the JayHackers, and associated hands-on training. Pending administrative approval, the certificate will be offered starting in the Fall 2019 term.

Project: Formal Approaches to the Ontology and Epistemology of Resilience

A new research seminar examining privacy issues from an interdisciplinary perspective is being led by John Symons. The seminar will focus first on philosophical notions of privacy and integrate concepts from social sciences and computer science.

Project: Scalable Trust Semantics and Infrastructure

Two undergraduate seminars are offered as introductions to the mathematical basis for establishing trust in computing systems. An undergraduate Honors Symposium on the mathematical origins of computer science is available to all incoming freshmen in the KU Honors Program. The purpose of the symposium is to introduce the roots of mathematical systems useful for establishing trust in computing systems in a manner that is accessible to a general audience. A weekly undergraduate research seminar examining tools and techniques for software verification in support of trusted systems is offered to undergraduate students working on Lablet projects. The seminar uses topics from our remote attestation work as examples of systems for verification.

COMMUNITY OUTREACH

The KU Lablet held the Securing the Internet of Things Workshop and Advisory Board Meeting October 2-3 on the University of Kansas Edwards Campus in Kansas City. The workshop recruited regional companies and government agencies with the goal of building a security community in the Midwest. Keynote presentations were given by Brigadier General Jennifer Buckner, Army G-3/5/7 Director of Cyber, Aaron Weissenfluh, CISO Cboe Global Markets, and Brian McClendon, former Vice President of Research at Google and Uber. A panel of advisory board members discussed the promises and pitfalls of the ever-growing Internet of Things, and KU Lablet researchers presented updates on their research. A student poster competition provided a forum for regional researchers to present their work. With over 120 attendees, the workshop was a success and plans are already being made for next year.

A key element of our community building is engagement with an external Advisory Board. We recruited members from a cross-section of industry and government representing CPS domains that include finance, healthcare, infrastructure, security, and manufacturing. Current Advisory Board includes members from Cerner, Garmin, Sprint, Koch Industries, Honeywell FM&T, Cboe Global Markets (formerly BATS) and The Federal Reserve. At the inaugural Spring 2018 advisory board meeting we introduced the board to Lablet research and the concept of Science of Security, and the topics of our initial projects. Our discussion topics related to applicability of the research and identifying gaps in our approach. At the Fall 2018 meeting we discussed outreach, workforce development and research directions. Specific topics included supply chain as a battlefield, Internet of Things, and machine-to-machine communication, as well as training and support for Lablet research.

Project: Cloud-Assisted IOT Systems Privacy

- Fenjun Li and Bo Luo, "Exploring IoT-Cloud Systems to Secure IoT Applications", invited talk at The CPS Security and Education Workshop, UNC Charlotte, July 26, 2018.

- Bo Luo, "Supporting Data Privacy in Cloud-Assisted IoT Applications," invited talk School of Information Science and Technology, University of Science and Technology of China, June 19, 2018.

- Fengjun Li, “Privacy-Preserving Classification for IoT Applications,” invited talk at HotPrivacy Day of IEEE Symposium on Privacy-Aware Computing, Washington D.C., September 26, 2018.

Project: Formal Approaches to the Ontology and Epistemology of Resilience

- John Symons, “Incompressibility and reversibility in prediction,” International Symposium on Forecasting (University of Colorado Boulder) June 18 2018

- Computational Methods and the Future of Science (https://cps-vo.org/group/CMFS) will be held during Q2 2019 at KU and will feature a keynote and technical presentations focusing on the intersection of science and computation. Our Lablet resilience work will feature prominently in the program

Project: Scalable Trust Semantics and Infrastructure

- Perry Alexander, “Trust and Proof,” invited presentation at the 2018 KCNSC Trust Consortium, Department of Energy’s Kansas City National Security Campus, Grandview, MO, June 5-6, 2018.

- Sarah Helble, Aaron Pendergrass, Pete Loscocco, Perry Alexander, Adam Petz, Paul Rowe, John Ramsdell, “Principles of Layered Attestation,” invited presentation at the High Confidence Software and Systems Conference, Annapolis, MD, May 7-9, 2018.

- David Hardin, Perry Alexander, Konrad Slind, “VeriCores: Cyber-Instrumenting Devices Built from Verified Components,” invited presentation at the High Confidence Software and Systems Conference, Annapolis, MD, May 7-9, 2018.

- Lablet researchers work jointly with MITRE, Johns Hopkins University Applied Physics Lab (JHUAPL) and NSA Information Assurance Research in the development of trust semantics.

KU THE UNIVERSITY OF KANSAS

Vanderbilt University

NSA Lablet at Vanderbilt to make sure America keeps moving after hacks

by Heidi Hall May. 7, 2018, 4:10 PM

Xenofon Koutsoukos, professor of computer science, computer engineering and electrical engineering, is heading up the new NSA Lablet. (Steve Green/Vanderbilt)

According to one of the widely accepted definitions, Cyber-Physical Systems (CPS) are engineered systems where functionality emerges from the networked interaction of computational and physical processes. Complex CPS abound in modern society and it is not surprising that many of these systems are safety and mission critical, thus making them a target for attacks. Even under normal conditions, CPS face complex issues cross-cutting many disciplines with significant implications on essential system functions. Adding cyber-attacks in all their insidious variety creates a massive challenge that cannot be neglected due to the potential consequences. Because of its significance, security and resilience have attracted considerable attention in many CPS application domains. Because of the heterogeneity and complexity, methodologies that improve CPS security are very diverse with different objectives, specifications, and constraints resulting in a broad body of knowledge. Research efforts are starting to use scientific methods and results to shape technology, practice, and policy in protecting systems from attackers, detecting intrusions, and recovering from compromises. However, scientific methods remain underutilized and they do not adequately address the involved interdisciplinary socio-technical aspects. Beyond the complex structure and interactions, security and resilience properties emerge from complex interrelationships between engineered systems and humans; they are not explained by understanding the individual elements of the system, and are highly dynamic in response to changing environment and circumstances. What is needed is a Systems Science of Secure and Resilient CPS which brings together interdisciplinary research with the goal of identifying, exploring, and understanding patterns of complexity which cross disciplines and application domains.

The Vanderbilt University (VU) Lablet aims at developing the principles governing secure and resilient CPS in adversarial environments and using these principles for system design and management. System approaches require a mix of methods and tools. The Lablet has four projects for 2018: Analytics for Cyber-Physical System Cybersecurity, Foundations of CPS Resilience, Multi-Model Test Bed for the Simulation-based Evaluation of Resilience, and Mixed Initiative and Collaborative Learning in Adversarial Environments.

The projects build upon our strengths in 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. Details on these projects can be found below in Fundamental Research and in the Fall Quarterly meeting section.

VU Sub-Lablets are Massachusetts Institute of Technology, University of California, Berkeley, and University of Texas at Dallas. The VU Lablet is led by Principal Investigator (PI) Xenofon Koutsoukos.

PI Xenofon Koutsoukos

Fundamental Research

Project: Analytics for Cyber-Physical System Cybersecurity
PI: Nazli Choucri
Sub-Lablet: Massachusetts Institute of Technology
Hard Problem: Policy-Governed Secure Collaboration

Mounting concerns about safety and security have resulted in an intricate ecosystem of guidelines, compliance measures, directives and policy reports for cybersecurity of all critical infrastructure. By definition, such guidelines and policies are written in linear sequential text form that makes them difficult to integrate, or to understand the policy-technology-security interactions, thus limiting their relevance for science of security. This project addresses the Hard Problem of Policy-Governed Secure Collaboration related to cyber-physical security of critical infrastructure, and focuses on a generic and fundamental capability, namely smart grid power systems central to modern critical infrastructures. The challenges are to develop a structured system model from text-based policy guidelines and directives in order to identify major policy-defined system-wide parameters, situate vulnerabilities, and map security requirements to security objectives. In addition, advance research on how multiple system features respond to diverse policy controls to strengthen the security of fundamentals in cyber physical systems. The goal of this project was to develop text-to-analytics methods and tools focusing on CPS domains such as smart grids.

We articulated the details of the activities for the first year and the foundational tasks for the overall project. The core data base consists of the "raw text" of policy documents, directives and guidelines--a powerful source of directives based on the extensive research and scientific inquiry from a wide range of institutions, stakeholders, and agencies. We have also defined the policy e-landscape and reviewed the entire e-policy ecology to concentrate on the core. We identified key NIST policy documents selected as the core of our CPS investigation. These documents cover both generic directives pertaining to large scale domain-independent critical infrastructure and specific documents related to the smart grid.

Designed as a multi-method approach, our research is anchored in transforming the text-based policy-defined operational features of NIST cybersecurity policies and guidelines into a structured model of system-features and information flows for use as a platform for system-wide policy impact analysis of cybersecurity directives, and exploration of "malicious pathways".

We have identified the policy relevant ecosystem, formalized rules for extracting structured data from text-based policy materials, and identified relevant linkages for representing and implementing cybersecurity measures. We are currently identifying "missing elements" and constructing internally consistent structures to represent, organize, metricize, and manage text-based materials essential for application and development of cybersecurity analytics.

We have sequestered the key policy documents that are related to smart grid cybersecurity. (NISTIR 7628, Guidelines for smart grid cybersecurity, is the central document). A detailed study is being conducted to identify the key pieces/elements for implementation of cybersecurity measures. The team has begun formalizing the rules for extraction of data from text and has completed a validation study to test the scalability, portability, and transposability of the rules to other policy domains.

Project: Foundations of CPS Resilience
PI: Xenofon Koutsoukos
Hard Problem: Resilient Architectures

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. The proposed approach integrates redundancy, diversity, and hardening methods for designing either passive resilience methods that are inherently robust against attacks or active resilience methods that allow responding to attacks.

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. We introduce a framework for quantifying cybersecurity risks and optimizing the system design by determining security investments in redundancy, diversity, and hardening. To demonstrate the applicability of our framework, we present a case study in water distribution systems. Our numerical evaluation shows that integrating redundancy, diversity, and hardening can lead to reduced security risk at the same cost.

Attacks in CPS which manipulate sensor readings can cause enormous physical damage if undetected. Detection of attacks on sensors is crucial to mitigate this issue. We study supervised regression as a means to detect anomalous sensor readings, where each sensor's measurement is predicted as a function of other sensors. We show that several common learning approaches in this context are still vulnerable to stealthy attacks,

which carefully modify readings of compromised sensors to cause desired damage while remaining undetected. We model the interaction between the defender and attacker as a Stackelberg game in which the defender chooses detection thresholds, while the attacker deploys a stealthy attack in response. We present a heuristic algorithm for finding an approximately optimal threshold for the defender in this game, and show that it increases system resilience to attacks without significantly increasing the false alarm rate.

Networked and distributed CPS increase attack surfaces and cybersecurity risks. Distributed diffusion is a powerful algorithm for multi-task state estimation which enables networked agents to interact with neighbors to process input data and diffuse information across the network. Compared to a centralized approach, diffusion offers multiple advantages that include robustness to node and link failures. We consider distributed diffusion for multi-task estimation where networked agents must estimate distinct but correlated states of interest by processing streaming data. By exploiting the adaptive weights used for diffusing information, we develop attack models that drive normal agents to converge to states selected by the attacker. The attack models can be used for both stationary and nonstationary state estimation. In addition, we develop a resilient distributed diffusion algorithm under the assumption that the number of compromised nodes in the neighborhood of each normal node is bounded and we show that resilience may be obtained at the cost of performance degradation. We evaluate the proposed attack models and resilient distributed diffusion algorithm using stationary and non-stationary multi-target localization.

Finally, CPS are subject to typical cyber-attacks such as code injection attacks. With the tightly coupled nature of cyber components with the physical domain, these attacks have the potential to cause significant damage if safety-critical applications such as automobiles are compromised. Moving target defense techniques such as Instruction Set Randomization (ISR) have been commonly proposed to address these types of attacks. However, under current implementations, an attack can result in system crashing which is unacceptable in CPS. CPS necessitate proper control reconfiguration mechanisms to prevent a loss of availability in system operation. Our work addresses the problem of maintaining system and security properties of a CPS under attack by integrating ISR, detection, and recovery capabilities that ensure safe, reliable, and predictable system operation. Specifically, we consider the problem of detecting code injection attacks and reconfiguring the controller in real-time. The developed framework is demonstrated with an autonomous vehicle case study.

Project: Multi-Model Test Bed for the Simulation-Based Evaluation of Resilience
Lead PI: Peter Volgyesi
Co-PI: Himanshu Neema
Hard Problems: Security Metrics and Models, Resilient Architectures

We developed the Science of SecUre and REsilient Cyber-Physical System (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 was to extend these design and simulation capabilities for better understanding the security and resilience aspects of CPS systems. These improvements include support for the design of experiments (exploring different parameters and/or strategies), alternative CPS domains (connected vehicles, railway systems, smart grid), incorporating models of human behavior, and executing multistage games.

In planning the overall architecture and priorities of the new testbed, we rely on Veins as an integrated simulation capability (for the highway traffic and connected vehicle scenarios), WebGME as our collaborative web-based front-end, and Jupyter Notebook for executing and evaluating simulation runs. The new architecture is a significant departure from our previous SURE testbed, thus we bring-in and re-use previously developed elements selectively, instead of continuing the development of the previous testbed. One of the major design and implementation tasks is to develop programmable attack and mitigation strategies (courses of actions) in the Veins simulation engine. This capability had been provided by the C2 Windtunnel platform in the SURE testbed, previously. We have built the initial infrastructure of the testbed, and although the current version does not provide developed scenario models, all architectural elements above are included and integrated. In parallel with the development of the web-based testbed, we developed a simple emergency vehicle-based scenario using Veins/OMNeT++ directly. The scenario investigates how smart/connected traffic lights can improve the time needed for emergency vehicles to travel through a realistic street network (the Vanderbilt campus). The next step is to adopt this scenario in the web-based framework.

The development of the multi-model testbed included multiple experiments in two CPS domains. In the transportation domain, the experiments incorporate vehicle-to-infrastructure (V2I) technology to minimize the response time of emergency vehicles. In the power domain, the focus is on vulnerabilities in the electricity market infrastructure and how they can affect the power grid. Models in both domains are integrated with cyber-attacks and courses-of-actions models for performing security and resilience studies. These experiments will be refined and integrated in the web-based collaborative design environment and will be provided as security research scenarios.

As part of our framework's metrics-driven evaluation capability, we are developing a cyber-attack library and a language for their systematic incorporation in security and resilience experiments. We call this language Courses-of-Action (COA) modeling. The cyber-attacks include Denial-of-Service (DOS) attack, Packet Delay Attack, Data Corruption Attack, and Data Integrity Attack. These attacks will be configurable so that they can be deployed in any of the key network nodes during any time-interval in the simulation with any values for configuration of attack parameters. Currently, we have already implemented the DOS and Delay Attacks. We are continuing to implement the rest of the planned cyber-attacks as well as develop the COA modeling language to utilize these cyber-attacks in the form of security and resilience experimentation scenarios.

PUBLICATIONS

Project: Foundations of a CPS Resilience

- Bradley Potteiger, Zhenkai Zhang and Xenofon Koutsoukos, "Integrated Instruction Set Randomization and Control Reconfiguration for Securing Cyber-Physical Systems," *Symposium and Bootcamp on the Science of Security, HotSoS 2018*, Raleigh, NC, April 10-11, 2018.
- Jiani Li and Xenofon Koutsoukos, "Resilient Distributed Diffusion for Multi-task Estimation," *The 14th International Conference on Distributed Computing in Sensor Systems (DCOSS 2018)*, Bronx, NY, June 18-20, 2018.
- Himanshu Neema, Bradley Potteiger, Xenofon Koutsoukos, Gabor Karsai, Peter Volgyesi, and Janos Sztipanovits, "Integrated Simulation Testbed for Security and Resilience of CPS," *The 13th ACM/SIGAPP Symposium on Applied Computing (SAC 2018)*, Pau, France, April 9-13, 2018.
- Amin Ghafouri, Yengeniy Vorobeychik, and Xenofon Koutsoukos, "Adversarial Regression for Detecting Attacks in Cyber-Physical Systems," *27th International Joint Conference on Artificial Intelligence and 23rd European Conference on Artificial Intelligence (IJCAI-ECAI 2018)*, Stockholm, Sweden, July 13-19, 2018.
- Himanshu Neema, Bradley Potteiger, Xenofon Koutsoukos, CheeYee Tang, and Keith Stouffer, "Metrics-Driven Evaluation of Cybersecurity for Critical Railway Infrastructure," *National Symposium on Resilient Critical Infrastructure, Resilience Week 2018*, Denver, CO, August 20-23, 2018.
- Aron Laszka, Waseem Abbas, Yevgeniy Vorobeychik, and Xenofon Koutsoukos, "Synergistic Security for the Industrial Internet of Things: Integrating Reduncancy, Diversity, and Hardening," *IEEE International Conference on Industrial Internet (ICII 2018)*, Bellevue, WA, October 21-23, 2018.
- Akos Ledeczi et al., "Teaching Cybersecurity with Networked Robots," SIGCSE 2019, Accepted for publication.

EDUCATIONAL

Project: Analytics for Cyber-Physical System Cybersecurity

- The project research design and its implications were included in the curriculum in the course at entitled Cybersecurity offered jointly by the Department of Political Science and the Sloan School of Management at Massachusetts Institute of Technology (MIT).

Project: Foundations of a CPS Resilience

- We developed Roboscape, a collaborative, networked robotics environment that makes key ideas in computer science accessible to groups of learners in informal learning spaces and K12 classrooms. RoboScape is built on top of NetsBlox from Vanderbilt University, an open-source, networked, visual programming environment based on Snap! that is specifically designed to introduce students to distributed computation and computer networking. RoboScape provides a twist on the state-of-the-art of robotics learning platforms. First, a user's program controlling the robot runs in the browser and not on the robot. There is no need to download the program to the robot and hence, development and debugging become much easier. Second, the wireless communication between a student's program and the robot can be overheard by the programs of the other students. This makes cybersecurity an immediate need that students realize and can work to address.
- In 2018, we organized two summer camps on CPS security with 24 students in grades between 7 and 12. The curriculum is based on RoboScape. Details on the summer camps can be found in Section 3.
- We documented the technology behind RoboScape and the hands-on curriculum. Also, we evaluated the summer camps looking into the learning gains of the students.

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

- Two undergraduate students participated during the summer months. The students acquired knowledge to design, build and execute test scenarios targeting the Vanderbilt campus street network using the Veins platform.

We participated in the 2018 National Tennessee Valley Corridor Summit held in Oak Ridge, TN on May 29-31, 2018 to present the research efforts of the Lablet.

COMMUNITY OUTREACH

Project: Analytics for Cyber-Physical System Cybersecurity

- Our research design and current work was presented at the weekly meeting of the research/consortium on "Cybersecurity at MIT Sloan" focusing on proof-of-concept for the overall design. This is important because we are adopting a multi-method approach that must meet multiple types of critical reviews.

Project: Foundations of a CPS Resilience

- Our research was presented in the following conferences: HotSoS 2018, DCOSS 2018, IJCAI-ECAI 2018, and National Symposium on Resilient Critical Infrastructure.

- We participated in the 9th annual Computational Cybersecurity in Compromised Environments (C3E) workshop including research posters focusing on security and resilience problems in CPS with learning-enabled components.

- We had a technical meeting about security and resilience of CPS with Dr. James Peery, Associate Laboratory Director of Global Security and Kendal Card, Division Director, DOE-In Programs, Global Security, Oak Ridge National Laboratory.

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

- Our research was presented and the testbed was demonstrated at the Fujitsu System Integration Laboratory in Tokyo, Japan in September 2018.

- We had a poster presentation at the C3E Workshop in Atlanta, GA, September 2018.

Science of Security Quarterly Meetings

Winter 2018 Quarterly and **Kickoff Meeting**

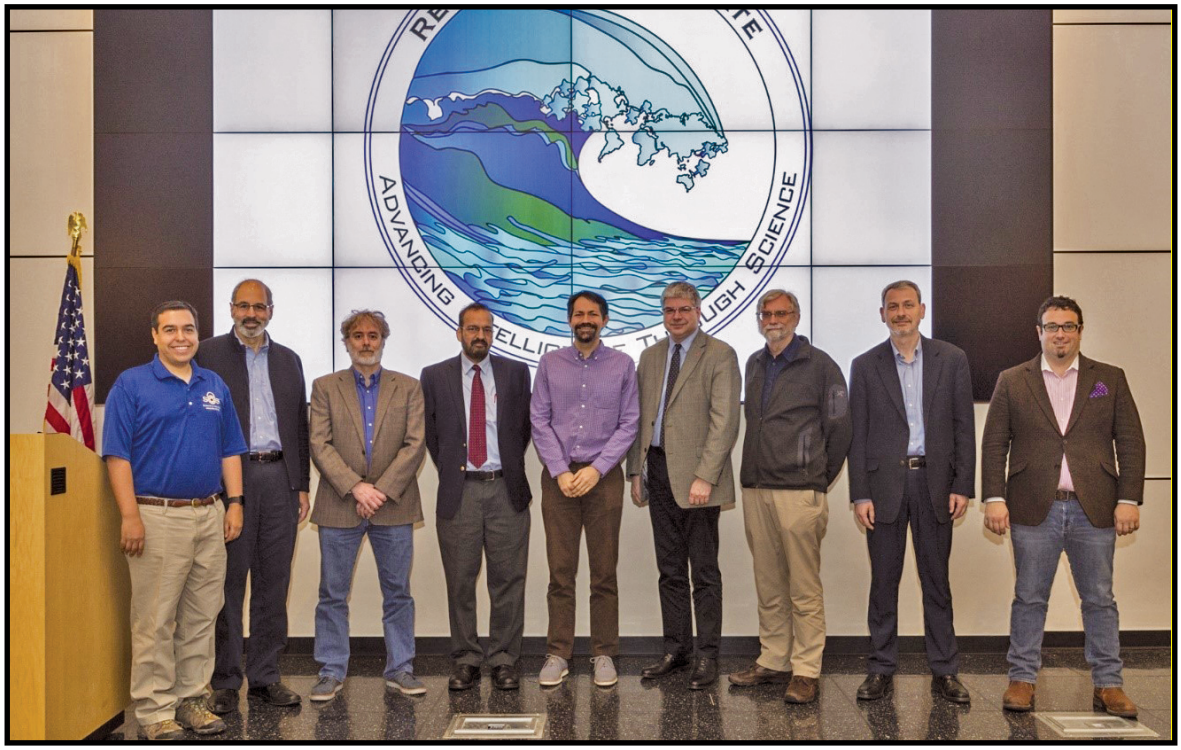

Lablet PIs and Co-PIs meet with NSA to for kickoff of 3rd generation of Lablets
Pictured left to right: Dr. Adam Tagert (SoS Technical Director), Dr. Bill Scherlis (CMU), Dr. Perry Alexander (KU), Dr. Munindar Singh (NCSU), Dr. Jonathan Aldrich (CMU), Dr. Bill Sanders (UIUC), Dr. David Nicol (UIUC), Dr. Xenefon Koutsoukos (VU), Dr. Serge Egelman (ICSI)

The Science of Security and Privacy (SoS) initiative held its kickoff meeting for the third generation of Lablets. The meeting initiated relationships between NSA and Lablet researchers. One of the objectives for the new generation of Lablets is to expand collaborative engagements by taking advantage of the synergy between scientific and operations perspectives in order to increase the impact of the research projects. The SoS research effort addresses some of the most significant cybersecurity research challenges aligned against the five Hard Problems. Sixty-five attendees met on 13-14 March 2018 at the Laboratory for Telecommunications Sciences in College Park, MD to kick off the effort.

The two-day meeting presented several NSA researcher perspectives and introduced the twenty new Lablet research projects. The Lablets are 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).

Adam Tagert, SoS Technical Director, welcomed the attendees and introduced the NSA SoS initiative leadership team and the Lablet Principal Investigators (PIs).

CMU: Bill Scherlis, Jonathan Aldrich
ICSI: Serge Egelman
KU: Perry Alexander
NCSU: Laurie Williams, Munindar Singh
UIUC: William Sanders, David Nicol, Sayan Mitra
VU: Xenofon Koutsoukos

Dr. Tagert noted that the third generation of Lablets has the same goals as the prior two generations, and he emphasized building up the foundational aspects of SoS anchored in scientific methods, models, and approaches. He is expecting research breakthroughs and the development of new technologies and tools under the initiative. NSA SoS leadership is seeking better engagement between NSA and the Lablets, and wants to facilitate tech transfer arising from Lablet research.

Summaries of all the presentations are provided below, and selected presentations can be found at https://cps-vo.org/node/36719/browser

NSA Research Perspectives

Information Assurance Research Overview
George Coker (NSA—Information Assurance Research Group)

Dr. Coker's presentation, "Cybersecurity: Effects at Scale" focused on a science-based approach to cybersecurity. He described the evolution of Information Assurance Research that has led to a focus on cyber, and noted that resilient and scalable solutions are now the crux of the issue. He noted that reversing the asymmetric advantage attackers now enjoy will require the achievement of defensive effects at scale which, in turn, necessitates a science-based approach in order to understand if we're making a difference at scale. See https://cps-vo.org/node/54404

Privacy
Travis Breaux (NSA—Secure Systems Architecture and Analysis)

Dr. Breaux identified the concepts of privacy to include confidentiality and secrecy as well as personhood and autonomy, and also addressed self-determination and self-discovery. He focused on the project of Protected Autonomy which shifts the focus to data integrity, noting that assured decision making depends on trusted algorithms for making predictions and curating information for human consumption. One of the research questions being addressed under this project is how to compose multiple non-symbolic programs (ML models) into a single simulation. He identified technical challenges as realistic datasets and algorithms, simulation infrastructure design, and mitigation design and evaluation. Other projects include Privacy Enhanced Architecture (use-based privacy and private information retrieval); Compliance Assistance (querying policy to improve legal and engineering coordination, emergent compliance—discovering rules from norms, and organizational practices to design privacy into software) and Mobile and IoT privacy. See https://cps-vo.org/node/54406

Cyber-Resiliency
Jim Holt (NSA—Adaptive Cyber-Defense Systems)

The speaker's presentation focused on what his team is trying to accomplish and how the Lablets can participate. In addressing autonomous cyber defense, he noted that cyber attackers have an asymmetric advantage over defenders and that another detector, sensor, or tool doesn't help; rather, the challenge is how to change the balance of power. To address the challenge, the team hypothesizes that the balance can be changed through increasingly autonomous cyber defense and the strategic use of deception in cyber defense. Human decision-makers can't respond quickly enough, but they can achieve speed and scale through autonomous cyber defenses. They acknowledge that autonomy is a spectrum, but that humans can delegate decision-making and actions to the system. Strategic use of deception deals with crafting answers to attackers to influence their decision-making since virtually every success an attacker is able to have is possible because the networks provide correct answers to attacker questions.

A key concept in this strategy is cyber resilience: anticipate, withstand, recover, and evolve. The speaker addressed building autonomous cyber defenses and then demonstrating their effectiveness scientifically. Challenges in doing so include the following: system complexity; defining success; differing missions, values, goals, and environments; undefined, unbounded, and evolving threats; and the fact that full-loop execution requires many components which are interdependent. He addressed some of the goals including well-defined measures of success, reproducible experiments, and collaboration. The challenges to implementing the approach include incomplete, uncertain, and/or untrusted data and testing. The decision-making piece includes looking at tools and techniques in AI, ML, planning systems, game theory, robotics and optimization, with the eventual choice likely being a hybrid of these elements. Their plan is to start simply and evolve, building simple versions to demonstrate the framework's extensibility. They have also considered building an open-source, shareable testbed, and the speaker addressed multiple ways for the Lablets to participate and collaborate. See https://cps-vo.org/node/54407

Cyber-Physical Systems/IoT
Raj Pal (NSA—Trust Mechanisms)

Dr. Pal's presentation was entitled "Building towards a Trustworthy IoT Ecosystem" and he said that while the IoT team has been focusing on applied research, they are looking to partner with the Lablets to expand foundational research. He noted that the ideal system would have end-to-end trust, root of trust, be remotely attestable, and have trustworthy integrity verification. He defined IoT as a set of network-capable products when platform integrity cannot be verified with confidence, and said it was important to be able to measure these devices in a trustworthy way. He believes that the research challenge, in a resource and functionally constrained platform, is how to overcome hardware and functional barriers and incorporate the mechanisms identified for richer platforms into the constrained device. Relevant Information Assurance Research motivations included emerging technologies to further the mission and blending trusted and untrusted devices. He identified IoT research interests as 8,16, and 32 bit platforms; trusted computing base; trustworthy integrity verification; memory separation; and P2P networks. The desired impact of the research, he concluded, was to develop solutions to shape the security landscape and influence industry. See https://cps-vo.org/node/54405

Lablet Project Presentations

Hard Problem: Policy-Governed Secure Collaboration

Uncertainty in Security Analysis (UIUC)
Frank Nguyen

The speaker noted that there are many models in security analysis of computer systems, all of which require information about devices, interconnections, services, configurations, attacker, and defender. The problem, he said, is that in practice, the information about the model is incomplete, which leads to either making simplifying assumptions or explicitly modeling the incomplete information as an input uncertainty. The goal is to develop techniques for expressing uncertainty in the input of the security models, and for assessing uncertainty in the model output. He talked about the early work done in this area (presented at HoTSoS 2017) including how topological

uncertainty networks impact reachability analysis and Extended Uncertain Graph (EUG). Theoretical results show that EUG is capable of describing any joint distribution of edge existence and that uncertainty analysis in EUG is tractable if the Boolean functions are monotone. The technical approach includes formalisms for expressing uncertainty in model input, analysis techniques for assessing model output uncertainty, the UQ framework for scientific evaluation of outcomes, and demonstration on large-scale real-life attack graphs. He tied the research to Policy-Governed Secure Collaboration by noting that the attacker's ability to harm a system depends on the policy used to protect it—there might be uncertainty in knowing exactly what policy is used. The Year 1 milestone is to develop the formalisms. See https://cps-vo.org/node/54408

Analytics for Cyber-Physical System Cybersecurity (VU)
Nazli Choucri, Massachusetts Institute of Technology

This research is focused on analytics to better understand the structure of policy, and the focus is to construct methods-sequence to extract value for policy guidelines for cyber security. The issue was cast as a policy problem: guidelines encourage passive compliance rather than active performance, and policy documents are framed as "stand alone" and unconnected to related documents. The research uses NIST reports on the cybersecurity of the smart grid, and researchers employ analytic methods to capture full value of cybersecurity policies and guidelines. The overall objective is to integrate smart grid cybersecurity policies with a research approach that undertakes a multi-method modular investigation of cybersecurity policy documents in order to create coherent and verifiable analytics for SoS. Professor Choucri identified three dimensions of SoS contributions:

- Policy analytics: replicable methods for analysis of systems and enterprise-wide cybersecurity
- Research: multi-methods for deep analysis of cybersecurity, policies and framework
- Education: demonstrate use of multi-methods for dynamic analysis of cybersecurity

See https://cps-vo.org/node/54401

Operationalizing Contextual Integrity (ICSI)
Serge Egelman

Professor Egelman presented work on mobile device apps that has led up to what they plan to do in the future, addressing privacy as contextual integrity. He noted that inappropriate data flows violate contextual information norms; contextual information norms are modeled using data subjection, data sender, data recipient, information type, and transmission principle (constraints). In questioning what this means for user-centered design, he suggested that an app should only provide notice when reasonable privacy expectations are expected to be violated. He described studies done on permission requests when a phone was inactive (a training exercise), addressed the use of ML to detect when context has changed from expected data use to unexpected, and then described a second experiment done in real-time that confirmed earlier findings. 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)

Principles of Secure Bootstrapping for IoT—NCSU
Ninghui Li, Purdue University

Professor Li noted that this research builds upon work begun several years ago, citing the motivation as the fact that IoT devices need trust and secure communication—trust between devices, and trust between device and users. Constraints, however, limit options, and deployment scenarios determine resource availability, including power supply, computing resources, and serviceability. The research goal is to develop a lexicon and principles to model the different IoT security bootstrapping scenarios and tools to help developers. He described a five-step research plan:

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

Metrics include the number and importance of protocols classified by the framework, the number of vulnerabilities, and the percentage of failed protocols. The success criteria include being able to see the developed lexicon and develop the most important IoT bootstrapping tool.

Contextual Integrity for Computer Systems—ICSI
Michael Tschantz

Dr. Tschantz described the overall goal of the research as converting the philosophical theory of contextual integrity into terms computer scientists can use. He noted that there is no agreement on what a context is: philosophers and computer scientists have different understandings, with philosophers focusing on abstract spheres of life, and computer scientists focusing on the concrete. The goal is to develop models of context and contextual integrity that meet computer scientists on their own truth. Relevant research questions include accounting for privacy in the design of multi-use computer systems that cut across contexts; modeling the adaptation of contexts to changes in technologies; and determining how contextual integrity relates to differential privacy. The current organizing hypothesis is that contexts are defined by a purpose. He noted that the privacy norms of a context promote the purpose, and that purpose restrictions are ubiquitous. He proposed several possible models including game models, Markov decision process models, partially observable Markov decision process models, and multi-agent influence diagrams. Some of the

challenges are that contexts don't exist in a vacuum, contexts might be in competition, privacy is multifaceted, and people often disagree. He identified potential outcomes as progress on defining privacy, further accountability for big data systems that cut across contexts, and enabling policy governed privacy with respect to collaboration.

Obsidian: A Language for Secure-by-Construction Blockchain Programs—CMU
Joshua Sunshine, CMU
Jon Bell, George Mason University

This research is focusing on Etherium and Hyperledger platforms and is also focusing on smart contracts. The researchers are designing a program language—the goal is to do human-centric language design so that resultant programs are secure—and it will be Obsidian language. They noted that smart contracts have forced ordering, and Typestate language enforces ordering constraints. They addressed secure collaboration in scientific research, noting the wide variety of artifacts and the desire for a decentralized mechanism to share across organizations. They also addressed the issue of complex access controls expected by different parties. A strawman approach is a trusted third party, but it is not ideal because of the temporal nature of saving. They noted that BitLedger allows researchers to share and that it uses blockchain system for access controls. They reported on controlled experiments of programmers, and the speed and security of writing Obsidian language smart contracts. They are also looking for other blockchain applications and addressed collaboration as key to their approach.

Scalable Trust Semantics and Infrastructure—KU
Perry Alexander

Professor Alexander noted that although KU is a new Lablet, the university has a history working trust issues with Information Assurance Research and predecessor organizations. He described the criteria for when you should trust a system as the following: you know its identity; you know it's built from good parts; you know it's behaving as expected. He also addressed semantic remote attestation--he presented a simplistic model and then explained why it's more complicated than the model presented. In developing a science of trust there are five tasks:

- Semantics of trust
- Measurement, attestation, and appraisal
- Roots of trust
- Attestation protocols
- Implementing and scaling infrastructure

This research is currently focused in two areas: development of attestation protocol semantics (under way now with Information Assurance Research, MITRE, John Hopkins University Applied Physics Lab); and soundness and sufficiency of measurements.

Governance for Big Data—ICSI
Serge Egelman

In introducing the topic, Professor Egelman suggested that the risk in governance for big data is that access control does not capture privacy requirements. He addressed sensitive inferences and reidentification, noting that it is difficult to redact sensitive information from rich data sets and that often sensitive data can be reidentified using additional information outside the data set or proxies. He suggested that Machine Learning will find such correlations automatically; binary allow/deny access control fails to capture this well. In discussing limiting sensitive inferences, he pointed out several related issues, including differential privacy, encryption and access control, and fairness issues. A new data governance approach focuses on accountability and relates more to accounting and auditing. This project aims to synthesize computer science abstractions with governance goals. 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.

Designing for Privacy—ICSI
Serge Edelman

The project focuses on designing for privacy holistically: from "privacy by design" to "privacy with design", i.e., designing with privacy throughout whole life cycle. Professor Egelman noted that design interventions for privacy can occur at a lot of stages and levels, and that the goal of the project is to develop a new toolbox of techniques and help designers understand when best to apply tools. He addressed defining privacy in contextual, situational, and relational ways, and identified its dimensions as 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. He also plans to hold privacy design workshops to address engineering practices, methods, and tools, bringing together practitioners, researchers, and policy-makers.

Hard Problem: Resilient Architectures

Foundations of CPS Resilience—VU
Xenofon Koutsokos

Professor Koutsoukos addressed the need 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 unanticipated attacks. The foundation of CPS resilience includes developing principles and methods for designing and analyzing resilient CPS architectures that deliver required service utility in the face of compromised components; integrating redundancy, diversity, hardening methods for designing passive resilience methods that are inherently robust against attacks; and developing active resilience methods that allow response to attacks including optimal control and reconfiguration. He discussed automated and connected vehicles within the context of the resiliency of intelligent transportation models. The research seeks to develop models to form hypotheses for simulations. Model components include diversity, redundancy, and hardening, and integrating those components for designing passive and active resilience methods (passive—robust against attacks; active—allow responses). He addressed how to improve structural robustness in networks citing the need for more than redundancy by adding diversity and hardening. He discussed game-theoretic formulation to find optimal resiliency and optimal defense strategy in the face of attacks and used examples of attack and defense in the transportation network. System models include configuration, attack, detection, mitigation, and responsive attack. The goal is to develop a model that allows finding optimal resilient configurations of CPS by integrating redundancy, diversity and hardening in the face of strategic attacks. He concluded by noting that considering attacks in CPS in all their insidious

variety creates a massive challenge can't be neglected due to potential consequences. See https://cps-vo.org/node/54403

Coordinated Machine Learning-Based Vulnerability Discovery and Security Patching for Resilient Virtual Computing Infrastructures—NCSU
Helen Gu

The presentation focused on Docker security and addressed attack surfaces and vulnerabilities in Linux kernel, Docker engine, and container applications. The existing approach is static security analysis and scheduled patching. In the researchers' experiments, this approach fails to detect 90% of vulnerabilities, displays high false alarms, and shows memory inflation caused by unnecessary security patching. Their proposal 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.

Model-Based Explanation for Human-in-the-Loop Security—CMU
David Garlan

Professor Garlan provided context for the research by noting that automation is becoming increasingly important for modern systems, and many systems require combinations of automated and human involvement to handle security attacks. The problem is how to create effective coordination. The solution is simple in explanation but difficult in practice--the system needs to understand what humans can do, and humans need to understand the system. Decisions must be made about which tasks are to be allocated to the system vice humans, and humans must be able to trust in automated actions. Automation is improved by learning based on what humans do. Prior research included the adoption of a control systems view of system autonomy and led to the development of RAINBOW framework; earlier work also looked at humans as actuators who effect changes. Current research addresses putting the human in the planning area. A key idea associated with this work is to use formal models for planning as the basis of human-understandable explanation. Technical challenges include explaining a plan that is computed from a probabilistic system model and determining the basis for selecting the best alternative. He concluded by noting that system resilience and security can be enhanced through automation, but autonomous decision making is often opaque. We need better transparency through an explanation of the models used for planning, and we can inform system autonomy by allowing a system to learn by example from expert user behavior.

Predicting the Difficulty of Compromise through Modeling how Attackers Discover Vulnerabilities—NCSU
Andy Meneely, Rochester Institute of Technology

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. Professor Meneely noted that a literature review to classify attack surface definitions led to six clusters of definitions which differ significantly (methods, avenues, flows, features, barriers, and vulnerabilities). He further discussed 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. Future activities include incorporating risky systems calls, architectural decisions, risky developer activity and human-in-the-loop. The researchers want to develop metrics that are useful and improve the metric formulation based on qualitative and quantitative feedback.

Formal Approaches to the Ontology and Epistemology of Resilience—KU
John Symons

Professor Symons began by identifying the epistemic challenge as "what is the best way to understand resilience?" and the ontological challenge as "what is resilience and how does it emerge?" He noted that this work contributes to the establishment of interdisciplinary Science of Security by focusing on its most important concept at the fundamental level—formalism—and attention to neglected aspects. He said that with respect to cybersecurity, the view is that the network model is valuable but incomplete. The speaker addressed the definition and aspects of resilience, noting that a system can be said to be resilient if it is prepared for attack or disruption. maintains its identity, isn't compromised to the point of not being itself, bounces back, and learns from past disruptions or attacks and adapts. He also pointed out non-network aspects of resilience including the resilience of the mechanisms underlying functions, the functions themselves, the distinction between robustness (static) and resilience (dynamic), and the conditions underlying the emergence and persistence of the systems in question. With respect to the ontological aspect (the nature of resilience), he stated that the emergence of resilient norms, for example, is not amenable to network theoretic treatment but essential to security. Plans for foundational research for science of resilience include existing foundational research and exploring the formalism. He identified open issues as: what are the constraints and factors that allow for resilience to emerge; how do we understand the role of emergent norms in the Science of Security; trust (roots of trust); and common knowledge. The researchers plan to run a series of cross-disciplinary seminars and build on the KU network model for work.

Hard Problem: Metrics

Multi-model Test Bed for the Simulation-based Evaluation of Resilience—VU
Peter Volgyesi

Professor Volgyesi described the existing cloud-based testbed environment for CPS developed under the Science of SecUre and REsilient CPS (SURE) project, and proceeded to discuss directions for future research. Areas to be explored include:

- New CPS domains (smart grid, IoT)
- Streamlined infrastructure for the Traffic CPS
- Different abstraction levels
- Hardware in the loop
- RF domain
- Transactive energy domain

He summarized the program goals as follows: integrate proven best-of-class simulators for CPS domains; add cyber security aspects (attack/defense programs); multiple levels of

abstractions; collaborative design environment with versioning and libraries; and cloud-based simulation and analysis.

Securing Safety-Critical Machine Learning Algorithms—CMU
Matt Frederickson

Professor Frederickson noted that Machine Learning is ubiquitous and that it works in many applications, sometimes outperforming humans. He discussed the Deep Neural Network (DNN) model for image classifying and addressed the example of an adversary that can change the features (pixels in images) that are given to the model and thereby change the outcomes (evasion attack). He raised the questions of whether attacks work if they have to be physically realizable and inconspicuous, and whether attacks can be robust to training and model selection. He presented a target attack centered around face recognition, addressing impersonation, dodging, and implementing attacks with physical changes. The challenge, he noted, is building models that are resilient to physical attacks. In addressing vulnerability, the researchers have looked at which parts of the DNN model were most susceptible to attack. They are seeking to leverage explainable features in classification to make models more resilient.

Hard Problem: Scalability and Composability

An Automated Synthesis Framework for Network Security and Resilience—UIUC
Matthew Caesar

This project builds on earlier work and is focused on building a rigorous method for Science of Security, developing techniques for performing and integrating security analyses to automatically and rigorously study hypotheses about the end to end security of a network. The Automated Synthesis Framework (ASF) goal is a new network architecture for resilience with a focus on network data flow security. The approach is to leverage network synthesis to automate experiments and then apply results. Professor Caeser identified the following three tasks:

- Network control syntheses—develop algorithms and systems that perform automated synthesis
- Network software analysis and modeling—develop frameworks for writing secure network control programs
- Resilient and self-healing network applications

With respect to the technical approach, the ASF consists of a network model, controller, policy, verification engineering, and correction engine. The project is representing network state with a policy model, and the speaker cast the problem as an optimization problem.

A Monitoring Fusion, and Response for Cyber Resilience—UIUC
Mohammad Noureddine

This project continues work done earlier, and the speaker identified the three components of the research as:

- Monitor deployment and compromise detection--monitor placement was done in an earlier phase; the new phase is looking at dealing with monitor compromise
- Rich data fusion for improved detection--prior work started at the host level and incorporated more diverse data sources; since they don't know whether correlation chains are malicious or administrative, they added new data sources from outside the network to address that question
- Automated response and recovery—the motivation is to lessen the burden on system administrators and enable response by designing autonomous agents to monitor the activity and respond; earlier work dealt with lateral movement and modeled zero-sum game, and formulated the same problem as a control theory problem, while the new work addresses puzzle difficulty selection and applies science

In the future, they plan to develop adaptive techniques to combat large-scale volumetric attacks with the goal being to push insights from control and game-theory into the reactive security realm.

Cloud-Assisted IoT Systems Privacy—KU
Fengjun Li

Professor Li noted that the privacy problem is amplified in IoT because of the long and complex value chain and the large number of stakeholders included in data processing. The goal of this research is to develop a privacy threat analysis and protection framework to provide a systematic methodology for modeling and mitigating privacy threats in cloud-assisted IoT systems. Challenges include identifying which information is considered privacy and needs to be protected since privacy protection is subjective; not all users are aware of privacy risk, and there is privacy leakage due to big data analytics. The speaker addressed privacy threats including information disclosure, identifiability, profiling, and information linkage. She identified Privacy-Enhancing Technologies (PET) as a potential solution, but raised the issue of how to select and combine appropriate PETs to address identified privacy threats with acceptable performance, within hardware, software, and data constraints. The research plan is a pilot project focused on privacy-preserving classification for cloud-assisted IoT applications. The desired research outcomes are a privacy threat analysis framework and a privacy protection framework.

Side-Channel Attack Resistance—KU
Heechul Yun

Professor Yun addressed the needs for Intelligent CPS and System On a chip (SOC). The speaker noted that micro-architectural side-channels in advanced embedded computing hardware are serious security threats in CPS and can compromise spatial and temporal isolation needed to implement secure and safe CPS. The project will investigate new abstractions, OS, and architecture designs for side-channel attack-resistant computing platforms for CPS. The project goal is to develop micro-architectural side-channel attack-resistant OS and architecture enhancements. By focusing on critical memory, the high cost of supporting strong isolation can be minimized. Tasks include critically and side-channel aware OS-level memory management on existing hardware, and new abstractions in both hardware and OS.

Spring 2018 Quarterly
In lieu of a Spring Quarterly, the SoS community met in April at HotSoS 2018 hosted by the NCSU Lablet in Raleigh, NC.
See Section 3 for full HotSoS 2018 details.

Summer 2018 Quarterly

University of Illinois at Urbana-Champaign

The Summer 2018 Quarterly Science of Security and Privacy (SoS) Lablet meeting was held at the University of Illinois at Urbana-Champaign (UIUC) July 31-August 1, 2018, and was hosted by Bill Sanders, Co-Principal Investigator (PI) at UIUC. This was the first Quarterly meeting since the kickoff of the new contract, and included all of the Lablets awarded contracts in 2018: Carnegie Mellon University (CMU), the International Computer Sciences Institute (ICSI), University of Kansas (KU), North Carolina State University (NCSU), UIUC, and Vanderbilt University (VU). Since research under this contract is not yet mature, the meeting focused on issues and case studies associated with transition to practice and outreach. In welcoming attendees, Adam Tagert, SoS Technical Director, noted that the first day would deal with technology transfer since one of the goals of the SoS program is to have more impact and influence on moving research into practice. The second day will concentrate on outreach, promoting rigorous research methods, and increasing participation in STEM.

Nadia Carlsten, Program Manager for Transition to Practice (TTP) in the Cyber Security Division (CSD) of the Homeland Security Advanced Research Projects Agency in the Department of Homeland Security Science and Technology organization (DHS S&T), spoke on how the TTP program at DHS identifies promising federally-funded cybersecurity research and accelerates transition from the laboratory to the marketplace through partnerships and commercialization. The goal of the TTP program is to bridge the "Valley of Death" in order to get federally funded technologies to market where they can make an impact, specifically by identifying mature technologies that address an existing or imminent need; increasing utilization through partnerships and commercialization; and improving the long-term ability of federally-funded labs to efficiently transition technology. The TTP program guides researchers through a proven maturation process and connects the researchers with investors, technologists, and government to create better solutions through partnerships. Transitions include startups, commercial products, open source solutions, and government-wide use. Dr. Carlsten noted that the "value is being able use TTP as a one stop shop" which lowers risk considerably. The TTP office spends significant time scouting out technologies that are ready and engaging stakeholders (researchers, investors, users) to get them to interact and understand the benefits of the interaction. Following initial selection of the technology, the TTP process includes training for researchers on how to market their technology, market validation, test and evaluation, pilots, outreach, utilization, and licensing. As a result of the TTP program, $118M has been leveraged which has led to 40 technologies, 15 commercial transitions, and 19 pilots, as well as other achievements. She concluded by noting that the office is specifically focused on transition, not R&D, that they are not limited to DHS-funded technologies, and that they address any cyber-related area. Since commercialization is a priority, the TTP approach is flexible and offers multiple paths to achieve that objective. The theme of technology transfer was continued in presentations that described three successes stemming from Lablet research.

Lujo Bauer, CMU, addressed successful technology transfer based on one project that measured privacy risk and another that provided a method for measuring password strength. To provide decision-makers and privacy advocates a way to understand privacy-related tradeoffs, the first project, measuring privacy risk, sought to develop a repeatable method to measure privacy to help make decisions about sharing data. The speaker described the research approach and findings, noting that the results were built into a prototype by Pacific Northwest Lab using the algorithm computed. The second project measured password strength and looked at the issue of making passwords harder to guess without making them too hard to remember. The research addressed the issues of how to tell if an idea for new scheme is a good idea or not and what to measure. The speaker discussed the methodology used in the research, which involved CMU's Security Behavior Observatory, and pointed out that the results influenced new NIST guidelines, including a de-emphasis on length in favor of complexity, blacklists, feedback to users, Open Source NN password strength estimator, and meter.

The second presenter, Matt Caesar, UIUC, described Lablet work that led to the creation of a startup company called Veriflow. Veriflow was described as a "science-based security company," whose goal is to make networks secure and provide a rigorous, automated mathematical method to test complex systems. Such efforts are designed to prevent catastrophic failure and provide rigorous formal verification and continuous network verification to check all network-wide data flows. To support users,

Technology transfer panel discussion, left to right: Nazli Chouri, Perry Alexander, Bill Sanders, Michael Tschantz, Lujo Bauer, Munindar Singh, Adam Tagert

Veriflow builds a comprehensive model of the network and performs experiments motivated by users' hypotheses. Based on Lablet research that started in 2012, Veriflow secured $10M in first-round funding and currently has 30 people employed at two sites, and sales and/or pilots at multiple companies. Veriflow's success led to the creation of a new market segment by Gartner. Using formal verification and formal logic for reasoning about security, the use of Veriflow has led to the discovery of breaches most of the time it's deployed. He concluded by noting that industry is discovering the benefits of science of security with more rigor and new market segments.

Ehab Al-Shaer, University of North Carolina at Charlotte (UNCC), described a Lablet project at NCSU that automated cyber attack response with provable guarantees of success in mitigating the attacks. The CLIPS/Active SDN project started with the initial Lablet funding and involved the development of a flexible/expressive policy specification and a provably correct policy refinement engine to enable safe and efficient construction and execution of a course of action workflow with analysis ad reconfiguration. Transition to practice was funded by NSA to develop and demonstrate CLIPS/Active SDN case studies at Johns Hopkins Applied Physics Lab (APL). It was implemented and tested on virtual SDN using Mininet and OpenDaylight, then deployed on a real APL testbed. He concluded by noting that TTP is costly but rewarding: it directs the research to the right challenges; provides feedback on design; and provides balance between theory and engineering.

Following the presentations, Adam Tagert moderated a panel discussion addressing technology transfer. Panelist represetning Lablets and Sun-Lablets provided technology transfer case studies — some that worked and some that did not. Panelists Nazli Choucri (Massachusetts Institute of Technology (MIT), Vanderbilt University Sub-Lablet), Perry Alexander (KU), Bill Sanders (UIUC), Michael Tschantz (ICSI), Lujo Bauer (CMU), and Muninder Singh (NCSU) described approaches and efforts that varied widely. Each panelist provided a short introduction about their technology transfer initiatives.

Nazli Choucri talked about three collaborators and discussed the difference between targeted and non-targeted transition, noting that non-targeted efforts seem to be more effective.

Perry Alexander pointed out that what is now the KU Lablet started as a tech transfer lab supported by the state of Kansas. The goal of the program was to get research out of the lab; success was measured by the number of startups and jobs created. He provided some lessons-learned from the KU experience, including documentation and the competition for funding, talent, and time.

Bill Sanders noted that UIUC/ECE has about 30 active startup companies and that technology transfer is part of the UIUC culture. He cited the need for industry "pull" vice technology "push". He expressed the hope that SoS wouldn't lean too far toward tech transfer since the science is still critical and we can't pretend to be able to solve all the problems.

Michael Tschantz noted that he has only experienced technology transfer failures, attributing them to legal issues and ultimate lack of interest.

In discussing the prerequisites for tech transfer, Lujo Bauer said it was necessary to acknowledge the benefit to the recipient for tech transfer and to make sure it wasn't simply a personal research interest. He also stated that it was important to be close to stakeholders, to address scalability and adoptability from the outset, and to maintain a focus on the long horizon.

Munindar Singh stated that he believes the discussion about tech transfer is sometimes too narrow, thinking simply in terms of startups, since tech transfer can occur by IP being shared among companies. He addressed factors influencing tech transfer, the different types of tech transfer contributions, and ways of promulgating tech transfer in academia.

The moderated panel discussion addressed other means of technology transfer including transferring ideas through classes and curricula. One of the panelists noted that it is often difficult to introduce new ideas into curricula, but once it's done the students provide a multiplier effect; he also pointed out the value of faculty-to-faculty transfer, and while it's not routinely or easily done, it works well once accomplished. The panel discussed industry days as a means of tech transfer, and because of mixed results debated whether industry research labs would be a better target. While some industry engagements have worked well if the participants are highly technical, one panelist believes that for tech transfer to work, senior managers from industry also need to be involved. The panelists addressed workforce development, noting that the demand for students is huge and that students want a real problem to work on. One panelist made the point that the best form of tech transfer is transferring students. In responding to the question of how research results are being transferred

back to NSA, both panelists and government members of the audience focused on students again, one describing access to students who are developing deep expertise in areas of interest as a force-multiplier. One of the panelists noted that students are particularly drawn to research that is responding to real problems. In closing, Dr. Tagert asked the panelists what NSA could do to help the Lablets have more impact from research results. One suggested providing unclassified descriptions of classified problems; another asked that NSA stimulate their imaginations; another focused on the importance of researcher-to-researcher contacts. Finally, one of the panelists noted that they need to be given problems that advance science and engineering and have an impact, describing that as the "sweet spot" for science, engineering and tech transfer.

The second day of the Quarterly focused on outreach, with presentations on NSA outreach initiatives and a panel discussion on Lablet outreach activities.

Brad Martin addressed NSA's academic outreach strategy. This talk included NSA's vision and an overview of its efforts in multiple fields including, but not limited to, STEM and Intelligence Analysis, Language, and Cybersecurity education in support of NSA's mission. He addressed outreach within NSA and with other agencies, K-12 institutions, colleges and universities, as well as industry and state and local governments. He discussed the four strategic thrusts of the academic outreach program: 1) cultivate academic relationships; 2) influence curriculum; 3) broaden and deepen skills capacity; and 4) tackle hard mission problems. He emphasized the fact that NSA is making a greater effort to coordinate its activities internally. He noted that since NSA uses contracts rather than NSF-style grants, they focus on sponsoring research in which they have a strong interest that will allow them to dialogue with the researchers. He talked about the presence of on-campus NSA labs—the Laboratory for Telecommunications Sciences at University of Maryland College Park and the Laboratory for Analytic Sciences at NCSU—and a visiting professors program. Other activities he addressed included the Forum on Cyber Resilience that NSA sponsors along with the National Academy of Sciences, NIST and NSF, cooperation with standards bodies looking for input, the National Physical Sciences Program, the Research Experience for Undergraduates (REU) program, and the CyberCorps program with NSF. He also noted the mathematical sciences program which has been underway for over forty years. With respect to efforts addressing NSA hard mission problems, he pointed to a statistics advisory group and 8-12 week long summer camps involving cleared people and academics to focus on specific challenges.

Capt. Tina McAfee, USAF, addressed NSA's participation in the International Science and Engineering Fair (ISEF), the largest and most prestigious high school science fair whose roots go back 100 years. Follow-up ISEF activities include building a STEM pipeline, connecting the research with NSA researchers and mission needs, and growing intelligence relationships. (Details on NSA's involvement in ISEF 2018 can be found in Section 2 of the Annual Report.)

Adam Tagert talked about other SoS initiatives including NSA's activities at RSA where SoS was one of nine topics at the NSA booth, the 2018 Best Scientific Cybersecurity Paper Competition (details in Section 3), a new award for the best conference paper at SOUPS 2018 (details in Section 2), and the second Cyber-Physical Systems Summer Camp held at Vanderbilt University **(details in Section 3).**

Ahmad Riley presented details on other academic outreach programs including the On-Ramp program and the Meyerhoff Scholars Program. The On-Ramp program, which is ending in 2018, is a program with University of Maryland Baltimore County (UMBC) begun in 2009 to increase representation of underrepresented groups. UMBC has used the funds provided by the On-Ramp program for the Meyerhoff Scholars program which has supported over 200 students. The program has provided a comprehensive array of support services for students' academic, professional, and interpersonal growth. One element of the program is the Summer Bridge Program, mandatory attendance at an accelerated six-week residential program the summer preceding students' first college semester. As of Spring 2018, 65% of NSA UMBC scholars (35% female) have enrolled or have completed MS or PhD programs. Since 2011, NSA analysists have been regularly engaged with students to discuss STEM careers and provide career advice and mentoring.

There was another panel discussion that addressed successes and failures in outreach activities associated with promoting rigorous research methods and increasing participation in STEM. Moderated by Adam Tagert, the panelists were Perry Alexander (KU), Xenefon Koutsoukos (VU), Lindsey McGowan (NCSU), Stephanie Rosenthal (Chatham University, CMU sub-Lablet), Jana Sibestik (UIUC), Michael Tschantz (ICSI), and Andrea Whitesel (UIUC). They covered a variety of programs including youth summer camps and internships, early interest generation, industry community days, weekly seminars, the appointment and use of a corporate advisory board, and programs that build working toys. All of the panelists agreed that programs that interest and stimulate students need to begin early. One of the panelists addressed education outreach and developing a curriculum on privacy in age-appropriate manners for K-12; another talked about a program that included major outreach to the pre-college audience; and another addressed a CPS summer camp for high school students that uses robotics to teach security. Future outreach ideas included a Cyber Patriots camp for K-12 students; a cyber competition/hackathon; a SoS Lablet summer school; expanding a summer intern program from 8 to 10 weeks; additional workshops and conferences; teacher professional development opportunities; and a MOOC (Massive Open Online Course) for high school students. Several panelists noted that in order to stimulate interest in graduate degree programs, recruiting needs to be done in the freshman and sophomore years. One panelist noted the particular difficulty of recruiting graduate students later than their sophomore year because of offers from industry.

In discussing some of the challenges of working with younger students, panelists noted that it requires a lot of perseverance and suggested working with existing after-school programs and using teachers as a force multiplier. Additional topics of discussion included experience with MOOCs, the pros and cons of on-line learning, and successful outreach initiatives that the panelists and participants experienced personally. Finally, in responding to the question of what could motivate students to work for NSA, panelists cited the value of the programs identified by the NSA speakers as well as the challenging work. They suggested expanding personal contacts between NSA researchers and students, and cited the success of the Senior Executive Academic Liaison (SEAL) program in engaging potential hires on campus.

Fall 2018 Quarterly

Carnegie Mellon University

The Fall 2018 Quarterly Science of Security and Privacy (SoS) Lablet meeting was held at Carnegie Mellon University (CMU) on October 29-30, 2018, and was hosted by Bill Scherlis, Co-Principal Investigator (PI) for CMU. The focus of the Quarterly was to introduce new Lablet research projects funded for 2019 and to enable better engagement between NSA and the Lablets. The Quarterly also addressed the history and future of the five Hard Problems.

Enabling Better Engagement with the Lablets

NSA representatives described NSA Information Assurance Research activities and challenges. Brad Martin first provided a general overview of the Research Directorate and then focused on Information Assurance Research. In addition to investing in foundational research at the six Lablets, he noted efforts to grow the Science of Security community, citing as examples the Best Scientific Cybersecurity Paper Competition, the SoS Virtual Organization, sponsorship of awards at the Intel International Science and Engineering Fair (ISEF), Hot Topics in the Science of Security: Symposium and Bootcamp (HotSoS), and participation in the Special Cyber Operations Research and Engineering (SCORE) Subcommittee. He added that his organization worked with both the National Research Council Forum on Cyber Resilience and JASON. He provided examples of boutique analyses including Behavior (in select domains, performing exhaustive testing and producing evidence that every possible input results in the expected output); Designs in support of Defense (creating evidence-based tools and technique that analyze and verify the properties of standards, specifications, and systems to eliminate exploitable vulnerabilities); and Systems in support of Defense (allowing for the detection of concerns or assurance of designed properties through evidence-based approaches resulting in counter-examples or proofs). Specific areas in boutique analysis include scalability and composability, in policy and protocol work, as well as finding vulnerabilities in systems. With respect to Centaur-styled analyses, he cited Continuous Automated Patching (creating automated tools that can create a patch and apply the patch, all in real time); Partitioning of Remediated Vulnerabilities (applying reasoning capabilities that can determine the likelihood that a particular vulnerability will be discovered); and Automated Discovery of Vulnerabilities (extending vulnerability discovery tools to help secure government systems).

PI Bill Scherlis welcomes attendees to the Fall Quarterly meeting.

A representative from the Adaptive Cyber Defense Systems office addressed autonomous cyber defense. He provided a high-level definition of autonomous cyber defense as applying automation to achieve needed speed and scale in cyber defense. The following research is under way in this area:

- Research to develop an autonomous cyber defense system that reasons and responds to mitigate the effects of advanced cyber attacks at machine speed and enterprise scale.
- Research to extend Artificial General Intelligence (AGI) reasoning prototype to the cyber domain; focus on simultaneous reasoning and learning despite insufficient knowledge and resources.
- Research to apply reinforcement learning (RL) for autonomous decision-making; the goal is to train an autonomous system to make complicated, sequential decisions under conditions of uncertainty.

He further addressed work being done in adaptive data fusion including enabling adaptive data fusion and situational understanding of events from diverse and disparate data sources; developing new approaches to enable adaptive fusion analytics that address practical constraints with current Computer Network Defense (CND) systems; and quantifying benefits of adaptive analytics to balance with costs of gathering more information. He discussed applying results from Human Subjects Research to ongoing work in order to develop adaptive reasoning and response for autonomous cyber defense, noting that the work is being shared with cyber operators, partners, and researchers. He cited analysis of data from a rigorous experiment conducted using skilled industry red-teams within a fully instrumented environment to determine the effectiveness of deception (network decoy system) on cyber adversaries. The extensive data collected will be analyzed in 2019 and 2020 to quantify the effectiveness of decoy systems at delaying or detecting cyber attack, deception effectiveness with foreknowledge, effectiveness of false information about network deception, and persistence of effect on users following network deception. He also examined other areas of related research including Adversarial Machine Learning, Supply Chain, IoT, and Secure Wearable Authentication.

Eric Clemons, NSA, briefed on cybersecurity challenges from the operations perspective and noted that he wanted to have operations more involved in Lablet activities. He provided an overview of NSA cybersecurity partnerships with academia which are coordinated through the Industry and Academic Engagement office (IAE). These initiatives include academic liaisons with 250 colleges and universities, although he noted that only three of the Lablets (UIUC, NCSU and CMU) currently have academic liaisons. The goal of the academic liaison program is to build and foster mutually beneficial relationships with academia that address critical disciplines of interest, and he wants to identify NSA academic liaisons for the remaining Lablets and expand the current Lablet projects to leverage campus expertise in other fields such as business, social sciences, law, etc. Dr. Clemons volunteered to serve as the temporary academic liaison for the other three Lablets that lack current representation. Another IAE program, the NSA Codebreaker Challenge, also offers an opportunity for Lablet engagement. The Challenge consists of a series of cybersecurity tasks that are worth a varying amount of points based upon their difficulty, and he proposed using the challenge approach as a template for future SoS Lablet partnership. Under the current model, the Lablets work on Research Directorate Hard Problems, while he envisages a future model under which operations engages with research to identify real world problems for the Lablets to tackle. Dr. Clemons briefly addressed the DoD Cyber Strategy and its goals before returning to how the Lablets can assist NSA cybersecurity operations. Dr. Clemons urged the Lablets to engage the IAE office, understand the connection between the DoD Cyber Strategy and academia's role in promoting cybersecurity hygiene throughout U.S. college and university networks, and leverage current (and establish new) research relationships with other NSA academic partners. The discussion that followed his presentation touched on the issues of the long-term horizon associated with research vice operations' short-term horizon, the need for early wins to both satisfy operational requirements and validate academic research, and the potential for seeing that science makes a difference in operations.

New Lablet Research Projects

Project: Secure Native Binary Executions
Lablet: University of Kansas
PI: Prasad Kulkarni
Hard Problem: Resilient Architectures

This goal of this research is to build tools and techniques that will allow users to know the security level of their packaged software and enable them to add security to it. In the current environment, if a user wants to run binary software, he doesn't know how secure the software is or how it incorporates language/compiler level security checks, and there is no way for the user to update the software to enforce the desired level of security against attacks. In providing some background about the problem, Professor Kulkarni addressed memory errors, noting that they are one of the oldest issues in computer security, and then showed how memory errors are exploited. He talked about the characteristics of static binary rewriters and Dynamic Binary Translators (DBT), and said that they will initially use DBT to increase usability and flexibility. He addressed several issues with DBT performance, concluding that performance loss can be significant, and while translation cost is a major factor for latency, it is not significant for throughput. He is proposing Eager Translation to reduce translation cost using the premise of translating code ahead of time in multiple translation threads concurrently. One of the major challenges with this approach is predicting the blocks to translate eagerly. Other questions associated with the use of DBT include identifying what other factors affect DBT latency and throughput and how best to build a security API. In order to enhance the security of software binaries, he discussed several techniques to secure the binary against various security threats, and defined two goals to that end: 1) adapt existing security techniques to DBTs; and 2) assess the security level of source binary. He summarized by noting that in order to meet the research goal of developing a high-performance framework for client-side security assessment and enforcement for COTS binary software, the researchers will build a high-performance DBT platform, build a security API in the DBT to facilitate exploration of security techniques at runtime, build techniques to provide the user with best runtime security within a given performance budget, and build techniques to assess security level of incoming COTS binary.

Project: Resilient Control of Cyber-Physical Systems with Distributed Learning
Lablet: University of Illinois at Urbana-Champaign (UIUC)
Sub-Lablet: University of Texas at Austin (UT Austin)
PI: Sayan Mitra
Co-PIs: Geir Dullerud, UIUC;
Sanjay Shakkottai, UT Austin
Hard Problem: Resilient Architectures

Professor Mitra said that this research was motivated by the complex interaction of dynamics and decision making, noting that the integration of hundreds of components expose Cyber-Physical Systems (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. He pointed out that CPS models are often complex and not analytical, proprietary, and use machine learning. In going beyond model-based approaches, the researchers will exploit models when it makes sense but otherwise get (gracefully degrading) guarantees with blackbox executables, and will focus on the marginal benefits of models/executable fidelity in security and risk reduction; he provided an example of how the approach would work. He presented analyses beyond model-based approaches, including stochastic analysis, blackbox-whitebox analysis, and learning-based optimization/analysis, describing the three types of analyses as the researchers' best-guess approaches. He addressed each of these analyses in more detail, providing examples of applications for the approach, experimental findings, and advantages and potential shortfalls. The plan for Year 1 is to develop tools for online or offline detection and analysis of CPS with model fidelity sensitivity. Questions to be addressed include the cost (sample complexity) of answering a verification/monitoring query while learning the model from coarse simulations, and whether the extra cost associated with incremental analysis is worth the resiliency gained and risk reduced. Applications for this approach include engine control systems, autonomous vehicles, and spacecraft control systems (for which they have ongoing collaborations with AFRL and Boeing). The techniques to be used are data-driven verification, robust control of stochastic hybrid systems, smoothed analysis, and blackbox optimization.

Project: Mixed Initiative and Collaborative Learning in Adversarial Environments
Lablet: Vanderbilt University
Sub-Lablet: University of California, Berkeley
PI: Claire Tomlin

Co-PI: Shankar Sastry
Hard Problem: Human Behavior

This research project focuses on a game theoretic approach to learning dynamic behavior safely. Professor Tomlin addressed learning dynamic behavior safely through reachable sets, probabilistically safe planning around people, and safe policy gradient reinforcement learning. In addressing reachability as a game, she discussed disturbance (attempts to force system into unsafe region) and control (attempts to stay safe), and provided three examples using drones demonstrating collision avoidance, fast and safe planning, and safe policy gradient reinforcement learning. She also addressed fundamental issues with gradient play in games, noting that machine learning algorithms are increasingly being implemented in competitive settings, but an understanding of the behaviors (convergence, optimality, etc.) of these algorithms in such settings is sorely lacking. One of the goals of the research is to characterize the limiting behavior of machine learning algorithms deployed in competitive settings, and she provided the methodology to address this issue. The key takeaways from the results were that gradient-based learning algorithms will: almost surely escape saddle points; converge almost surely to one of finitely many differential Nash equilibria in potential games; and for generic vector fields, converge to one of finitely man linearly stable periodic orbits. As far as conclusions and proposed work, she noted that with respect to learning dynamic behavior, and specifically adversarial behavior, worst case behavior is in perception modules and how it affects dynamics. The fundamental issues with gradient play revolve around whether gradient descent is the right backbone for learning algorithms in games, and open questions include whether there are classes of algorithms that preclude limit cycles in their dynamics and how fast machine learning algorithms can escape saddles in competitive settings.

Project: Reasoning about Accidental and Malicious Misuse via Formal Models of User Expectations and Software Systems

Lablet: NCSUPI: Munindar Singh
Co-PIs: William Enck, Laurie Williams
Hard Problem: Policy-Governed Secure Collaboration

The goal of this research project, which focuses on mobile apps, is 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 implementation to identify misuses sensitive to usage and machine context. Professor Enck began his presentation by talking about early work analyzing mobile apps using a manual application, TaintDroid, and then looked at how and whether mobile app analysis has improved since then. He suggested that there are lots of static and dynamic program analysis tools and that test input generation has significantly improved, but that there has not been much progress on user expectations. He explored expectation context, noting that existing systems make security decisions using context, including temporal context, environment context, execution context, but lack "expectation context," and proposed the following research questions: What would expectation context look like? Where would it come from? How accurate can it be? Where can it be used? He provided an expectation example and then described the research approach as using a decision engine based on norm logic to supplement existing analysis tools and norm types to inform expected behavior. Challenges in this approach include how to identify the norms and acquire the context. After providing an example of expectation context from App GUIs and describing associated issues, he highlighted the findings including the types of most frequent requests as well as surprising requests including third party passwords and sensitive information which are often disclosed to advertisers. He described preliminary efforts in case studies on payment in mobile apps and misuse case discovery. He concluded his presentation by reiterating the goal of the project as helping security analysts seeking both to avoid vulnerabilities and identity malicious functionality, and noting that the research is focused on logical reasoning, norm/misuse identification and extraction, and context extractions.

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

Lablet: CMU
PI: Laurie Cranor
Hard Problem: Human Behavior

The Security Behavior Observatory (SBO) project continues work that has been underway at CMU since 2014. Hana Habib focused her presentation on "Away from Prying Eyes: Analyzing Usage and Understanding of Private Browsing," a paper published at SOUPS 18 which was based on data collected via the SBO. (See https://www.usenix.org/conference/soups2018/presentation/habib-prying). She noted that while users have private browsing modes to prevent collection of their browsing history, many have misconceptions about private browsing thinking it prevents all tracking from websites and advertisers, and that it prevents search engines from recording searches, both of which are untrue and may put users at risk. CMU researchers combined empirical data with a survey to address what private browsing looks like in the wild and how reported browsing behavior relates to observed browsing behavior. She provided some background on the SBO, noting that has instrumented (with permission) home Windows users and has more than 500 participants of which over 200 are active, allowing for empirical observation and a scientific analysis of behavior. For the paper, they conducted an empirical analysis of browsing behaviors collected over 3 years and compared normal browsing and private browsing activities. They then followed up with a survey to SBO participants and MTurk to identify common use cases for private browsing and address why privacy browsing is used and users' understanding of private browsing. The findings showed that private browsing is typically interspersed with normal browsing and that it was used for multiple types of searches that contained a higher proportion of sensitive activities. The survey revealed that while SBO users reported using private browsing, such activity was not observed on the SBO; SBO users reported using private browsing on other devices including mobile. Overall, empirical data was consistent with reported data, and participants provided a number of reasons for using private browsing, although there were major misconceptions about anonymity and cookies. Based on the study, researchers developed a number of design recommendations. They concluded that private browsing is a small percentage of all activities, is used for practical purposes as well as security, and participants had many misconceptions. Plans for future SBO work include studying day-to-day privacy and security behaviors to determine if behaviors match stated levels of privacy and security concerns and whether this information can be used to design protective measures that don't overwhelm users. Another area for future study will be to evaluate privacy information-seeking behaviors including

what kinds of privacy and security information users seek online and what motivates them to seek such information.

Project: Scalable Privacy Analysis

Lablet: ICSI
PI: Serge Egelman
Hard Problem: Policy-Governed Secure Collaboration/Privacy Emphasis

Professor Egelman presented the new project which is focused on the automatic evaluation of the privacy behaviors of Android apps using tools (Instrumented Android and Haystack) to monitor data flows. Researchers are primarily interested in data that's already on the phone and protected by the Android permissions system, specifically personal information and persistent identifiers. They are applying the research to examining compliance with the Children's Online Privacy Protection Act (COPPA) which regulates how mobile apps, games, and websites are allowed to collect and process personal information from children under the age of 13. The Act prohibits certain data collection processes and requires parental permission for others. Their research has found that of the 5855 free "designed for families" apps, 57% are in potential violation of COPPA. Research using the tools has found that potential violations often arise from third-party services included with apps (the SDKs that the developers use). The potential violations persist as a result of platform providers not enforcing their own terms. Professor Egelman reported that with respect to ads, 75% of apps transmit the ad ID along with a non-resettable identifier, a practice prohibited by COPPA. When the research team reported this fact to Google, he said that Google was uninterested. Additionally, 19% share identifiers or personal information, and he noted that companies owned by Google are doing this in violation of Google's own policies. He presented an example of a specific app that researchers found to be in violation of COPPA, and, following a New York Times article on the violations that was based on Professor Egelman's research, Google removed the app from its Play store. He discussed current work which is dealing with uncovering deceptive practices, vulnerability disclosure, and pluggable modules to deobfuscate. With respect to next steps, he addressed plans to study compliance with the European Union (EU) General Data Protection Regulation (GDPR) including whether third parties are listed in privacy policies, whether apps behave differently in the EU, and whether companies honor access and deletion requests. Researchers are also going to study paid versus free apps, looking at the popular assumption that paid apps provide greater privacy—preliminary findings suggest this is not true with the only real difference being paid apps don't have ads. They will perform longitudinal studies focused on before and after GDPR, enforcement actions, and general privacy trends.

Laurie Williams presents her research

Project: Development of Methodology Guidelines for Security Research

Lablet: NCSU
Sub-Lablet: University of Alabama
PI: Jeff Carver, University of Alabama
Co-PI: Laurie Williams
Hard Problem: N/A

The goal of this research is to aid the security research community in conducting and reporting methodologically sound science through the development, refinement, and use of community-based research guidelines and the characterization of the security literature based upon those guidelines. Professor Williams noted that this research began several years ago, so the effort has actually been going on for some time. The first of the research questions being addressed in the project deals with community values, specifically, what elements of methodologically sound research practice and reporting do members of the SOS Lablets and the larger security community identify and value? She reported that results of their research had been received with a wide range of responses: from "okay" to vehement dislike. She further reported that calling the research a case study elicited negative responses, and there was a discussion about how the study could be otherwise characterized. Another question deals with research types and asks how methodologically sound research differs across research types, with research types being identified as theory, algorithm, empirical, and proof. With respect to guidelines, the researchers are looking at to what extent papers from the Lablets and the larger security community conform to research guidelines reflecting community values modulated by research types and how that changes over time. They are also looking at the Lablet papers asking to what extent Lablet researchers differ from the larger research community in terms of their papers conforming to the community-based research guidelines. In addressing the methodology evaluation, they ask what impact structured training and guidance on research methodology may have for Lablet approaches to performing scientific security research. As background she addressed the analysis of indicators in scientific research as reported in the 2015 ACM CCS and IEEE S&P proceedings (presented as "Characterizing Scientific Reporting in Security Literature: An analysis of ACM CCS and IEEE S&P Papers" at HotSoS 2017; see https://dl.acm.org/citation.cfm?id=30355305&picked=prox) which had the goal of aiding security researchers in establishing a baseline of the state of scientific reporting in security as found in that material. She described the research questions the review sought to answer, an overview of the study, the methodology, and the results. The research plan for this project includes community engagement, guideline development, pre-publication manuscript feedback, and assessment of publications.

Hard Problems

Bill Scherlis offered a personal perspective on the seven-year history of the Science of Security project and the five Hard Problems. He noted that the selection criteria for the Hard Problems were: a high level of technical challenge; significant operational value; the likelihood of benefitting from emphasis on scientific research methods and improved measurement capabilities; and potential to identify synergetic common features. The focus on methods is reflected in the "science" of security and privacy, particularly with respect to metrics and human behavior. He said that he believes the framework has been useful, and the emphasis on models has given us a way to assess our research. Less successful aspects have been a failure to successfully connect with other research and operational communities, including the larger cybersecurity community, and initially not enough connection with mission needs. He suggested potential new Hard Problems relating to AI engineering and Machine Learning, CPS and IoT, and the Cloud.

Following Professor Scherlis' presentation, Adam Tagert moderated a Lablet panel discussion on the "History and Future of the 5 Hard Problems," which addressed Hard Problems coverage, relevance, the role of privacy, and completeness. There was a lively debate among the panelists and the audience. Some of the key points were as follows:

- The Hard Problems were not designed to be comprehensive; the creation of the Hard Problems helped the PIs and the Lablets come together as a group.
- As an organizing structure, the Hard Problems have been a success, but not significantly as a measure of progress toward a goal. Hard Problems are never solvable, but the framework is useful. The Hard Problems are a framework, not a finish line.Resilience may be an example of where the framework is driving progress; not so with metrics.
- Metrics is a Hard Problem, but it is also inside all of the others. The ability to measure the consequences of an action is important; measuring vulnerability complexity and discovery of vulnerabilities would be a good metric. Perhaps we should think about "degree of confidence" as a metric.
- AI and ML were not part of the Hard Problems when they were identified, but have gained in importance since then--do we need to have new Hard Problems or revise definitions? AI and ML are found within multiple Hard Problems.
- Since projects often cut across Hard Problems, does that mean the Hard Problems aren't aligned well with the research? While it may be a plus to cut across Hard Problems, does that make it difficult to talk about successes?
- Since the Hard Problem titles still seem relevant, pthe definitions need to be revised.
- Since PIs agreed early on that they wouldn't try to cover everything in identifying Hard Problems, the current ones should continue.
- Perhaps societal impact emphasis is lacking, so we may want to emphasize outreach activities.
- Given the huge body of results from the program, we need to make sure the community is aware of some of the results, particularly when they have an impact on mission, in a context where it really matters.

Section 2

Promoting Rigorous Scientific Principles

The 2018 NSA Research Directorate ISEF Winners along with Dr. Adam Tagert, SoS Technical Lead, and Dr. William Christian, Informatics Research Deputy Chief

Through their sponsorship of fundamental research at the six Lablets as described in Section 1 of this report, the Science of Security and Privacy Initiative (SoS) ensures a focus on promoting rigorous scientific principles. Over the past several years, the SoS initiative has used two other means to promote rigorous scientific principles, specifically, the Annual Best Scientific Cybersecurity Paper Competition and sponsorship of awards at the Intel International Science and Engineering Fair (ISEF). This year they added another initiative: sponsorship of two additional Best Paper Awards, one at the Hot Topics in the Science of Security: Symposium and Bootcamp (HotSoS), and the other at the Symposium on Usable Privacy and Security (SOUPS).

For the 6th Annual Paper Competition, there were 28 submissions which brought the total number of submissions to more than 225 during the six years of the competition. This year's winning paper was "How Shall We Play a Game? A Game-theoretical Model for Cyber-warfare Games" by researchers from Carnegie Mellon University (CMU) and University of California, Santa Barbara. The competition also included two papers that addressed the philosophical question "What is a science of security?", and the distinguished experts noted the authors' work in helping to shape and mature the science of security discipline which, in turn, promotes rigorous scientific principles.

This was the fourth consecutive year that SoS has sponsored prizes at ISEF, and this year there were seven winners in the SoS category: a First Place, three Second Places, and three Honorable Mentions. The Research Directorate presented other awards bringing the total number of awards to 23, and the winners were invited to NSA to present their research. Based on their research, two of the ISEF winning students were invited to attend the Computational Cybersecurity in a Compromised Environment (C3E) Workshop in Atlanta, Ga in September which brought together top academic, commercial, and government experts to examine new ways of approaching national cybersecurity challenges based on rigorous scientific principles.

The winners of the HotSoS and SOUPS Best Paper competitions were selected by the respective program committees based on evaluation criteria that SoS personnel developed with them. The winning papers were automatically submitted to the 7th Annual Best Scientific Cybersecurity Paper Competition.
Details on the Paper Competitions and the Intel ISEF can be found in the following pages.

Best Scientific Cybersecurity Paper Competition

The NSA sponsors an Annual Best Scientific Cybersecurity Paper Competition established in 2012 to promote and encourage rigorous research methodology in cybersecurity. Every December NSA invites nominations of papers published within the year that show outstanding contribution to cybersecurity science. Eligible papers must be published in peer-reviewed journals, magazines, or technical conferences and may come from any field of cybersecurity research. Nominations describe the scientific contribution of the paper and explain why a paper merits the award. Nominators may not be an author or co-author of the nominated paper. Strong nomination statements are used as part of the criteria for evaluating paper submissions. The papers are reviewed by a set of Distinguished Experts on the basis of scientific merit, significance of the work reported, and the degree to which the paper exemplifies how to perform and report scientific research in cybersecurity. The Distinguished Experts are leaders from various fields of cybersecurity research who volunteer their time to the competition. They provide individual assessments to the NSA Research Directorate. The NSA Director of Research selects the winning paper concluding the review process. The winning paper author(s) are invited to NSA to be awarded the Best Scientific Cybersecurity Paper and to present their research to the intelligence community.

In 2018 there were 28 submissions, bringing the total number of submissions to more than 225 during the six years of the competition. Papers nominated span the breadth of the cybersecurity field. This year the paper competition itself made an appearance in two of the papers nominated. These papers discussed the philosophical concept of a security science and used the paper competition as an example. The authors of these two papers were invited to the 2019 Hot Topics in the Science of Security: Symposium and Bootcamp (HotSoS) to discuss their perspective on science.

The following individuals served as Distinguished Experts for the 6th annual competition:

- Professor L. Jean Camp, Indiana University
- Dr. Robert Cunningham, Lincoln Laboratory
- Dr. Whitfield Diffie, Cybersecurity Advisor
- Dr. Dan Geer, In-Q-Tel
- Dr. John McLean, Naval Research Laboratory
- Professor Stefan Savage, University of California, San Diego
- Mr. Phil Venables, Goldman Sachs
- Professor David Wagner, University of California, Berkeley
- Dr. Jeannette Wing, Columbia University

The winner of the 6th Annual Best Scientific Cybersecurity Paper Competition was **"How Shall We Play a Game? A Game-theoretical Model for Cyber-warfare Games"** by Tiffany Bao, Yan Shoshitaishvili, Ruoyu Wang, Christopher Kruegel, Giovanni Vigna, and David Brumley. These researchers were at Carnegie Mellon University and University of California, Santa Barbara. (See https://ieeexplore.ieee.org/document/8049648). This paper was originally accepted at the 30th IEEE Computer Security Foundations Symposium (CSF '17).

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

From left: Mr. Phil Venables, Distinguished Expert; Dr. Deborah Frincke, NSA Research Director; Dr. Jean Camp, Distinguished Expert

The paper was selected because it exemplifies outstanding scientific research, is technically sound, and is well written. The authors developed a cyber-warfare strategy based on strong scientific methods, and this new approach performs better than what was previously known. The reviewers particularly liked that the game theoretic model was reflective of the physical world with a realistic set of assumption and attributes, which is refreshing to see in game theory papers. The paper is noteworthy in the validation effort to test the effectiveness of the game theory strategy. The team applied their game theory strategy to the 3rd place finisher at the DARPA Cyber Grand Challenge. Validation of research with real world situations is important in science and helps build confidence in that results apply to real-life situations. The attributes of this paper made it well-deserving of winning the 6th Annual Best Scientific Cybersecurity Paper Competition.

Paper Competition Award Ceremony
Pictured left to right:
Prof. Giovanni Vigna, Prof. David Brumley, Prof. Tiffany Bao, Dr. Deborah Frincke, NSA Research Director, Prof. Yan Shoshitaishvili, Prof. Ryoyu Wang

Awards Ceremony

The 6th Annual Best Scientific Cybersecurity Paper Competition Awards Ceremony honoring the winning authors was hosted by the Research Directorate at NSA on 30 November 2018. Five of the six authors and two of the Distinguished Expert reviewers attended.

Dr. George Coker, Chief, Information Assurance Research, welcomed the attendees noting that the Best Scientific Cybersecurity Paper Competition is the culminating event for the SoS program. This year's winning paper, he said, has provided new insights in cyber operations with new approaches for cyber strategies.

Dr. Deborah Frincke, NSA Research Director, addressed progress in cybersecurity science, noting that it is a field in which some of the learning comes by doing. She pointed out that investing in cybersecurity science is one way of demonstrating its importance, but that the science of security is still not recognized as much as it needs to be. What set this paper apart, she said, was its emphasis on scientific method. The authors investigated a model without trying to put too much into it—it was neither too simplistic nor too complex. Another positive attribute was that it asked key questions about cyber warfare and made reasonable assumptions about the number of parties involved.

Following presentations by Dr. Frinke to the five authors and the two Distinguished Experts, Professor Bao gave a presentation on their research.

Professor Bao presents the research to the attendees

Professor Bao began by describing the case study for their paper, specifically the application of their game theory strategy to the 3rd place finisher at the DARPA Cyber Grand Challenge. The authors (three of whom were members of the 3rd place team) were able to demonstrate that the application of their model would have achieved a better result than the original, concluding that "strategy matters." Noting that NSA has said it discloses 91% of the zero-day vulnerabilities it discovers following a deliberation process, Professor Bao said that because decisions are made in an ad-hoc way, parties need a systematic approach for decision-making since humans are the bottleneck for strategic decision-making. The paper explores the research question of augmenting the human decision-making process with automated techniques rooted in game theory. She discussed the challenge of developing a model which is comprehensive and an algorithm which is efficient, and the tradeoff between the two goals. She contrasted their work with previous work, noting that their model overcomes the limitations of previous work by considering multiple actions over time and uncertainty about other parties by proposing a feasible algorithm for approximating optimal decision-making. Their model is Partial Observation Stochastic Game (POSG), a multi-player, single round game with multiple actions over time (Stochastic) and uncertainty about the opponent (incomplete information game). She provided details on the development of the model, the algorithm (including computing Nash equilibrium), and how the model was applied to the CGG case

study. Their model challenges previous results (at least one player should attack), by showing scenarios where attacking is not optimal for either player. She addressed future work such as identifying and developing strategic techniques, including ricochet attacks, patch-based exploit generation, and patch deployment. She concluded by noting that because their model is more comprehensive and computationally tractable, it helps to reason about zero-day decision-making and the model is a start for systematically investigating zero-day vulnerabilities.

Following Professor Bao's presentation, Dr. Carl Landwehr moderated a Question and Answer session with the awardees. A lively dialogue between the authors and the audience raised a variety of issues including whether the authors considered consequences in domains other than cyber and differences among players' patching techniques and capabilities that might affect timing of exploits. Another topic was how this model would apply to war-gaming exercises and whether the model could shape rules for future war-gaming exercises. In addressing different values of vulnerabilities or their severity, speakers noted that being able to assign more quantitative values could help answer some of the other questions raised by the model. Dr. Landwehr noted that one of the things he liked about the paper was that it provided a much more applied version of game theory. One of the most interesting exchanges came in response to the question: "What did you wish you knew?" Professor Bao acknowledged wishing the authors had more information from NSA about zero-day vulnerability discoveries and responses, specifically how people decide, what the process looked like, and whether it was manual or automated. Professor Brumley talked about wanting information on different strategies and timing for disclosing zero-day vulnerabilities, and Professor Vigna said that understanding the adversarial model from an NSA perspective would have been desirable. Professor Bao summarized their views by saying that they would like to know if NSA agreed with the model—they'd like to be corrected if they're wrong. The ceremony capped a successful 6th paper competition. The 7th Annual Best Scientific Cybersecurity Paper Competition opened for nominations on December 15, 2018 for papers published in 2018.

Professors Bao and Wang answer questions from the audience

Notable Philosophy of Science Papers

Two papers were nominated for the paper competition that addressed the philosophical question "What is a science of security?" The Distinguished Experts noted the authors' work in helping to shape and mature the science of security discipline and wanted to recognize the contribution of this work to the SoS community even if the papers are generally not the type of contribution the competition recognizes. As such, the authors were invited to further discuss their perspectives at the HotSoS meeting to be held in April 2019. The first invited paper "SoK: Science, Security and the Elusive Goal of Security as a Scientific Pursuit" by Cormac Herley and Paul van Oorschot, examines what has been done in science of security and puts it in context with historical science to offer observations and insights. The authors propose eleven constructive suggestions on how the discipline can improve and learn from the development of other disciplines. This paper was originally published in 2017 IEEE Symposium on Security and Privacy. The second paper, "Practicing a Science of Security: A Philosophy of Science Perspective" by Jonathan Spring, Tyler Moore and David Pym was published at the 2017 New Security Paradigms Workshop. The authors examined purported serious obstacles to the practice of a science of security and found that they are either misguided or can be overcome.

Intel International Science and Engineering Fair

NSA Special Award Winners Recognition Ceremony

Led by the Science of Security and Privacy (SoS) initiative, the NSA Research Directorate participated as a Special Award Organization to sponsor awards at the Intel International Science and Engineering Fair (ISEF) held in Pittsburgh, Pennsylvania from May 13-18, 2018. Intel ISEF is a program under the Society for Science and the Public (the Society) with roots dating back to World War II, while the Society itself has been in existence for almost 100 years. At this competition, more than 1700 high school students from 81 countries earned finalist status as a result of their performance at a local, regional, or national level ISEF-affiliated fair. Participants are judged on various criteria to include skill set, creative ability, presentation clarity and scientific thought. The finalists this year competed for nearly $5 million in awards and scholarships.

For the fourth consecutive year, SoS sponsored awards at Intel ISEF in the field of Science of Security; this was the second year for the award in the field of Mathematics. The SoS initiative is dedicated to developing a scientific discipline focused on cyber security. The Research Directorate utilizes ISEF to spark student interest in research to protect cyberspace interactions in an increasingly interconnected world. These awards were created to encourage high school students to pursue scientific research in cybersecurity and related fields. This year, the NSA Research Directorate awarded a total of $13,000 to 23 ISEF participants, more than doubling the number of awards and the amount presented in 2017. In addition, new awards were sponsored in the following categories: Material Science, Cyber Pioneer, Physical Science, and Future of Computing. The SoS initiative will again award prizes at ISEF 2019 in Phoenix, AZ.

Furthermore, the NSA Research Directorate partnered with the Office of Director of National Intelligence (ODNI) to host an Intelligence Community (IC) booth at the ISEF exhibit hall. This exhibition allows NSA researchers an opportunity to personally interact with thousands of students, parents, and other attendees to inspire interest in cybersecurity and STEM. At the booth, NSA researchers had an authentic ENIGMA machine that was a huge success and led to a steady stream of eager students and curious attendees. The NSA researchers educated participants on career opportunities that exist within the Agency and the cyber security profession as a whole. For the first time, researchers also participated in three symposia sessions with FBI and ODNI to discuss the vast array of STEM careers available in the IC for students and attendees interested in employment. Moreover, Dr. William Christian, Deputy Chief of NSA's Informatics Research Office, gave an inspirational symposia on Sense-Making and Critical thinking. His presentation discussed challenges associated with answering difficult questions in a world of overflowing data – and the kinds of answers one might prefer.

NSA researchers showcase authentic WWII ENIGMA machine to eager students

The journey to find the NSA Research Directorate Award winners was a long adventure for the NSA judges. These individuals were tasked to review over 1500 presentations (papers, posters, and supplemental information) to determine if a project demonstrated excellence while being relevant to the NSA mission. After an initial review, the judges narrowed the top projects down to 50 and awarded each with a NSA pin. To select the top projects, the NSA judges conducted student interviews and reexamined all the student materials. Eighteen individual projects and two group projects were selected as the 2018 NSA Research Directorate Award winners.

NSA Research Directorate Awards

Science of Security

- 1st Place: Vivek Bhupatiraju of Lexington MA - "accAAD: Efficient Append-Only Authenticated Dictionary for Transparency Logs"
- 2nd Place: Deepti Vaidyanathan of Baton Rouge, LA - "Using Two-Mode Squeezing for Room-Temperature Photon-Number-Resolving Detection"
- 2nd Place: Suha Hussain of Ozone Park, NY - "A New Method for the Exploitation of Speech Recognition Systems"
- 2nd Place: Divya Amirtharaj of Portland, OR - "Utilizing Blockchain to Revolutionize Privacy and Security of Medical Records"
- Honorable Mention: Nicole Meister of Ellicott City, MD - Honorable Mention: "Improving Robustness of X-Ray Synchotron Image Analysis Using Deep Learning and Data Augmentation"
- Honorable Mention: Daniel Santiago of Anasco, Puerto Rico - "On the Validity of Composite Logical Functions"
- Honorable Mention: Shrya Pingali of Salt Lake City, UT - "Using Machine Learning to Optimize Key-Length Prediction for Polyalphabetically Encrypted Text"

Mathematics

- 1st Place: Bryan Gopal of Chandler, AZ - "A Novel Accelerator for Machine Learning Algorithms"
- 2nd Place: Franklyn Wang of Falls Church, VA - "Monodromy Groups of Indecomposable Rational Functions"
- 2nd Place: Jim Kong of Owings MD, William Longsworth, and Nathan Hayes of Dunkirk MD – "Iago: A Study of Neural Networks, Othello, Difficulty, and Intelligence"
- Honorable Mention: Isha Puri of Chappaqua, NY - "A Scalable and Freely Accessible Machine Learning Based Application for the Early Detection of Dyslexia"
- Honorable Mention: Emil Geisler of Bountiful, UT - "Combinatorics on Path Connections of a Rectangular Graph"

Material Science

- 1st Place: Kevin Meng of Plano, TX - "Vehicle Action Prediction with Artificial Intelligence: An Innovative Way to Transform Advanced Driver Assistance Systems from Reactive to Proactive"
- Honorable Mention: Jacob Wu of Princeton, NJ - "Spray and Stick: A Novel Agent for Pesticide Adhesion"

Cyber Pioneer

- 1st Place: Eshan Chhabra of Plano, TX - "Untapped Static: A New Paradigm for Energy Harvesting Integrating a Cost-Effective Electrostatic-Based Generator with Supercapacitors to Optimize Energy Storage and Energy Harvesting Efficiency"
- Honorable Mention: Aditya Singh of Ponte Vedra, FL - "Edge Detection in the Line of Sight"

Physical Science

- 1st Place: Sharmi Shah of Colonia, NJ - "Speech Intelligibility Analysis of Sound-Modulated Laser Signal Countermeasures"

- Honorable Mention: Carissa Wu of Potomac, MD and Abhishek Allamsetty of Herndon, VA - "Procedural Determination of Novel Stoichiometric Topological Superconductors Through Surface and Pressure Effects"

Future of Computing

- 1st Place: Swagat Bhattacharyya of Morgantown, WV - "DIMOS: A Novel Low-Power, Fast Response Logic Gate Architecture"

- Honorable Mention: Cade Brown of Knoxville, TN - "Software Techniques for Rendering Fractals"

On August 2-3, 2018, the Research Directorate hosted 7 of the Intel ISEF 2018 winners and their chaperones at NSA. The two-day visit included presentations, a poster session, tours with briefings, and meetings with researchers and senior leadership, including Dr. Deborah Frincke, NSA Director of Research, and Mr. Neal Ziring, NSA Technical Director of Capabilities. The highlight of the visit was the research presentations and poster session where NSA researchers and students interacted. Multiple researchers remarked that the students and their research were highly impressive.

A synopsis of the presentations follows:

Vivek Bhupatiraju presented his research entitled "An Efficient Append-Only Authenticated Dictionary for Transparent Public Logs" which centered around designing a bandwidth-efficient public log. The first step was to introduce a novel cryptographic primitive called an append-only authenticated dictionary (AAD), and subsequently include accAAD, an efficient AAD based on bilinear accumulators. Using novel amortization and hash prefix tricks, all of accAAD's cryptographic proofs are polylogarithmic in both computation and bandwidth. This work will improve efficiency of the logs and reduce cost for clients.

Deepti Vaidyanathan's research was entitled "Using Two-Mode Squeezing for Room-Temperature Photon-Number-Resolving Detection" which focused on determining whether two-mode squeezing (TMS) can be used for photon-number-resolving-detection (PNRD) to generate an accurate photon count. She described her experiment and noted that the findings indicate that her proposed PNRD setup would be able to detect different FOC states more than a PNRD setup that did not use TMS.

Suha Hussain presented her research, "A New Method for the Exploitation of Speech Recognition Systems." Her research focused on a way to exploit voice recognition systems by leveraging neural networks, an approach that had not been noted previously. She developed and evaluated an algorithm that demonstrated that neural networks in speech recognition systems are a significant vulnerability and deserve study to determine additional vulnerabilities and possible defenses.

ISEF Student visit to NSA
Pictured left to right: Sharmi Shah (Physical Science), Vivek Bhupatiraju (SoS), Eshan Chhabra (Cyber Pioneer), Dr. Frincke, NSA Research Director, Daniel Santiago (SoS), Shriya Pingali (SoS) Ddepti Vaidyanathan (SoS), and Suba Hussain (SoS)

Conference Distinguished Paper Awards

This year marked the first time that the Science of Security and Privacy (SoS) initiative sponsored Best Paper Awards, one at the Hot Topics in the Science of Security: Symposium and Bootcamp (HotSoS), and the other at the Symposium on Usable Privacy and Security (SOUPS). As is the case for the Best Scientific Cybersecurity Paper Competition, papers were evaluated on the basis of scientific merit, significance of the work reported, and the degree to which the paper exemplifies how to perform and report scientific research in cybersecurity. Generalizability, rigor of research, and clarity of writing are also considered. Both the HotSoS and SOUPS winning papers received an automatic nomination into the 7th Annual Best Scientific Cybersecurity Paper Competition. The HotSoS and SOUPS program committees selected the winning papers based on evaluation criteria that SoS personnel developed with them.

The Hot Topics in Science of Security (HoTSoS) Best Paper Award recognizes the paper that exhibits outstanding achievement in science. Papers are selected by the HoTSoS Program Committee. The HotSoS 2018 winning paper was "Robustness of Deep Autoencoder in Intrusion Detection under Adversarial Contamination" by Pooria Madani and Natalija Vlajic of University of York. URL: https://cps-vo.org/hotsos18/madani. The researchers examined how to make the machine learning algorithm more robust to adversarial attacks.

Dr. Adam Tagert, SoS Technical Director, presents Pooria Madani with the HotSoS 2018 Best Paper award

The Symposium on Usable Privacy and Security (SOUPS) Distinguished Paper Award was selected by the conference organizers and included a cash prize. The SOUPS 2018 winning paper was "I've Got Nothing to Lose": Consumers' Risk Perceptions and Protective Actions after the Equifax Data Breach by Yixin Zou, Abraham H. Mhaidli, Austin McCall, and Florian Schaub from the School of Information, University of Michigan. URL: https://www.usenix.org/conference/soups2018/presentation/zou. The winning researchers developed mental models to understand how consumers perceived the Equifax data breach and their response behaviors. The models can be used to help craft improved communications strategies to improve consumer responses to data breaches.

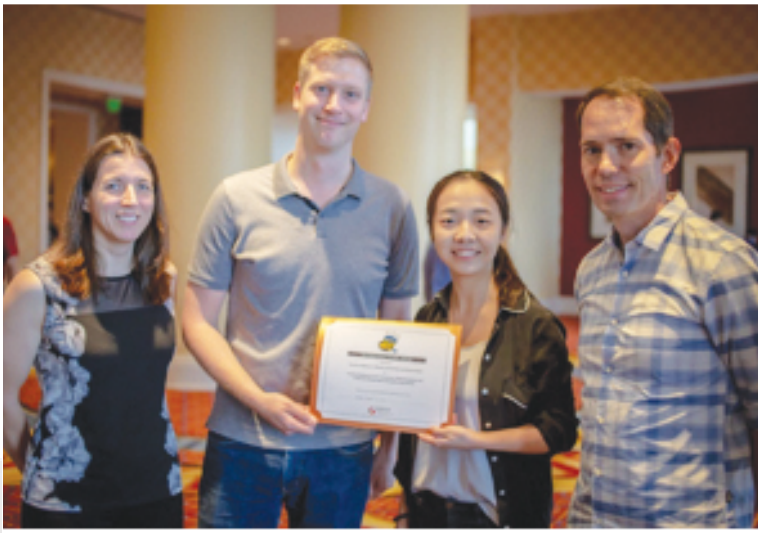

Authors Florian Schaub and Yixin Zou receive the SoS Best Paper Award at SOUPS 2018

Introduction to the National Security Agency's Annual Paper Competition

Adam Tagert, Ph.D. – National Security Agency – actager@tycho.ncsc.mil

SCIENCE OF SECURITY AND PRIVACY

The Competition

Goal

NSA intends the competition to be the premier award for recent scientific foundational cybersecurity research. The goal is to promote rigorous research methods by identifying and highlighting excellence.

Origin

2000s – Community discussions on need for a more scientific approach to cybersecurity
2011 – National Security Agency establishes its Science of Security Initiative
2012 – Establishment of Annual Best Scientific Cybersecurity Paper Competition
2018 – Sponsorship of conference paper awards
First is SOUPS '18

Process

Dec – Mar	Papers nominated by public on sos-vo.org
Apr – July	Papers reviewed by NSA and Distinguished Experts
Aug / Sept	Announcement
Oct	Authors visit NSA

Eligibility

- Published in appropriate calendar year
- Published in a peer-reviewed venue
- Nominated by someone other than an author or co-author
- Open to International authors, nominators and venues

Distinguished Experts

L. Jean Camp
Robert Cunningham
Whitfield Diffie
John McLean
Dan Geer
Angela Sasse
Stefan Savage
Phil Venables
David Wagner
Jeanette Wing

What is Science?

Question that the security community is actively discussing.
Paper Competition evaluates on whether:

- Authors discover something about the cyber world using scientific methods,
- Authors indicates how the results do or might extend beyond current case,
- Authors uses and documents a sound methodology in the research,
- Authors expresses assumptions and limitations,
- Result has foundational or systemic impact,
- Authors' writing is clear and well organized

Papers in Usable Security

2013

Winner :
The Science of Guessing: Analyzing an Anonymized Corpus of 70 Million Passwords by Joseph Bonneau

2014

Honorable Mention :
Rethinking SSL Development in an Appified World by Sascha Fahl, Marian Harbach, Henning Perl, Markus Koetter, and Dr. Matthew Smith

2015

Honorable Mention:
Increasing Security Sensitivity with Social Proof: A Large-Scale Experimental Confirmation by Sauvik Das, Dr. Adam D.I. Kramer, Prof. Laura Dabbish and Prof. Jason Hong

2017

Winner:
You Get Where You're Looking For: The Impact of Information Sources on Code Security by Yasemin Acar, Michael Backes, Sascha Fahl, Doowon Kim, Michelle L. Mazurek, Christian Stransky

Highlighting the Winner of 2017

About – the root causes of security vulnerabilities introduced by software developers.

Study – lab experiment of android developers testing the effects of reference materials on functionality and security on developed code.

Results
› Using Stack Overflow produced more functional code but less secure code
› Using official documentation led to more secure code but less functional
› Books were the best balance, but nearly 0 developers freely chose to use a book

Reason for Selection Include:
› Scientifically confirming anecdotal thought
› Work to validate lab result
› Through discussion of limitations
› Applicability of Results

One expert reviewer hailed the paper as, "*one of the greatest pieces of research in our field in recent years. It shows what dramatic improvements can be made when researchers take the effort to trace the root of the problem and come up with a solution that makes it easy for the people involved to choose the secure option within the other demands they face.*"

https://sos-vo.org/papercompetition

Dr. Tagert's poster presented at SOUPS 2018

Growing the Science of Security

The Science of Security and Privacy initiative sponsored the 2018 Hot Topics in the Science of Security: Symposium and Bootcamp (HotSoS 2018). HotSoS was held in cooperation with the Association for Computer Machinery (ACM), and hosted over 120 participants from government, industry, and academia at North Carolina State University for the two-day workshop in April 2018. In addition to HotSoS, the SoS-VO is a longstanding initiative designed to grow the Science of Security community. The Science of Security Virtual Organization (SoS-VO) now includes 1500 members and continued to provide a centralized location for cybersecurity research, events, and news. New outreach initiatives in 2018 included Conference Distinguished Paper awards, co-sponsorship of a Cyber-Physical Systems Security Summer Camp, advertisements in scholarly journals, and attendance at multiple conferences. As a result of all of these efforts, the SoS community continued to expand. The monthly Science of Security Reviews and Outreach newsletter has over 1500 subscribers and includes the curated publications of over 20,000 researchers. This year also saw a growing SoS presence on social media. Details on HotSoS, and other outreach initiatives are found in the following pages.

HotSoS

Hosted by the North Carolina State University (NCSU) Lablet, the 2018 Hot Topics in the Science of Security: Symposium and Bootcamp (HotSoS) was held April 10 and 11 in Raleigh, NC. This was the sixth time researchers have come together to interact and attend presentations demonstrating rigorous scientific approaches to prevent, detect, and mitigate cyber threats. A major continuing focus of the conference is the advancement of scientific methods in approaching the Hard Problems in cybersecurity. NCSU Lablet Co-Principal Investigators (PIs) Laurie Williams and Munindar Singh welcomed the audience to HotSoS 2018, noting that there were about 120 attendees from government, industry and academia for the two days of presentations in both research and industry tracks. They reported that the number of research paper and poster submissions from 15 universities worldwide demonstrates the growing interest in HotSoS and the growth in collaboration. The symposium and bootcamp included research papers, keynote and invited presentations, industry presentations, and tutorials. A panel discussion and poster sessions rounded out the agenda. Details on the presentations and posters are provided below.

HotSoS 2018 proceedings have been published by ACM and are available online at the ACM Digital Library at https://dl.acm.org/citation.cfm?id=3190619&picked=prox.

Opening Remarks

George Coker, Chief of NSA Information Assurance Research, welcomed the attendees and challenged them to advocate for the science of security. He pointed out that since cybersecurity is the intersection of multiple disciplines, we need to build on the science of those multiple disciplines to build the science of cybersecurity. He noted that the Science of Security initiative has grown to the Science of Security and Privacy initiative with six Lablets, all addressing the five Hard Problems, with one Lablet focused on privacy and another focused on Cyber-Physical Systems.

Research Papers

Nine refereed papers were selected for presentation out of twenty-nine submissions. The paper tracks were organized into three focus areas: Vulnerabilities and Detection, Secure Construction, and Applications and Risk Evaluation.

This was the first year for the Hot Topics in Science of Security (HoTSoS) Best Paper Award which recognizes the paper that exhibits outstanding achievement in science. Papers are selected by the HoTSoS Program Committee. The winning paper receives an automatic nomination into the Annual Best Scientific Paper Competition. The paper "Robustness of Deep Autoencoder in Intrusion Detection under Adversarial Contamination" by Pooria Madani and Natalija Vlajic of University of York received the 2018 HotSoS Best Paper Award.

Vulnerabilities and Detection

1. Robustness of Deep Autoencoder in Intrusion Detection under Adversarial Contamination
Pooria Madani and Natalija Vlajic, University of York

Intrusion Detection Systems (IDSs) generally use some Machine Learning (ML) algorithms. However, a sophisticated adversary could target the learning module of these IDSs in order to circumvent future detections. Consequently, robustness of ML-based IDSs against adversarial manipulation (i.e., poisoning) will be a key factor for the overall success of these systems. The authors presented a novel evaluation framework for performance testing under adversarial contamination, studying the viability of using deep autoencoders in the detection of anomalies in adaptive IDSs and their overall robustness against adversarial poisoning.

2. Understanding the Challenges to Adoption of the Microsoft Elevation of Privilege Game
Inger Anne Tøndel, Norwegian University of Science and Technology, and Tosin Daniel Oyetoyan, Martin Gilje Jaatun, and Daniela S. Cruzes, SINTEF Digital

In Norway, there is very low adoption of threat modeling for software security. This study used a card game, Microsoft Elevation of Privilege (EoP), to make threat modeling more fun and available to developers. The EoP card game helps clarify the details of threat modeling and examines possible threats to software and computer systems. The EoP game focuses on

spoofing, tampering, repudiation, information disclosure, denial of service, and elevation of privilege, and uses a simple point system that allows players to challenge other developers and to become the opponent's biggest threat. Results of the study suggest that using the game has the potential to improve security interest and awareness and may also be useful in training for threat modeling.

3. Reinventing the Privilege Drop: How Principled Preservation of Programmer Intent Would Prevent Security Bugs
Ira Ray Jenkins, Sergey Bratus, and Sean Smith, Dartmouth College, and Maxwell Koo, Narf Industries

The principle of least privilege requires that components of a program have access to only those resources necessary for their proper function. Defining proper function is a difficult task. The authors present the use of their Executable and Linkable Format (ELF)-based access control (ELFbac), a technique for policy definition and enforcement. ELFbac leverages the common programmer's existing mental model of scope and allows for policy definition at the Application Binary Interface (ABI) level.

Secure Construction

4. Secure MR: Secure MapReduce Computation Using Homomorphic Encryption and Program Partitioning
Yao Dong and Ana Milanova, Rensselaer Polytechnic Institute, and Julian Dolby, IBM

As customers upload data and computation to cloud providers, they typically give up data confidentiality. In this study, the speakers describe SecureMR, a system to analyze and transform MapReduce programs to operate over encrypted data. SecureMR uses partially homomorphic encryption and a trusted client. Their "secret sauce" is using a partially homomorphic encryption instead of fully homomorphic which reduces overhead. They evaluated SecureMR on a set of complex computation-intensive MapReduce benchmarks on Google Cloud with good results. In their evaluation, 89% required no conversions.

5. Integrated Instruction Set Randomization and Control Reconfiguration for Securing Cyber-Physical Systems
Bradley Potteiger, Zhenkai Zhang, and Xenofon Koutsoukos, Vanderbilt University

Cyber-Physical Systems (CPS) require proper control reconfiguration mechanisms to prevent a loss of availability in system operation. This presentation addressed the problem of maintaining system and security properties of a CPS under attack by integrating ISR, detection, and recovery capabilities to ensure safe, reliable, and predictable system operation. The authors consider the problem of detecting code injection attacks and reconfiguring the controller in real-time for an autonomous vehicle case.

6. Formal Verification of the W3C Web Authentication Protocol
Iness Ben Guirat, INRIA, and Harry Halpin, INSAT

The formal verification of protocols can set the science of security on firm foundations. The design validation of new protocols in an automated method allows protocol designs to be scientifically compared in a neutral manner. The authors demonstrate how formal verification can be used to analyze new protocols such as the Web Authentication Working Group (W3C) Web Authentication.

Applications and Risk Evaluation

7. Application of Capability-Based Cyber Risk Assessment Methodology to a Space System
Martha McNeil, Thomas Llanso and Dallas Pearson, Johns Hopkins University Applied Physics Laboratory

Cyber threats remain a growing concern that requires stakeholders to perform cyber risk assessments in order to understand potential mission impacts. The authors present an automated, capability-based risk assessment approach, compare it to manual event-based analysis approaches, describe its application to a notional space system ground segment, and describe the results. Their capability-based approach "BluGen" uses mission impact for every combination of mission, asset, data type and effect. Effectiveness is derived from Reference Catalog data, but it is currently immature. Risk plots are reusable. They seek metrics instead of expert systems with the objective of getting humans out of the loop. There are limitations in what they have produced and are testing it now, but they need more and better data.

8. Challenges and Approaches of Performing Canonical Action Research in Software Security
Daniela S. Cruzes, Martin Gilje Jaatun, and Tosin Daniel Oyetoyan, SINTEF Digital

The objective of this research is to develop a research-based model of security engineering for agile software development through science of security. Their methodology was to create software that can withstand a malicious attack by creating work processes to handle security issues in order to assure security will be addressed by the software team. Canonical Action Research (CAR) has well defined principles and is based on a 2004 work by Davison et al in the Information Systems Journal. The principles of this methodology include Researcher-Client agreement, cyclical process model, theory, change through action, and learning through reflection. Challenges to implementation include building trust, data collection, analysis of data, security, use of other theories (not just technical), and metrics to measure success in adding security.

9. Quantifying the Security Effectiveness of Firewalls and DMZs
Huashan Chen and Shouhuai Xu, University of Texas at San Antonio, and Jin-Hee Cho, Army Research Lab

The authors present a framework for investigating the security effectiveness of Firewalls and Demilitarized Zones (DMZs) in protecting enterprise networks. Their objective is to provide a systematic, fine-grained framework for modeling firewalls and DMZs by treating an entire enterprise network as a whole and by treating individual applications and operating system functions as "atomic" entities. They are accommodating realistic, APT-like attacks. Their global view, they assert, allows them to quantify the network-wide effectiveness of replacing one mechanism with an improved mechanism.

Presentations

Keynotes

1. Foundational Cybersecurity Research: Report of a Study by NASEM

Steve Lipner, Executive Director, SAFECode

From 2012-2014, a committee of the National Academies of Science, Engineering, and Medicine (NASEM) conducted a study at the request of NSA Information Assurance Research to look at the science of cybersecurity. Mr. Lipner was one of the members of the committee that included representatives of academia as well as cybersecurity practitioners from industry. The report reviewers were well-known academicians in the field. The study noted that despite investments, significant problems remained. Moreover, old approaches hadn't been adequate for a number of reasons including asymmetry, difficult routes for solution adoption, system complexity, and risk aversion. The study identified four broad aims for cybersecurity research: strengthening the scientific underpinnings of cybersecurity; integrating the social, behavioral, and decision sciences in security science; integrating engineering, operational, and life-cycle challenges in security science; and supporting and sustaining foundational research for security science. The speaker addressed the following institutional challenges and opportunities: demand science of security standards; support joint projects across disciplines; emphasize operational and lifecycle perspective in design and evaluation; and integrate with business cases to support adoption. He concluded by calling for scientific rigor, interdisciplinary approaches, and real-world applications as well as theory.

2. You've Got a Vuln, I've Got a Vuln, Everybody's Got a Vuln

Ari Schwartz, Venable LLP

The speaker addressed vulnerability disclosure policies and how they affect government, academic and private security research. Existing vulnerability standards focus on what vendors do when they receive notification of vulnerabilities, and he advocates for standards for Coordinated Vulnerability Disclosure (CVD) that would go beyond vendors. He noted that the Vulnerabilities Equity Process (VEP) had been reinvigorated following media leaks, and he addressed criticisms of and recommendations to improve the VEP. He noted the following issues for the future of CVD:

- How to adapt CVD
- How to encourage research in the right areas while limiting researcher liabilities
- Whether increased government hacking , with oversight, can make up for lost data from greated end-to-end encryption

3. Cyber Security for Aviation Weapon Systems

David Burke, Technical Director, Naval Air Systems Command (NAVAIR) Cyber Warfare Detachment

Dr. Burke noted that NAVAIR is the acquisition arm for Naval Aviation and is responsible for the cybersecurity posture for all naval weapons systems, systems which represent a diversity of both information and operational technologies. He said that NAVAIR wants to be able to take advantage of advances in cybersecurity, but that poses a difficulty given the number of legacy systems in the inventory. He also addressed the challenge of how to take people who understand military aviation and enable them to deal with cybersecurity challenges, discussing his efforts to train hackers--those who can make systems work better through tweaks and shortcuts. He believes that an opportunity for academic research is how to quantify risk in cybersecurity, specifically as risk relates to CPS.

4. An Access Control Perspective on the Science of Security

Ravi Sandhu, University of Texas at San Antonio

Dr. Sandu provided comparisons between cybersecurity and the physical sciences but proposed that cybersecurity is an inherently different science and shouldn't be compared to natural sciences. He believes that there are stronger parallels with medicine and perhaps economics. He argued against the traditional boundary between basic and applied research for cybersecurity and suggested teams that address both aspects. Echoing a point made by many of the other speakers, Dr. Sandhu noted that cybersecurity is asymmetric and that makes it unique. He discussed the evolution of Access Control and suggested that cybersecurity can learn from the Access Control environment. He concluded by emphasizing the need to combine basic and applied research, treating cybersecurity holistically, and drawing inspiration from other sciences while not depending on the comparisons.

Invited Presentations

1. Building a Virtually Air-gapped Secure Environment in Amazon Web Services (AWS)

Erkang Zheng, Phil Gates-Idem, and Matt Lavin, LifeOmic, Inc.

The speakers talked about the work their company is doing building a platform on top of the cloud dealing with health care information where, because of the sensitivity of data, security is critical. They addressed the ten principles on which their work was built, including assuming the cloud is secure, assuring no single point of compromise, engaging everyone, and automation. They addressed the unique aspects of their program, including lessons-learned and future development.

2. You Get Where You're Looking For: The Impact of Information Sources on Code Security

Michelle Mazurek, University of Maryland

The presented paper was the winner of the NSA 5th Annual Best Scientific Cybersecurity Paper Competition. Authors Yasemin Acar, Michael Backes, Sascha Fahl, Doowon Kim, Michelle Mazurek, and Christian Stransky were from Saarland University in Germany and the University of Maryland, College Park, in the United States. The paper was presented at the 2016 IEEE Symposium on Security and Privacy.

Author Michelle Mazurek gave a presentation on their research which was inspired by a common problem, specifically why software developers are writing programs that have security vulnerabilities. When software developers get "stuck", they often turn to resources such as Stack Overflow to find solutions. Unfortunately, many of the posted solutions are not necessarily secure. The researchers investigated how different information sources available to the developer influence the developer's abilities to

program quickly and securely. They explored developers' problem-solving choices, and the impact on the software ecosystem. They noticed that an unsettling number of Android apps used readily available, and insecure, code snippets. They studied 54 developers, both professionals and students, in Germany and the United States in a controlled laboratory setting where they had them write security and privacy-relevant code under time constraints. They examined four conditions: developers were allowed to use 1) any source (free choice); 2) Stack Overflow only; 3) official Android documentation only; and 4) books only. After describing their methodology of subjecting Android developers to various security-relevant tasks and varying their choices of resources, they reviewed their findings on the impacts to both functional correctness, and security correctness. The researchers found that Official API documentation is secure but hard to use, while informal documentation such as Stack Overflow is more accessible but often leads to insecurity. Interestingly, books (the only paid resource) perform well both for security and functionality, but are rarely used. While suggesting that project managers should "take developers offline and give them a book," they chose to explore a more practical solution. They concluded that Stack Overflow provides quick functional solutions, but is less secure, and they developed several ideas to integrate both aspects. They noted that while professionals tended to produce functional code more reliably, they were no better than the students at security.

3. Microarchitectural Attacks: From the Basics to Arbitrary Read and Write Primitives without any Software Bugs
Daniel Gruss, Graz University of Technology

The speaker used the analogy of a safe to explain how systems may give clues to an attack, and he provided multiple examples. He demonstrated the Meltdown attack which exploited out of order execution/flush and reload attacks. He also demonstrated the Rowhammer attack in which cells leak faster from proximate accesses. He noted that microarchitectural attacks have been ignored in the past. Finally, he suggested that what we have learned from these attacks is the opportunity to rethink processor design, to "grow up" as other fields have done, to find tradeoffs between security and performance, and to spend more time on identifying problems rather than mitigating known problems.

Panel Discussion

Four Cybersecurity Framework (CSF) Practitioners participated in a panel discussion moderated by Nikola Vouk. The NIST CSF was described as the result of a collaborative effort with government and industry to identify what matters, figure out how to protect it, know when things go bad, and know what to do to correct them.

- Jeremy Maxwell, Allscripts, a national provider of Health IT, led off with a description of his company's approach to incidence response and risk management, noting that the company went with an ISO approach rather than the CSF because they view ISO as a holistic program that helps in preparation for incidents.

- Andrew Porter, Merck Pharmaceuticals said that the analytics group in risk management of Merck was one of the participants who worked on the CSF. He noted that the analytics group is having varying levels of success with the framework within the company. The primary benefit has been development of a common language and the ability to communicate cybersecurity risks to their Board of Directors.

- Alex Rogozhin, BB&T Bank, described his team as looking for data intelligence and security and a data-driven way to assess risk. In searching for a better reporting framework to keep their leadership aware of issues, they chose CSF. Their main difficulty has been practical implementation. They are trying to be compliant, which does not necessarily also mean secure.

- Greg Witte, G2, said NIST was charged to standardize and coordinate cybersecurity and, as a result, developed the CSF. He noted that the framework doesn't provide a lot of guidance since NIST wanted to avoid squashing innovation, and the design criteria included being flexible, agile, and applicable to companies of many sizes. The CSF aims to drive discussions about the "what is and what should be" and is purposely designed not to be prescriptive.

Industry and Government Presentations

1. DevSecOps: Security at the Speed of Software Development
Larry Maccherone, Comcast

Mr Maccherone referenced Mr. Lipner's morning keynote presentation in addressing obstacles to adoption, noting that "bolt-on" security by security specialists won't scale, so security must be a primary concern during development. He defined DevSecOps (DSO) as empowered engineering teams taking ownership of how their product performs in production, including security. He described a three-part framework for adopting new practices and DSO culture change: 1) principles acceptable to lean/agile development teams; 2) make it easy for the DSO teams to both understand what the right thing is and actually do it; and 3) engage management.

2. HACSAW: A Trusted Framework for Cyber Situational Awareness
William Glodeck, Department of Defense

Mr. Gloduck discussed the DoD High Performance Computing (HPC) modernization program and the HPC Architecture for Cyber Situational Awareness (HACSAW). The HACSAW initiative is a multi-disciplinary, multi-year project that examines the applicability of HPC to cyber SA using the most comprehensive cybersecurity dataset available to the DoD R&D community. The goal of HACSAW is to meet mission-essential tasks and reduce barriers to data and computing resources.

3. Compliance as Code: Policy-Governed Automated Security Checkpoints
Nikola Vouk, Independent, and David González, near Form

The presentation centered on how to move away from the stage-gate model and work at the speed of development. Mr. Vouk noted that governance is now with stage gate at the end, and that there are not enough resources at that point: 1 security person per 75 developers. He suggested the need to change that to 1 teacher per 75 students. He proposes an automated governance workflow, still with some manual steps, though noting that manual steps needed to be minimized. Mr. González provided brief demos that introduced policies that can lead to visualization on a dashboard and introduce metrics.


HOTSOS
2018
5TH ANNUAL
HOT TOPICS in the SCIENCE OF SECURITY
NC STATE UNIVERSITY
MERCK


The End

Tutorials

1. Combinatorial Security Testing Course
Rick Kuhn, NIST, and Dimitris Simos, SBA Research

The tutorial explained the background, process, and tools available for combinatorial testing for security, including illustrations based on industry's experience with the method. Mr. Kuhn presented the basics of combinatorial testing: what it is, how it works, and why it works. He noted that software testing may be up to half of overall software development cost, and that there was still a need to estimate the residual risk that remains after testing. This talk formulated the problem of software security testing as combinatorial problems, citing the need for empirical data to inform assumptions. Characterizing Combinatorial Security Testing (CST) as large-scale software testing for security, he noted that CST can make software security testing more efficient and effective than conventional approaches. Dr. Simos gave examples of how CTS is used in real world scenarios, developed models to demonstrate CTS, provided cased studies, and addressed experimental evaluation using different frameworks for different scenarios.

2. Applying the Framework for Improving Critical Infrastructure Cybersecurity
Greg Witte, G2

Mr. Witte provided a history of the Cybersecurity Framework (CSF), pointing out that an Executive Order made CSF applicable to all sectors, and that the framework was developed in partnership among industry, academia, and government. While there are multiple frameworks to leverage for cybersecurity, the CSF establishes a common language within organizations and among external partners, providing a good way to organize and communicate. He described the three components in CSF (core, profile, and implementation tiers) and how they combined to provide a holistic approach. He further described the seven steps in implementing CSF and addressed the changes that have been implemented in the updated version. Finally, he described resources that can be used to help implement the framework.

Posters

There were 57 cybersecurity research posters submitted to the program committee for consideration, and 18 were selected for presentation at HotSoS. One poster was awarded Best Poster. The Best Poster award is designed to encourage scientists across multiple disciplines to address the fundamental problems of security in a principled manner. The winning poster displayed the best combination of scientific rigor, clarity of presentation, and global impact.

Best Poster

1. What Proportion of Vulnerabilities Can Be Attributed to Ordinary Coding Errors?

Rick Kuhn and Raghu Kacker, National Institute of Standards and Technology, and Mohammad Raunak, Loyola University

The key question the authors sought to address is the degree to which vulnerabilities arise from ordinary program errors, which may be detected in code reviews and functional testing, rather than post-release. Findings include the fact that high severity vulnerabilities are trending downward, declining about 15 percentage points in the last ten years. About two-thirds of this fraction has shifted to medium severity vulnerabilities. Implementation or coding errors account for roughly two thirds of the total. They consider the proportion of implementation vulnerabilities, rather than absolute numbers, because the number of vulnerabilities is partially a function of the number of applications released, which has increased over time. Implementation vulnerabilities for 2008-2016 are close to the 64% reported for 1998-2003. This high proportion of errors suggests little progress has been made in reducing vulnerabilities from simple mistakes, and that more extensive use of static analysis tools, code reviews, and testing could lead to significant improvement.

2. A Comparative Analysis of Manual Methods for Analyzing Security Requirements in Regulatory Documents
Sarah Elder and Anna Mattapallil, North Carolina State University

This presentation is designed to assist analysts in selecting an appropriate approach for developing security requirements from regulatory documents by comparing the output of approaches from academic publications with similar outputs from industry. Initial results show that there is wide variance in how information is aggregated from security regulations at the requirement level.

3. An Expert-Based Bibliometric for a Science of Security
Lindsey McGowen and Angela Stoica, North Carolina State University

The research objective was to develop a scalable bibliometric customized for the Science of Security that would address limitations of existing citation-based bibliometrics. Existing citation databases do not adequately capture conferences and workshops where security researchers often publish, nor are they adaptive enough to be used with emerging fields of study. Computer science databases such as SiteSeerX and dbpl fall short of capturing venues appropriate for disseminating multidisciplinary research. Any citation-based metric will be a lagging indicator for fields that evolve at an extraordinarily fast pace. Expert-based review is a preferred method for evaluating faculty in computer science, which may be usefully applied to evaluation of publications. Their expert-based method shows potential for developing custom bibliometrics for evaluating publication venues in emerging and multidisciplinary fields.

4. Cryptography in a Post-Quantum World
Katharine Ahrens, North Carolina State University

Quantum resilience makes lattice-based hard problems a leading candidate for implementation in future public key cryptographic schemes. Lattice cryptosystems can offer both encryption schemes (to securely transmit data from sender to receiver) and signature schemes (used for a receiver to verify that information actually originated from the claimed sender). This poster gives an overview of past attempts to approach a lattice hard problem known as the Shortest Vector Problem (SVP) in a class of ideal lattices generated using the cyclotomic integers, a type of mathematical object known as a ring. The poster includes preliminary results on the security of the SVP in ideal lattices generated in a ring previously unstudied, and discusses the practically of using the ring in place of the cyclotomic integers in some lattice cryptosystems.

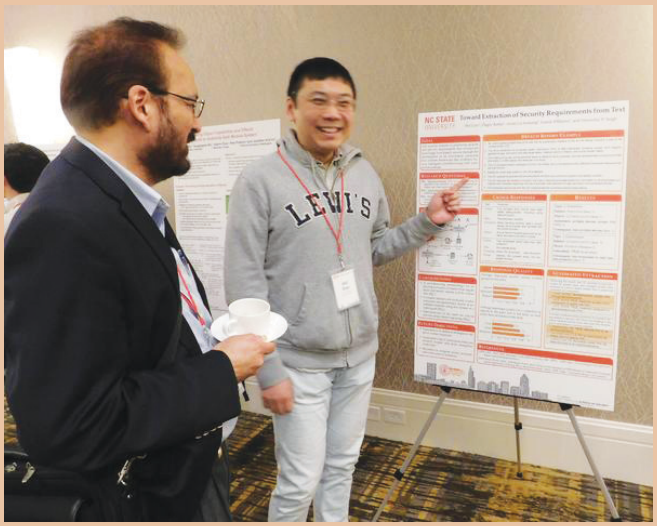


RansomAir Filled with Clouds
Dusko Pavlovic, University of Hawaii
• Problem: ransom attacks
• Insight: data are different
BACKUP
• Solution: cloud storage
• PROBLEM: cloud security
• INSIGHT: slow recovery ok
– crypto requires ongoing fast decryption
– data attacks are not ongoing
– data recovery is rarely needed
• SOLUTION: code, not cryp
Even lightweight encryption incurs slowdown
Coding incurs speedup
• SOLUTION: deletion channel security
SCORE

5. Detecting Monitor Compromise Using Evidential Reasoning
Uttam Thakore, University of Illinois at Urbana-Champaign

This poster demonstrates a data-driven technique to detect monitor compromise using evidential reasoning. Since hiding from multiple sensors is difficult for an attacker, combing alerts from different sensors by using Dempster-Shafer theory can identify potential attacks, monitor compromise, and compare the results to find outliers.

6. Ethics, Values and Personal Agents
Nirav Ajmeri, North Carolina State University

The research question addressed is "How can we engineer an ethical Socially Intelligent Personal Agent (SIPA) such that it understands its user's preferences among values, and reasons about values to make ethical policy decisions?" That question led to the development of Ainur, a framework for engineering value-driven, ethical SIPAs that can make value-promoting ethical decisions, especially in scenarios where the applicable norms conflict.

7. Exploring the Raspberry Pi for Data Summarization in Wireless Sensor Networks
Andrés Alejos, Matthew Ball, Conner Eckert, Michael Ma, Hayden Ward, Peter Hanlon, and Suzanne J. Matthews, United States Military Academy

Single board computers are good candidates for at-node data summarization tasks in a wireless sensor network. Reducing data transfer in a wireless sensor network is critical for energy efficiency and improved latency. This poster shows the viability of a wireless sensor network composed of Raspberry Pis for video and audio summarization tasks. Contributions include a novel sensor and gateway node design and a user interface implemented as an Android App.

8. Hourglass-Shaped Architecture for Model-Based Development of Safe and Secure Cyber-Physical Systems
Muhammad Umer Tariq and Marilyn Wolf, Georgia Institute of Technology

This proposed approach is inspired by the hourglass-shaped architecture of the Internet. It can support the goals of an integrated CPS theory and development methodology while taking into account the differences between the domain-specific skillset that control system engineers and embedded system engineers typically possess. This poster also outlines CPS-related safety and security concerns that the proposed hourglass-shaped architecture for networked CPS development must meet in order to address safety and security concerns.

9. How Bad Is It, Really? An Analysis of Severity Scores for Vulnerabilities
Christopher Theisen and Laurie Williams, North Carolina State University

In this presentation, a distribution of 2,979 vulnerabilities mined for Fedora 24 and 25 was analyzed using a high-medium-low evaluation rather than the usual binary vulnerability/no vulnerability method. The authors also verify the security vulnerabilities reported publicly as actual vulnerabilities and use keyword searches to identify bugs that should be included in vulnerability datasets.

10. Indirect Cyber Attacks by Perturbation of Environment Control: A Data-Driven Attack Model
Keywhan Chung, Zbigniew T. Kalbarczyk, and Ravishankar K. Iyer, University of Illinois at Urbana-Champaign

The indirect attack model targets a super computer by obfuscating the control of a CPS responsible for maintaining the operational environment. The authors' approach consists of four steps: data preparation, parameter analysis, inference of critical condition, and validation. Initial results indicate that their approach would have effectively identified two CPS-related incidents: a chilled water leakage at the construction site of a new building which could have caused an outage of the computing infrastructure, and a maintenance operation on the campus chilled water loop, which shut down a set of cabinets of the computer infrastructure.

11. Integrating Historical and Real-Time Anomaly Detection to Create a More Resilient Smart Grid Architecture
Spencer Drakontaidis, Michael Stanchi, Gabriel Glazer, Madison Stark, Caleb Clay, Jason Hussey, Nick Barry, Aaron St. Leger, and Suzanne J. Matthews, United States Military Academy

The authors developed a novel MapReduce algorithm to detect anomalies in historical grid data that leverages the cluster computing framework Apache Spark. The algorithm checks a sliding "window" of data for power fluctuations that meet the criteria of constraint and temporal anomalies described by Matthews and St. Leger. Experimentation was performed on a 36-core compute node on a supercomputer with a 1 million real measurements dataset collected from their test bed. Preliminary results show that the algorithm is capable of detecting constraint and temporal anomalies simultaneously.

12. Investigating TensorFlow for Airport Facial Identification
Nikolay Shopov, Mingu Jeong, Evin Rude, Brennan Nessaralla, Scott Hutchison, Alexander Mentis, and Suzanne J. Matthews, United States Military Academy

The authors describe a facial identification approach that can be deployed at airports. Their contributions include facial identification software built on top of Google's TensorFlow framework; a data collection scheme that can be implemented at airports nationally; and a user interface for collecting data.

13. Quantifying the Security Effectiveness of Network Diversity
Huashan Chen and Shouhuai Xu, University of Texas at San Antonio

This poster demonstrates a framework that quantifies the security effectiveness of network diversity in computer networks. The potential value of enforcing diversity in networks is well recognized, but security effectiveness of enforcing network diversity has not yet been quantified. In this work, the authors propose a systematic, fine-grained framework for modeling the diversification of software stacks in networks and quantifying network diversity security effectiveness using a suite of security metrics.

14. Quantitative Underpinnings of Secure Graceful Degradation
Ryan Wagner, David Garlan, Matt Fredrikson, Carnegie Mellon University

Defenders need a way to reason about and react to the impact of an attacker with existing presence in a system. It may not be possible to maintain one hundred percent of the system's original utility; instead, the defender might need to gracefully degrade the system, trading off some functional utility to keep an attacker away from the most critical functionality.

15. Ransomware Research Framework
Dan Wolf and Don Goff, Cyber Pack Ventures, Inc.

This research presented a series of joint efforts designed to produce a framework for studying ransomware. The seven contributors addressed detection, response, mitigation, consequences and attribution, as well as encryption, and an approach to modeling the problem from the behavioral viewpoint of criminology.

16. Toward Extraction of Security Requirements from Text
Özgür Kafali, University of Kent, Anne-Liz Jeukeng, University of Florida, Laurie Williams, Hui Guo, and Munindar P. Singh, North Carolina State University

The goal of this research was to produce improved security and privacy requirements that accommodate both social and technical considerations and incorporate knowledge from post-deployment artifacts such as breach reports. This framework combines crowdsourcing with automated methods to produce improved security and privacy requirements incorporating knowledge from post-deployment artifacts such as breach reports.

17. Understanding Privacy Concerns of Whatsapp Users in India
Jayati Dev, Sanchari Das, and Jean Camp, Indiana University

WhatsApp is a leading platform for mobile messaging with the largest user base being in India, yet research on Indian perspectives towards privacy and security in social networking platforms is sparse. WhatsApp incorporates features which pose privacy challenges, including Last Seen, Live Location, and personal profile information. The researchers implemented a survey, querying both privacy attitudes and privacy behaviors, with 213 Indian participants. They found the majority of participants reported that they actively use the privacy controls provided by WhatsApp to restrict access to their information. They provided visualizations of the raw results and initial recommendations.

18. Using Object Capabilities and Effects to Build an Authority-Safe Module System
Darya Melicher, Yangqingwei Shi, and Valerie Zhao, Wellesley College, Alex Potanin, Victoria University of Wellington, and Jonathan Aldrich, Carnegie Mellon University

The research team designed and implemented a capability-based module system that facilitates controlling the security capabilities of software modules. Their approach ensures that a software system maintains the principle of least authority and also allows for attenuation of module authority. This design is implemented as part of the Wyvern programming language.

HotSoS 2018 Program Committee Members

NCSU Lablet Co-PIs Munindar Singh and Laurie Williams served as General Co-Chairs for HotSoS 2018. Program Co-Chairs were Rick Kuhn of the National Institute of Standards and Technology (NIST) and Tao Xie from the University of Illinois at Urbana-Champaign.

HotSoS Program Committee Members:

Jonathan Aldrich, Carnegie Mellon University
Homa Alemzadeh, University of Virginia
Jean Camp, Indiana University
Amit Chopra, Lancaster University, United Kingdom
Daniela Cruzes, SINTEF Digital
Michel Cukier, University of Maryland, College Park
Christopher Gates, Symantec
Vincent Hu, National Institute of Standards and Technology
Limin Jia, Carnegie Mellon University
Özgür Kafalı, North Carolina State University
Sneha Kasera, University of Utah
Jonathan Katz, University of Maryland, College Park
Nadin Kökciyan, King's College London, United Kingdom
Constantinos Kolias, George Mason University
Carl Landwehr, George Washington University
Yves Le Traon, University of Luxembourg, Luxembourg
Emil Lupu, Imperial College London, United Kingdom
Aaron Massey, University of Maryland, Baltimore County
Sayan Mitra, University of Illinois at Urbana-Champaign
Pradeep Kumar Murukaniah, Rochester Institute of Technology
Christopher Oehmen, Pacific Northwest National Laboratory
Pete Rotella, Cisco
Sean Smith, Dartmouth College
Adam Tagert, National Security Agency
Claire Vishik, Intel Corporation, United Kingdom
Jeff Voas, National Institute of Standards and Technology
Tim Weil, Scram Systems
David Wheeler, Institute for Defense Analyses
Rebecca Wright, Rutgers University
Dinghao Wu, Pennsylvania State University
Xusheng Xiao, Case Western Reserve University
Zhi Xu, Palo Alto Networks
Danfeng Yao, Virginia Polytechnic Institute
Ting Yu, Qatar Computing Research Institute, Qatar

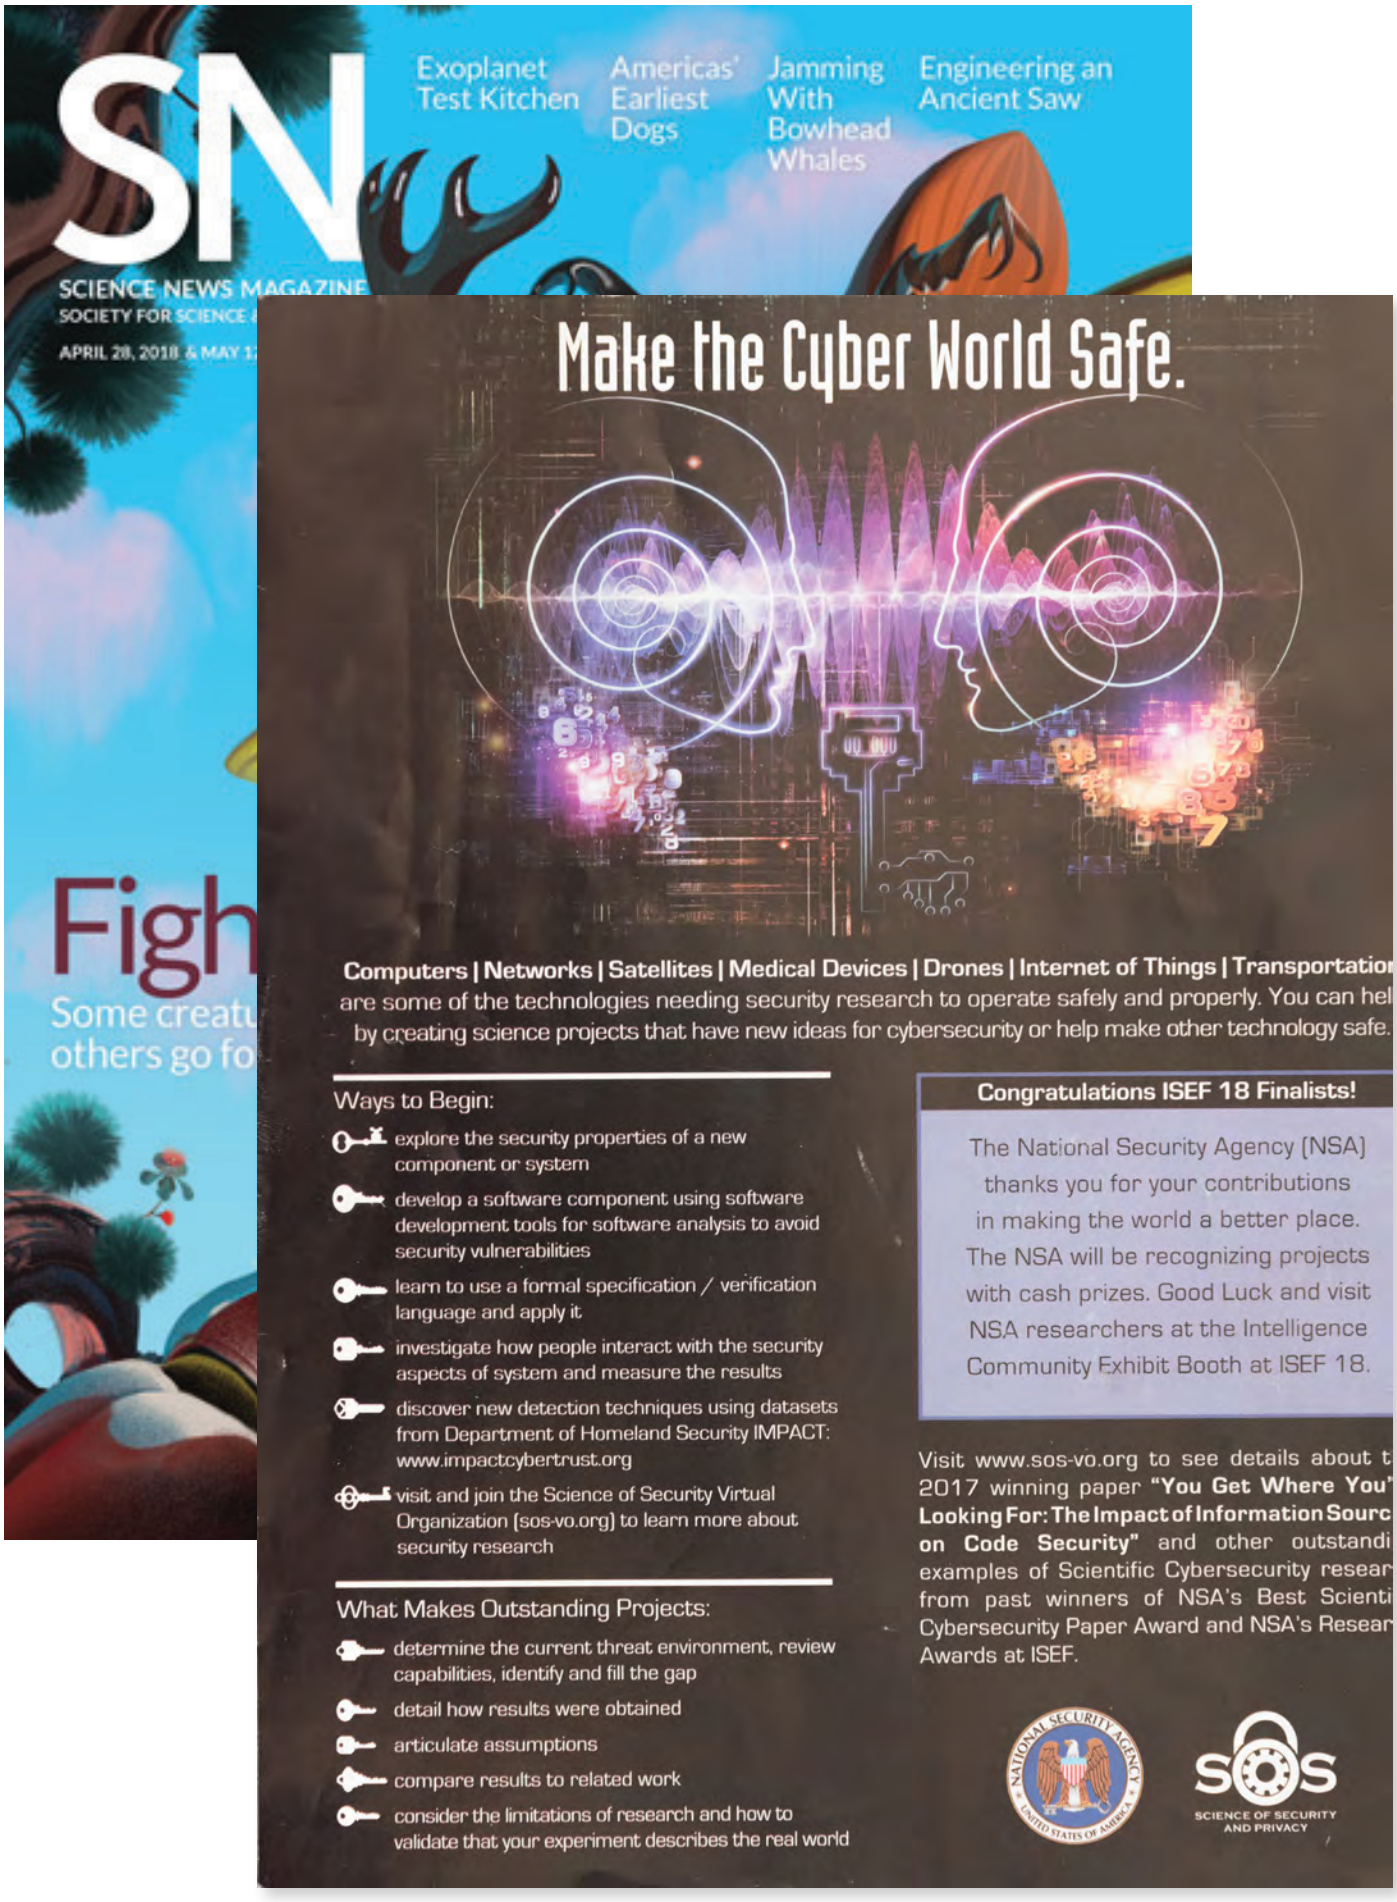
SN
SCIENCE NEWS MAGAZINE
SOCIETY FOR SCIENCE
APRIL 28, 2018 & MAY 1
Exoplanet Test Kitchen
Americas' Earliest Dogs
Jamming With Bowhead Whales
Engineering an Ancient Saw
Fight
Some creatu
others go fo
Make the Cyber World Safe.
Computers | Networks | Satellites | Medical Devices | Drones | Internet of Things | Transportatio
are some of the technologies needing security research to operate safely and properly. You can hel
by creating science projects that have new ideas for cybersecurity or help make other technology safe.
Ways to Begin:
explore the security properties of a new component or system
develop a software component using software development tools for software analysis to avoid security vulnerabilities
learn to use a formal specification / verification language and apply it
investigate how people interact with the security aspects of system and measure the results
discover new detection techniques using datasets from Department of Homeland Security IMPACT: www.impactcybertrust.org
visit and join the Science of Security Virtual Organization (sos-vo.org) to learn more about security research
What Makes Outstanding Projects:
determine the current threat environment, review capabilities, identify and fill the gap
detail how results were obtained
articulate assumptions
compare results to related work
consider the limitations of research and how to validate that your experiment describes the real world
Congratulations ISEF 18 Finalists!
The National Security Agency (NSA) thanks you for your contributions in making the world a better place. The NSA will be recognizing projects with cash prizes. Good Luck and visit NSA researchers at the Intelligence Community Exhibit Booth at ISEF 18.
Visit www.sos-vo.org to see details about t
2017 winning paper "You Get Where You'
Looking For: The Impact of Information Sourc
on Code Security" and other outstandi
examples of Scientific Cybersecurity resear
from past winners of NSA's Best Scienti
Cybersecurity Paper Award and NSA's Resear
Awards at ISEF.
NATIONAL SECURITY AGENCY
UNITED STATES OF AMERICA
SOS
SCIENCE OF SECURITY AND PRIVACY

SoS in the News


Taylor Johnson @taylorjohnson · May 8
Write up on our new NSA Science of Security Lablet, focusing on cyber-physical systems and internet of things! engineering.vanderbilt.edu/news/2018/kout...
@verivital @VUEngineering @NSAGov
Koutsoukos heads NSA Lablet to enhance America's post-hack resili...
Cyber-physical systems let you analyze Fitbit data on a smartphone. They tell your house to bump up the thermostat before you get home. They run t...
engineering.vanderbilt.edu


NSA/CSS
Follow
#NSA is proud to support the next generation of #STEM leaders at @Society4Science #IntelISEF2018. NSA Research will award $13,000 worth of scholarships today.
NSA/CSS
@NSAGov


NSA/CSS
@NSAGov
Follow
Congrats to the #soups2018 Distinguished Paper Award winners for their research on the consumer's perception of risks & actions following a major data breach. #NSA's Science of Security Initiative sponsored the award. bit.ly/2nBXmme


CMU School of Computer Science Retweeted
SEI News @SEInews · 12 Oct 2018
@NSAGov honors @CarnegieMellon & @SEInews researchers for papers submitted for the Best Scientific #Cybersecurity Paper Competition. Read the papers @ sei.cmu.edu/news-events/ne...
Messages
Search


Messages
Search
Carnegie Mellon ECE Retweeted
CyLab @CyLab · 16 Oct 2018
Congrats to CyLab's @thedavidbrumley and @__tiffanyb__ -- this year's winners of NSA's Best Scientific Cybersecurity Paper competition!
CyLab researchers win NSA's Best Scientific Cybersecurity Paper co...
CyLab researchers win NSA's Best Scientific Cybersecurity Paper competition - CyLab Security and Privacy Institute
cylab.cmu.edu


NSA/CSS
@NSAGov
Follow
Are you the next winner of #NSAs Best Scientific Cybersecurity Paper Competition? If you've recently submitted an academic paper to a peer-reviewed journal, magazine, or technical conference consider entering by March 30th for a chance to win. Learn more: bit.ly/2FcTZwy
BEST SCIENTIFIC CYBERSECURITY PAPER COMPETITION
NSA/CSS


Shawn P Riley shared a link.
December 12, 2018
Part of my efforts to share knowledge, as selected by Twitter vote, a video series on operational cybersecurity science and symbolic AI called Talking SoS. The series kicks off with a 15 minute overview of cyber threat intelligence.
YOUTUBE.COM
Talking SoS with Shawn Riley - A 15 minute overview of cyber threat intelligence
A video series on the Science of Security (SoS) with Cybersecurity Scientist Shawn Riley Recorded -...
Seen by 13
Like
Comment
Share


Ismail Khoffi
@KreuzUQuer
Follow
"Practicing a science of security - A Philosophy of Science Perspective": benthamsgaze.org/2018/01/16/pra ...
Related to twitter.com/CormacHerley/s ...
UCL InfoSec Group @uclisec
Practicing a science of security benthamsgaze.org/2018/01/16/pra...
8:28 AM - 16 Jan 2018

« IMMIGRATIONS AND CUSTOMS ENFORCEMENT OFFICERS ARREST 225 IN NEW YORK

IAEA USES DRONES TO RELEASE STERILE MOSQUITOES IN HIGH-RISK ZIKA AREAS »

NSA hosts researchers from six Science and Security Privacy initiative 'lablets'

Friday, April 20, 2018 by Kevin Randolph

The National Security Agency (NSA) recently hosted researchers from "lablets" taking part in its Science of Security and Privacy (SoS) initiative, which advances the design of trusted systems across disciplines ranging from computer science to behavioral science.

Since its launch in 2012, the NSA's SoS initiative has advanced security and privacy science as a recognized field of research. Most recently, the University of Kansas, Vanderbilt University, and the International Computer Science Institute joined were named SoS lablets. The three original lablets were Carnegie-Mellon University, the University of Illinois-Champaign and North Carolina State University.

"The NSA SoS lablets will focus on the discovery of formal underpinnings of the design of trusted systems which spans the disciplines of computer science, electrical engineering, mathematics, behavioral science, statistics, philosophy, public policy and physics," an NSA release stated.

The multidisciplinary lablets have launched 20 projects that address challenges in cyber-physical systems, cybersecurity metrics, policy-governed secure collaboration, privacy, resilient architectures, scalability and composability, and understanding and accounting for human behavior.

More than 300 universities across the country have applied to be SoS lablets. The scientific rigor or research projects, timeliness of challenges being addressed, and advancement of existing efforts to build security and privacy in the scientific community are among the selection criteria.

« IMMIGRATIONS AND CUSTOMS ENFORCEMENT OFFICERS ARREST 225 IN NEW YORK

IAEA USES DRONES TO RELEASE STERILE MOSQUITOES IN HIGH-RISK ZIKA AREAS »

Sign in or sign up to access your shots across devices and save your favorites forever.

EECS Graduate Student Pooria Ma...

just now expires in 14 days

Take Note!
Updated editing tools let you crop, highlight, and even add text to your shot.
Give them a try

EECS Webmail Login EECS Internal

EECS ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

Current Students Community Programs Activities Research Our Home Future Students

News

Home / News / Graduate News / EECS Graduate Student Pooria M ...

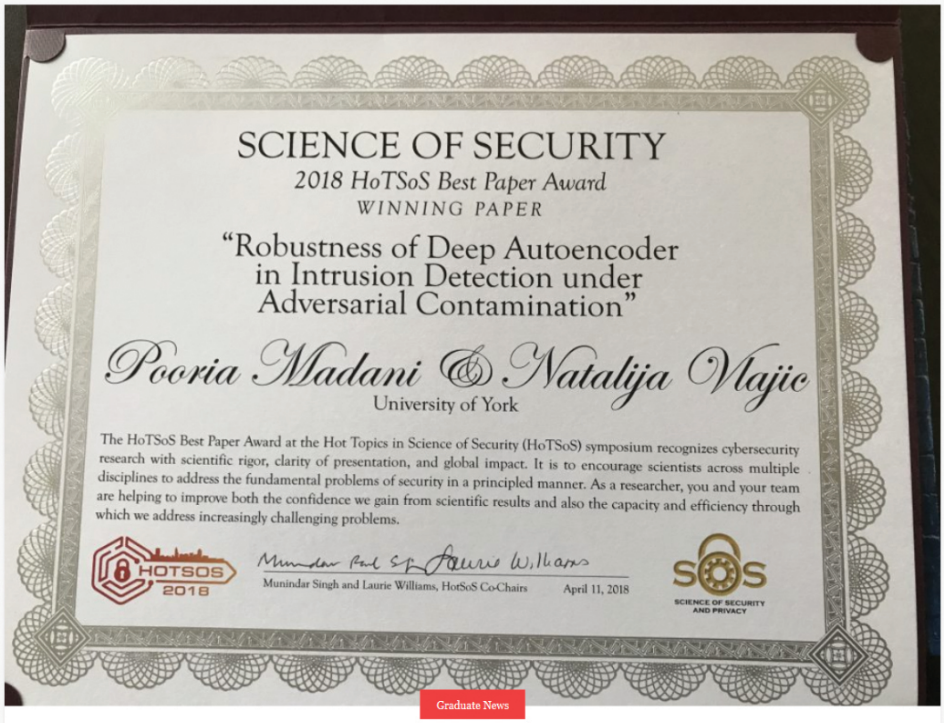

Graduate News

EECS Graduate Student Pooria Madani, and Prof. Natalija Vlajic, are recipients of 2018 HoTSoS Best Paper Award.

HotSos is a research conference centered on the Science of Security, which aims to emphasize the advancement of research methods as well as the development of new research results. The award winning paper by Madani and Vlajic studies the viability of using deep autoencoders in the detection of anomalies in adaptive Intrusion Detection Systems (IDSs), as well as their overall robustness against adversarial poisoning.

Apr 18 / Ouma Jaipaul-Gill / Comments Disabled

Outreach

In addition to the activities described elsewhere in this report, the Science of Security and Privacy (SoS) initiative used a variety of other means to expand awareness of Science of Security in 2018. Researchers working at the Lablets and Sub-Lablets and collaborators from elsewhere in academia, government and industry served as Science of Security ambassadors within their own organizations and at local and international symposia and conferences.

One of the primary means of outreach is the Science of Security Virtual Organization (SoS-VO) which was established to provide a focal point for the initiative's significant research results, activities, and artifacts. It emphasizes community development, information sharing, and interaction among researchers in the field. SoS-VO membership grew to over 1500 members in 2018, extending the SoS presence to universities, research centers, private companies, and government agencies worldwide. The SoS-VO provides a forum to discover resources, connect to others, and share and survey cybersecurity research. The goal of the SoS-VO is to help establish and support true collaboration in advancing cybersecurity science.

The SoS-VO provides information on SoS activities, to include Lablet research, Lablet quarterly meetings, HotSoS, the annual Best Scientific Cybersecurity Paper Competition, the Intel International Science and Engineering Fair (Intel ISEF), and other SoS activities. This year the SoS-VO built on its sponsorship of the Intel ISEF awards by posting information on the SoS-VO on how to get started on a cybersecurity science fair project. The SoS-VO enables its members to post research findings and publications done elsewhere, hosts chats, blogs, and forums, and provides information on upcoming events, position openings, calls for papers for conferences, and general cybersecurity news. In 2018 there were over 1500 unique news items posted on the SoS-VO.

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

The SoS Reviews and Outreach (SoS R&O) newsletter published 12 editions in 2018 to over 1500 subscribers. The purpose of the R&O is to highlight the research, news, and events that impact the SoS technical community. All materials included in the R&O are also available on or through the SoS-VO and are organized as follows:

- **Pub Crawl:** A summary, organized by Hard Problem, of publications that have been peer reviewed and presented at SoS-related conferences or referenced in current work. The topics are chosen for their usefulness for current researchers. There were approximately 2600 Pub Crawl items published in 2018 covering over 260 topics and representing the curated work of over 15,000 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 2018 approximately 800 news items were included in the R&O.
- **Upcoming Events**: Information on SoS-related conferences, symposia, and workshops.
- **Cyber Scene**: Material that provides an informative, timely backdrop of events, thinking, and developments that contribute to the technological advancement of SoS Cybersecurity collaboration and extend its outreach. This section explores other dimensions of cyber research beyond the academic, and also addresses US and international policy issues, proposed regulations here and abroad, congressional inquiries and testimony, and in-depth articles from non-technical publications.
- **Musings:** Brief articles on areas of concern or interest in areas of Science of Security.

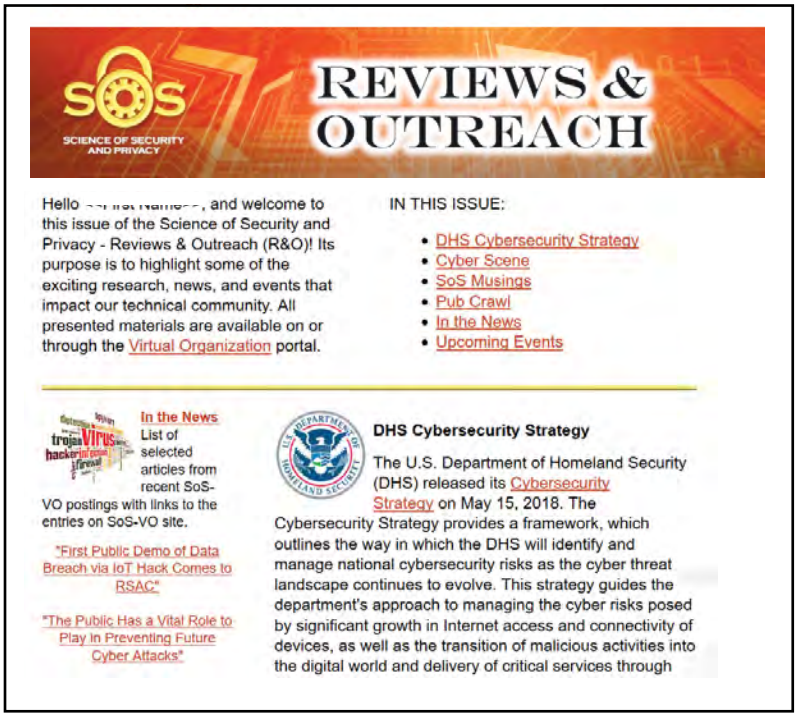

Sample SoS R&O

NSA booth at RSA

In addition to HotSoS, the Intel ISEF awards, the Annual Best Scientific Cybersecurity Paper Competition, and the new Conference Distinguished Paper Competitions, SoS participated in a number of conferences and workshops where SoS personnel shared information on the science of security and SoS initiatives. The SoS initiative had a booth as part of the NSA booth at RSA 2018 in San Francisco, CA from April 16-18, and was selected to be one of the top 10 programs showcased by NSA. SoS presented a poster at SOUPS on the Annual Best Scientific Cybersecurity Paper Competition. (The poster is shown elsewhere in this report.) SoS personnel had a booth at the NIST NICE K-12 Conference (National Institute for Standards and Technology National Initiative for Cybersecurity Education), December 3-4 in San Antonio, TX and participated at the CSforAll Summit, October 8-11 in Detroit, MI where they demonstrated the technologies developed at the CPS Summer Camp (see elsewhere in Section 3 for details on the Summer Camp). At the Society for Advancement of Chicanos/Hispanics and Native Americans in Science (SACNAS) conference in San Antonio on October 13, SoS personnel staffed the NSA booth and participated in a panel on the diversity of cybersecurity. SoS also briefed the National Academies of Science—Intelligence Community Studies Board on SoS activities.

Lablet summer programs geared towards K-12 students and SoS engagement at the Intel ISEF provided opportunities to engage the next generation of researchers. The incorporation of fundamental research findings in Lablet and Sub-Lablet undergraduate coursework increases emphasis on Science of Security principles in a wide range of related disciplines. 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.

Vanderbilt University CPS Summer Camp

Vanderbilt University's Institute for Software Integrated Systems (ISIS) offered a Cyber-Physical Systems (CPS) Security Summer Camp for middle and high school students in June 2018. ISIS partnered with NSA, the USAF Research Lab and the National Science Foundation to hold two sessions of the camp. This is the second summer for CPS camps and ISIS plans to offer them again in 2019.

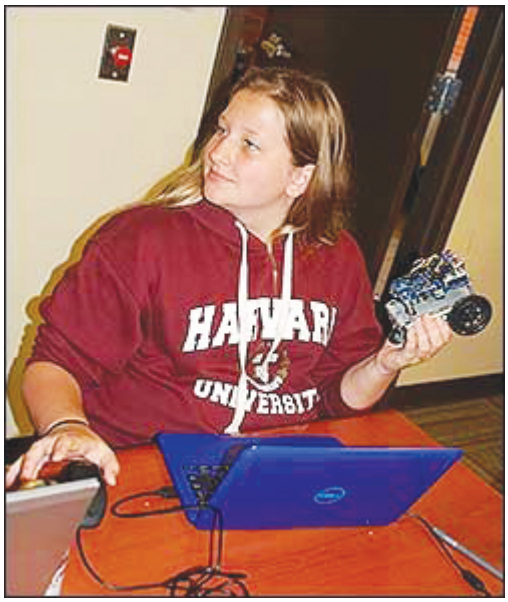

High School student programs Parralax robot

The camps are a free-of-charge, five-day experience for highly qualified local students interested in the growing field of CPS security. A teacher from each school is invited to attend as well. The goals of the program are to help students understand CPS security, increase diversity and interest in CPS, improve teaching methods for delivering content in CPS curricula, and prepare students for related coursework in college. The long-term goals of this initiative include establishing (or building) a pipeline that targets highly-qualified students who are interested in CPS-related engineering disciplines.

Campers learned to program Parralax robots using RoboScape NetsBlox*, a visual programming environment created by Professor of Computer Engineering Akos Ledeczi and his team. NetsBlox introduces programming basics and a high-level view of distributed computing. The team has worked with students as young as middle schoolers, and the June camps are another opportunity to work with young learners. Campers started by learning the fundamentals of programming and especially the programming of the NetsBlox platform and then moved to controlling the robot. By midweek, campers with little or no knowledge of computer programming had learned control structures, variables, data types and functions. They were able to write scripts and drive robot cars using the keyboard. They also wrote a self-driving program and participated in a tug a war--not between two robots, but two users on one robot; the winner was the one that directed the robot to their end of the court. They also learned about cyber security of cars and some of the known vulnerabilities and attacks in the context of modern cars. Cyber security of GPS satellites and receivers and of unmanned aerial vehicles were additional topics. By the end of the camp, they had learned how to encrypt messages. The camp concluded with another cyber attack and defense exercise.

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

Adam Tagert, the SoS Technical Director, and Capt. Tina McAfee, SoS Deputy Program Manager, attended one of the camps to get familiar with the material in the hopes of using the camp curriculum to support other SoS outreach activities for younger learners.

The technologies developed from the CPS Summer camps were demonstrated twice following the camps. The first time was at the CSforAll Summit, October 8-11 in Detroit, MI and then at the NIST NICE K-12 Conference (National Institute for Standards and Technology National Initiative for Cybersecurity Education), December 3-4 in San Antonio, TX. These materials served to underscore NSA's commitment to advancing cybersecurity.

Middle school student prepping for robot obstacle course

* RoboScape is a collaborative, networked robotics environment that makes key ideas in computer science accessible to groups of learners in informal learning spaces and K12 classrooms. RoboScape is built on top of NetsBlox from Vanderbilt University, an open-source, networked, visual programming environment based on Snap! that is specifically designed to introduce students to distributed computation and computer networking. RoboScape provides a twist on the state-of-the-art of robotics learning platforms. First, a user's program controlling the robot runs in the browser and not on the robot. There is no need to download the program to the robot and hence, development and debugging become much easier. Second, the wireless communication between a student's program and the robot can be overheard by the programs of the other students. This makes cybersecurity an immediate need that students realize and can work to address.

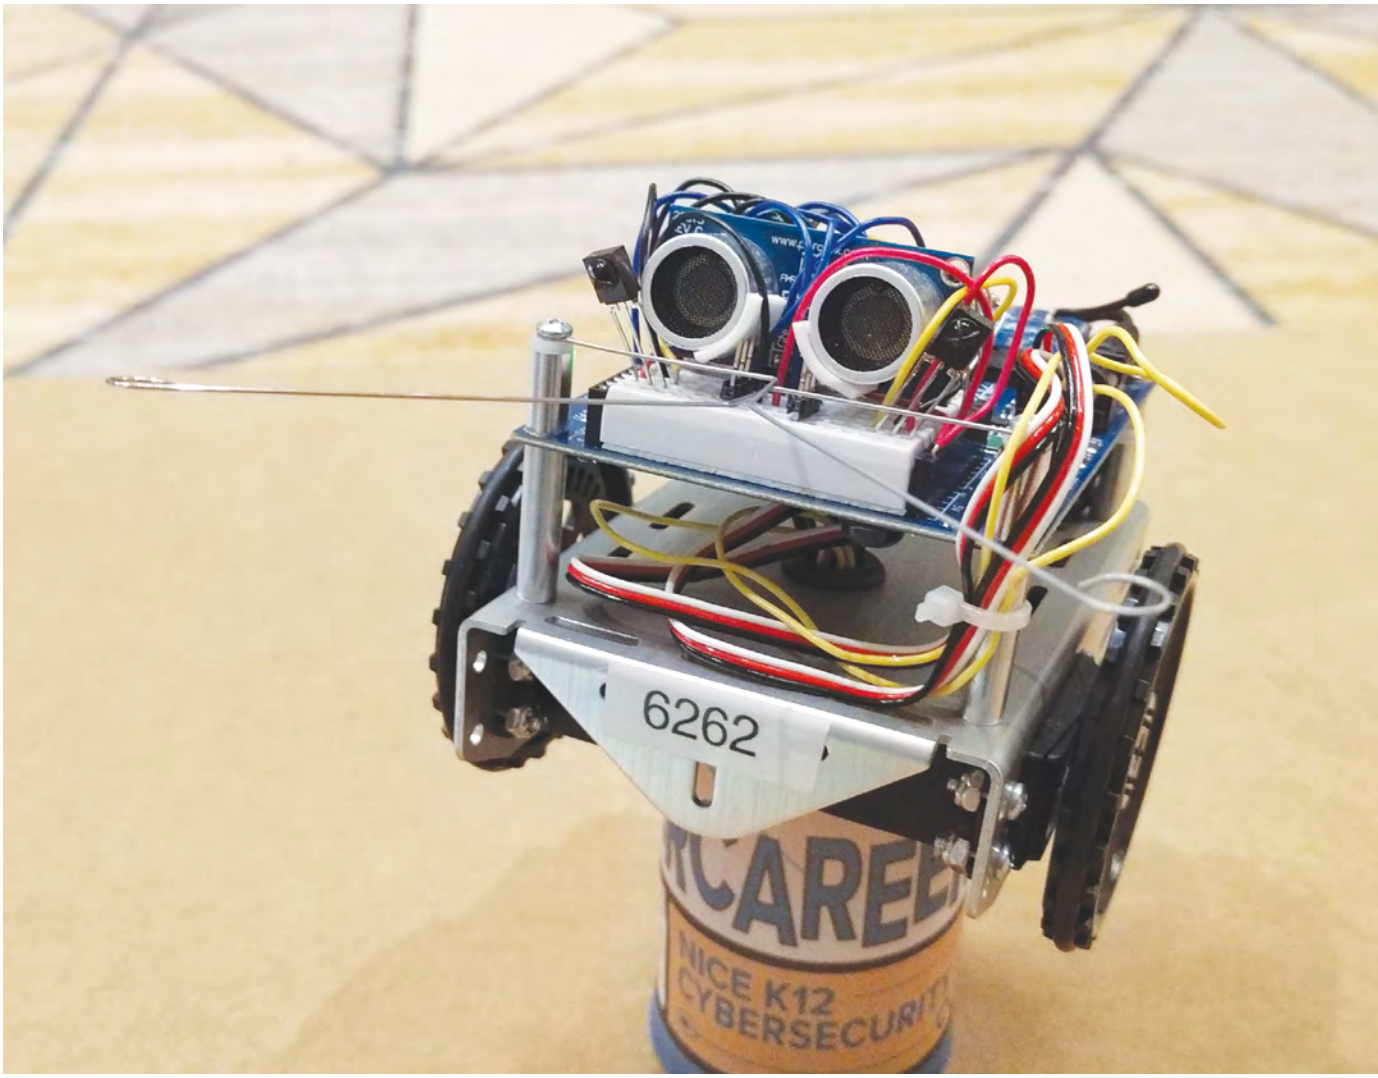


SN
IENCE NEWS MAGAZINE
OCIETY FOR SCIENCE & THE PUBLIC
NOVEMBER 24, 2018
Countries Line Up for Moon Visits
Appendix Linked to Parkinson's
Plugging Holes in Water Quality
Roaches Rebuff Zombie-Making Wasps
Mammo
Lives
Ancient elephant
giving up their se
ADVERTISEMENT
SOS
DEFENDING
OUR NATION
SECURING
THE FUTURE
LEARN & TEACH
Robot Hacking
CYBERSECURITY EDUCATION
SOS-VO.ORG/CPSS-EDU
Cyber-Physical Systems (CPS) Cybersecurity Curriculum
Free curriculum with supporting technologies to introduce students to the exciting field of cybersecurity. Enable K-12 educators and students to understand the methods of hacking and defensive solutions by experimenting on wireless controlled robots. Curriculum uses NetsBlox based on Scratch and Snap! for students with no programming experience to learn cybersecurity skills.
SEE US AT NICE K-12
Get better acquainted with Netsblox and the cybersecurity curriculum by visiting the NSA/Vanderbilt booth at the NICE K-12 Cybersecurity Education Conference. December 3-4, 2018 in San Antonio, Texas.
Participate in Educator Workshops Available in 2019
Learn to utilize the cybersecurity curriculum and introduce cybersecurity principles to students.
Be Involved with Cybersecurity Research All Disciplines Welcome
Want to help the direction of cybersecurity? Research and science fair projects are a great way to advance the field. Getting started is easy and outstanding projects are recognized by the National Security Agency with cash prizes at ISEF 2019 and regional science fairs.
VANDERBILT UNIVERSITY
NSA SCIENCE OF SECURITY AND PRIVACY


SOS
sos-vo.org