Science of Security & Privacy

2017 ANNUAL REPORT

Science of Security
and
Privacy (SoS) Initiative

Annual Report

2017

Research Directorate
Advancing Intelligence Through Science

Table of Contents

Executive Summary 3
Privacy

Section 1: Engaging the Academic Community for Foundational Research 10
Science of Security Lablets Progress on Hard Problems 2017 10
Scalability and Composability 11
Policy-Governed Secure Collaboration 13
Security Metrics and Models 14
Resilient Architectures 16
Understanding and Accounting for Human Behavior 18
Carnegie Mellon University 20
Fundamental Research 21
Publication 24
Educational 25
Community Outreach 25
North Carolina State University 26
Fundamental Research 27
Publications 30
Educational 31
Community Outreach 32
University of Illinois at Urbana-Champaign 33
Fundamental Research 34
Publications 37
Educational 39
Community Outreach 39
University of Maryland 41
Fundamental Research 41
Publications 44
Educational 45
Community Outreach 46
Science of Security Quarterly Meeting 48
Winter 2017 Quaterly: North Carolina State University 48
Spring 2017 Quaterly 50
Summer 2017 Quartley: Carnegie Mellon University 50
Science of SecUre and REsilient Cyber-Physical Systems (SURE) 54
Publications 57

SECTION 2: Promoting Rigorous Scientific Principles 58
Annual Best Scientific Cybersecurity Paper Competition 59
Intel International Science and Engineering Fair 61

SECTION 3: Growing the Science of Security 64
HotSoS 66
Research Papers 67
Keynotes 69
Posters 70
Tutorials 76
The Science of Security Virtual Organization (SoS-VO) 79
Outreach 80

digital copies are available at https://CPS-VO.org/group/sos/annualreport2017

Executive Summary

In 2017, the Science of Security and Privacy initiative (SoS) increased in both size and scope as members of academia, industry, and government researched and collaborated under the sponsorship of the National Security Agency (NSA) Research Directorate (RD). NSA RD secures the future by conducting ground-breaking research in a wide variety of science, technology, engineering, and mathematics areas, and in 2017 the SoS research they sponsored provided theories, models, methodologies, tools, technologies, and designs to safeguard interactions in cyberspace.

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

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

The SoS initiative engaged the academic community for foundational research in 2017 by continuing to sponsor four Science of Security Lablets (Carnegie Mellon University, North Carolina State University, the University of Illinois at Urbana-Champaign, and the University of Maryland) and 25 participating Sub-Lablets. The Lablets engaged in 42 foundational research projects during 2017, all of which were aligned with one or more 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

Some of the major achievements in 2017 included:

- A tool (JITANA) that addresses security vulnerabilities that arise from the interaction of multiple applications on a platform such as Android by enabling the analysis of the applications at classloading, rather than compilation, time. JITANA has been shown to be 4 to 30 times more efficient than a state-of-the-art approach, and it can effectively and efficiently analyze complex apps including Facebook, PokemonGo, and Pandora.
- A new method to simultaneously analyze very large numbers of compile-time configurations of C code at once, without looking at each configuration individually. This approach addresses the problem of security issues that arise from system evolution and configuration.
- The application of network analysis techniques to revision data from software repositories to identify unusual commits that are outliers and merit security review in order to address the problem of how to predict which code fragments within a system such as the Linux kernel are most likely to contain security vulnerabilities. This mechanism detected code commits that are actionable for developers to prioritize reviews, enabling the security review process to scale up more effectively.
- A new language that supports the tri-modular representation of firewall policies coupled with analysis of policies expressed in that language. The development enabled the specification of policies with fewer errors than otherwise possible.
- A two-layer stochastic game-theoretic model to effectively configure and share resources among collaborative Intrusion Detection Systems (IDS). In a case study of ten collaborative IDS, there was a performance improvement of 120%.
- A framework that factors out the complexity of implementing security policies from the complexity of implementing performance optimizations, in the context of SDN controllers. This development makes design of SDN resilient networks more achievable.

NSA's Science of Security

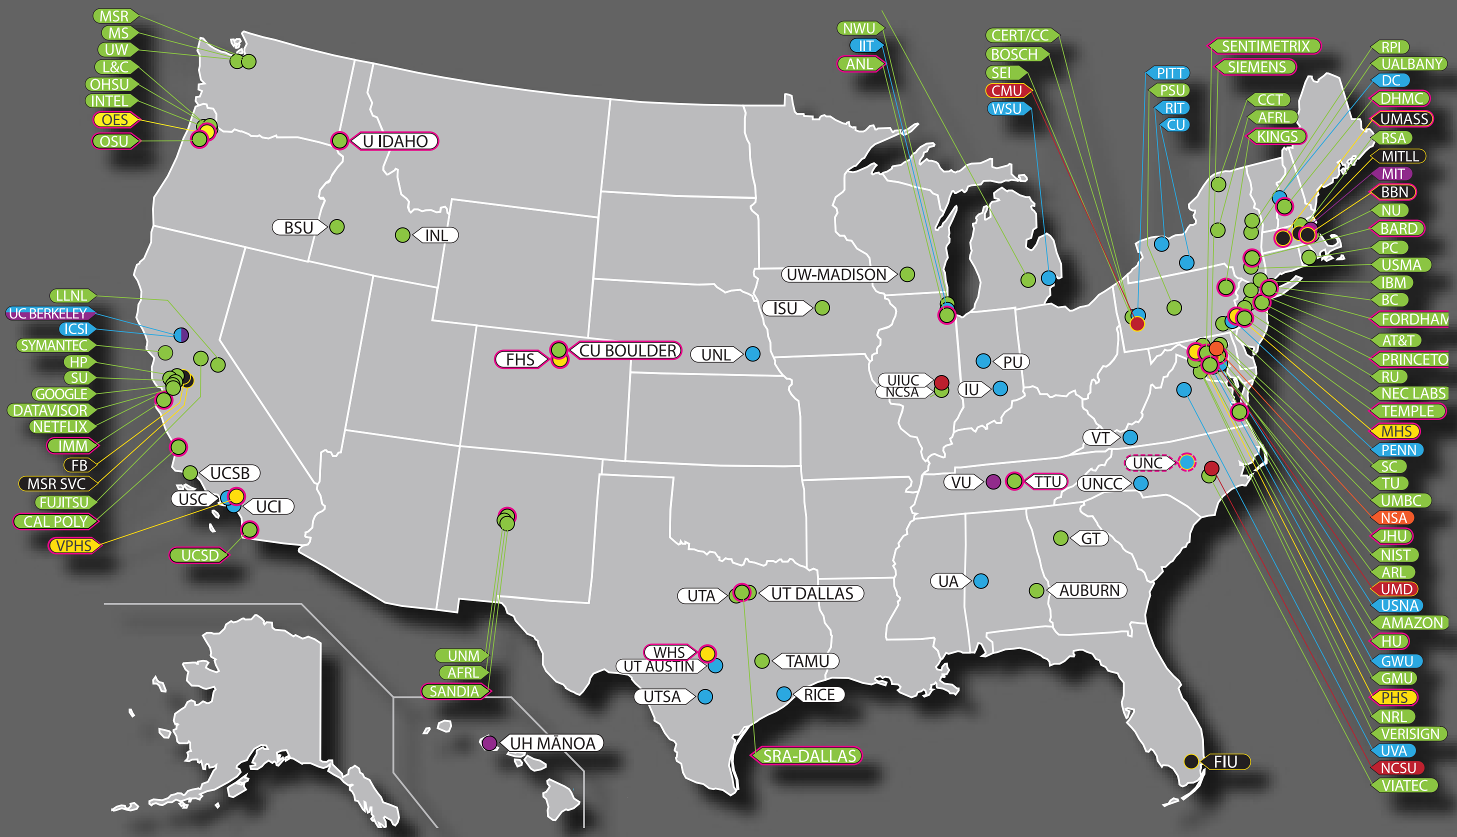

ABBREVIATIONS

🟢	AFRL	Air Force Research Laboratory
🟢	AMAZON	Amazon Web Services
🟢	ANL	Argonne National Lboratory
🟢	ANU	Australian National University
🟢	ARIEL	Ariel University
🟢	ARL	US Army Research Lab
🟢	AT&T	AT&T Labs Research
🟢	AUBURN	Auburn University
🟢	BARD	Bard College
⚫	BBN^{4th}	Raytheon BBN
🟢	BC	Brooklyn College
🟢	BIU	Bar-Ilan University
🟢	BOSCH	Bosch Research and Technology Center
🟢	BSU	Boise State University
🟢	BUAA	Beihang University
🟢	CAL POLY	California Polytechnic State University
⚫	CAMB^{1st}	University of Cambridge
🟢	CAS	Chinese Academy of Sciences
🟢	CCT	Clarkson University
🟢	CERT-CC	CERT Coordination Center
🔴	CMU$^{3rd \& 4th}$	Carnegie Mellon University
🔵	CU	Cornell University
🟢	CU BOULDER	University of Colorado, Boulder
🟢	CU-CANADA	Concordia University
🟢	DATAVISOR	DataVisor
🔵	DC	Dartmouth College
🟢	DHMC	Dartmouth-Hitchcock Medical Center
🟢	EPFL	École Polytechnique Fédérale de Lausanne
🟢	ETH ZÜRICH	ETH Zürich
⚫	FB^{3rd}	Facebook
🟡	FHS^{2015}	Fairview High School
⚫	FIU^{3rd}	Florida International University
🟢	FORDHAM	Fordham University School of Law
🟢	FRAUNHOFER IESE	Fraunhofer Inst. for Experimental Software Eng.
🟢	FUJITSU	Fujitsu Laboratories
🟢	GMU	George Mason University
🟢	GOOGLE	Google Headquarters
🟢	GT	Georgia Institute of Technology
🔵	GWU	George Washington University
🟢	HP	Hewlett Packard
🟢	HU	Hampton University
🟢	IBM	IBM, T.J. Watson Research Center
🔵	ICSI	International Computer Science Inst. at Berkeley
🔵	IIT	Illinois Institute of Technology
🟢	IMDEA	IMDEA Software Institute
🟢	IMM	Institute for Molecular Manufacturing
🟢	INL	Idaho National Laboratory
🟢	INTEL	Intel Labs
🟢	ISU	Iowa State University
🔵	IU	Indiana University
🟢	IUT	Isfahan University of Technology
🟢	JHU	Johns Hopkins University
🟢	JWGU	Goethe University
🟢	KINGS	King's College
🟢	KTH	KTH Royal Institute of Technology
🟢	L&C	Lewis & Clark
🟢	LEEDS	University of Leeds
⚫	LIX^{3rd}	LIX École Polytechnique Paris
🟢	LLNL	Lawrence Livermore National Laboratory
🟢	LU	Lancaster University
⚫	LUH^{2nd}	Leibniz Universität Hannover
🟡	MHS^{2016}	Methacton High School
🟣	MIT	Massachusetts Institute of Technology
⚫	MITLL^{4th}	Massachusetts Inst. of Technology Lincoln Labs
🟢	MS	Microsoft
🟢	MSR	Microsoft Research
⚫	MSR SVC^{1st}	Microsoft Research Silicon Valley
⚫	MU^{3rd}	Macquarie University
🟢	NCSA	National Center for Supercomputing Applications
🔴	NCSU	North Carolina State University
🟢	NEC LABS	NEC Laboratories America
🟢	NETFLIX	Netflix
🔵	NEWCASTLE (UK)	Newcastle University
🟢	NIST	National Institute of Standards & Technology
🟢	NOVA	Universidade Nova de Lisboa
🟢	NRL	Navy Research Laboratory
🟠	NSA	National Security Agency
🟢	NTU	Nanyang Technological University
🟢	NTUA	National Technical University of Athens
🟢	NU	Northeastern University
🟢	NUST	National University of Sciences & Technology
🟢	NWU	Northwestern University
🟢	NYUAD	New York University, Abu Dhabi
🟡	OES^{2015}	Oregon Episcopal School
🟢	OHSU	Oregon Health and Science University
🟢	OSU	Oregon State University
🟢	OVGU	Otto Von Guericke University Madgeburg
🟢	OXFORD	University of Oxford
🟢	PASSAU	University of Passau
🟢	PC	Providence College
🔵	PENN	University of Pennsylvania
🟡	PHS^{2015}	Poolesville High School
🔵	PITT	University of Pittsburgh
🟢	PKU	Peking University
🟢	POLYMTL	Polytechnique Montreal
🟢	PRINCETON	Princeton University
🟢	PSU	Pennsylvania State University
🔵	PU	Purdue University
🟢	QUEENS	Queens University
🔵	RICE	Rice University
🔵	RIT	Rochester Institute of Technology
🟢	ROYAL HOLLOWAY UNIV.	Royal Holloway, University of London
🟢	RPI	Rensselaer Polytechnic Institute
🟢	RSA	RSA Security
🟢	RU	Rutgers University
🟢	RUB	Ruhr University Bochum
🟢	SANDIA	Sandia National Laboratories
🟢	SC	Swarthmore College
🟢	SEI	Software Engineering Institute
🟢	SENTIMETRIX	SentiMetrix
🟢	SIEMENS	Siemens Corporation
🟢	SMU	Singapore Management University
🟢	SRA-DALLAS	Samsung Research America - Dallas
🟢	SRL	Simula Research Laboratory
🟢	SU	Stanford University
🟢	SWANSEA	Swansea University
🟢	SYMANTEC^{1st}	Symantec Corporation
🟢	TAMU	Texas A&M University
🟢	TEMPLE	Temple University
🟢	THU	Tsinghua University
🟢	TODAI	University of Tokyo
🟢	TTU	Tennessee Technical University
🟢	TU	Towson University
🟢	TU BRAUNSCHWEIG	Technical University Braunschweig
🟢	TU DARMSTADT	Technische Universität Darmstadt
🟢	TU KL	University of Kaiserslautern
🟢	TU/e	Eindhoven University of Technology
🔵	UA	University of Alabama
🟢	UALBANY	University of Albany
🟢	UCAS	University of Chinese Academy of Sciences
🟣🔵	UC BERKELEY	University of California, Berkeley
🔵	UCI	University of California, Irvine
🟢	UCL	Université catholique de Louvain
🟢	UCSB	University of California, Santa Barbara
🟢	UCSD	University of California, San Diego
🟢	UDS	Universität des Saarlandes
⚫	UE^{1st}	University of Edinburgh
🟢	UFAL	Universidade Federal de Alagoas
🟢	UFCG	Universidade Federal de Campina Grande
⚫	UFMG^{3rd}	Federal University of Minas Gerais
🟣	UH MĀNOA	University of Hawai'i at Mānoa
🟢	U IDAHO	University of Idaho
🟢	UiO	University of Oslo
🔴	UIUC	University of Illinois at Urbana-Champaign
🟢	UJN	University of Jinan
🟢	ULISBOA	University of Lisbon
⚫	UMASS^{4th}	University of Massachusetts, Amherst
🟢	UMBC	University of Maryland, Baltimore County
🔴	UMD$^{2nd \& 4th}$	University of Maryland
🟢	UMICH	University of Michigan
🔵	UNC^{4th}	University of North Carolina at Chapel Hill
🔵	UNCC	University of North Carolina at Charlotte
🟢	UNIMELB	University of Melbourne
🟢	UNIS	University of Surrey
🟢	UNIVR	University of Verona
🔵	UNL	University of Nebraska Lincoln
🟢	UNM	University of New Mexico
⚫	UNSW^{3rd}	University of New South Wales
🟢	UOFW	University of Waterloo
🟢	UON	University of Nottingham
🟢	UP	Universidade de Porto
🟢	UPB	University Politehnica of Bucharest
🟢	UPM	Universdade Politécnica de Madrid
🟢	UPV	Universitat Politècnica de València
🔵	USC	University of Southern California
🟢	USI	Università della Svizzera italiana
🟢	USMA	United States Military Academy
🔵	USNA	United States Naval Academy
🟢	USP	University of São Paulo
🟢	UT	University of Twente
🔵	UT AUSTIN	University of Texas at Austin
🟢	UT DALLAS	University of Texas at Dallas
🟢	UTA	University of Texas at Arlington
🔵	UTSA	University of Texas at San Antonio
🔵	UVA	University of Virginia
🟢	UW	University of Washington
🟢	UW-MADISON	University of Wisconsin
🟢	UWAR	University of Warwick
🟢	UWO	Western University
🟢	VERISIGN	Verisign Labs
🟢	VIATEC	Viatec Research
🟡	VPHS^{2016}	Villa Park High School
🔵	VT	Virginia Polytechnic Institute and State University
🟣	VU	Vanderbilt University
🟢	VUW	Victoria University Wellington
🟢	WANDOUJIA	Wandoujia Lab
🟡	WHS^{2016}	Westwood High School
🟢	WRUT	Worcław University of Technology
🔵	WSU	Wayne State University
🟢	ZIU	Zhejiang University

🔴 Lablet (4)

🔵 Sub-Lablet (25)

🟣 SURE (4)

🟢 Collaborator (138)

⚫ Paper Competition (17)
SUPERSCRIPT INDICATES YEAR

○ 2017 Additions (36)

🟡 ISEF Winners (6)
SUPERSCRIPT INDICATES YEAR

— CMU
— NCSU
— UIUC
— UMD

Privacy Research Network

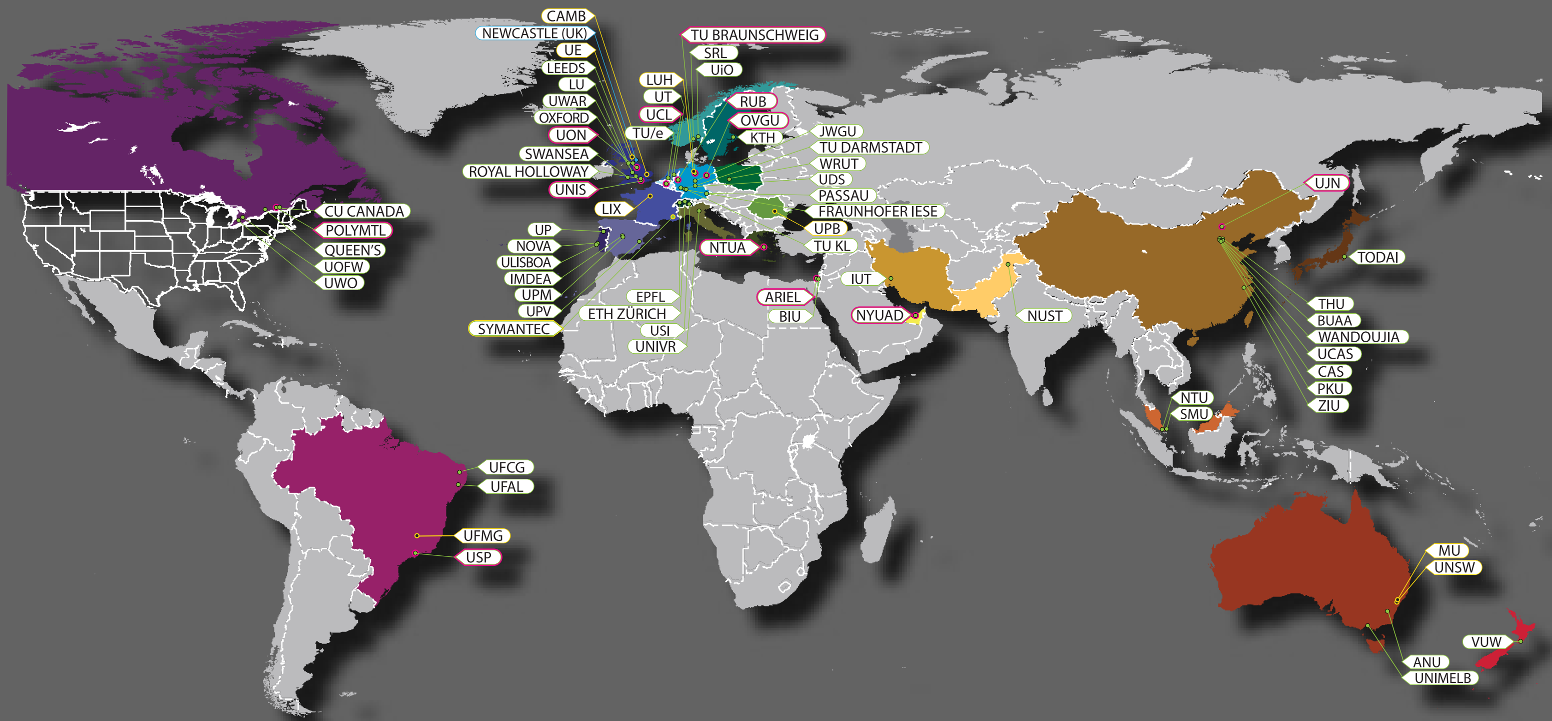

🔴 Lablet (4)

- CMU
- NCSU
- UIUC
- UMD

🔵 Sub-Lablet (25)

🟣 SURE (4)

🟢 Collaborator (138)

⚫ Paper Competition (17)
SUPERSCRIPT INDICATES YEAR

○ 2017 Additions (36)

🟡 ISEF Winners (6)
SUPERSCRIPT INDICATES YEAR

The Science of Security & Privacy Initiative at the National Security Agency Research Directorate promotes foundational cybersecurity and privacy science that is needed to mature the cybersecurity discipline and to underpin advances in cyberdefense. Beginning in 2012, one part of the initiative is to fund foundational research at "Lablets." With emphasis on building a community, each Lablet created partnerships with other universities called "Sub-Lablets." Science of Security researchers often freely collaborated with researchers in other institutions worldwide. In 2014, the SURE project was founded to investigate cybersecurity in the cyber-physical systems realm. Since 2012, the initiative also promotes rigorous research methods through its annual best **paper competition** by highlighting the best example of scientific cybersecurity research contribution. To also promote rigorous research at the high school level, SoS started recognizing outstanding achievement with a cash award at the Intel International Science and Engineering Fair (ISEF) since 2015. This map illustrates the expansion of the Science of Security & Privacy Initiative community.

- An approach that gathers evidence pointing to possible insider attacks based on the interaction of sequences of user interactions with resilient architectural behaviors. An evaluation showed that this approach could perfectly rank anomalous traces on an architecture with 240 components and 10% anomalous behaviors, and could maintain better than 95% perfect ranking of anomalous traces with up to 40% of simulated traces as anomalies.
- A study of how users' security beliefs, knowledge, and demographics correlate with their sources of security advice, as well as how these factors influence their security behaviors. Researchers produced results that can be used to improve the delivery of security advice in order to increase the likelihood of uptake of important security behaviors.
- An experiment to determine whether the effect of situational deterring cues in an attacked computer system influenced the likelihood of system trespassers engaging in active online behaviors on an attacked system, and whether this effect varies based on different levels of administrative privileges taken by system trespassers. Researchers found that a situational deterring cue reduced the probability of system trespassers with fewer privileges on the attacked computer system (non-administrative users) to enter activity commands. In contrast, the presence of these cues in the attacked system did not affect the probability of system trespassers with the highest level of privileges (administrative users) to enter these commands.
- A study of how to gather useful data about security-related user behaviors when accessing websites. Based on data from actual password use, researchers discovered that partial password reuse (e.g., a common stem) is rampant, which suggests potential security risks that go beyond reuse of entire passwords.

The SURE program (SecUrity and REsilience for Cyber-Physical Systems) is an additional research effort sponsored by the SoS initiative to develop foundations and tools for designing, building, and assuring cyber-physical systems that can maintain essential system properties in the presence of adversaries. SURE is led by Vanderbilt University and includes researchers at the Massachusetts Institute of Technology, the University of California, Berkeley, and the University of Hawaii.

In pursuit of its second goal, promoting rigorous scientific principles, the SoS initiative sponsored the 5th Annual Best Scientific Cybersecurity Paper Competition, where one paper was recognized as an outstanding scientific paper in cybersecurity research. The SoS initiative also sponsored awards at the 2017 Intel International Science and Engineering Fair (ISEF) to recognize the scientific achievements of high school students in the field of cybersecurity.

In addition to engaging with the academic community and promoting rigorous scientific research, the SoS initiative used other tools to expand the community to close to 200 institutions, 1300 Virtual Organization (VO) members, and over 700 published authors. The 2017 Hot Topics in the Science of Security: Symposium and Bootcamp (HoTSoS) provided a forum for researchers from many different disciplines to present new refereed papers and discuss scientific approaches to cybersecurity. The SoS VO (www.sos-vo.org) continued as the centralized, online location for researchers and all interested parties to engage in discussion and access the most current research in cybersecurity. NSA Research Directorate personnel, Lablet and Sub-Lablet researchers, and VO members all serve as ambassadors to promote a Science of Security.

This year, the title of Science of Security initiative transitioned to the Science of Security and Privacy initiative to reinforce the critical nature of privacy research.

2017 marked the conclusion of Lablet research funded under the most recent Broad Area Announcement (BAA). The three years of research by the four Lablets produced 330 publications and 59 tools; a total of 567 papers have been published since the SoS initiative began.

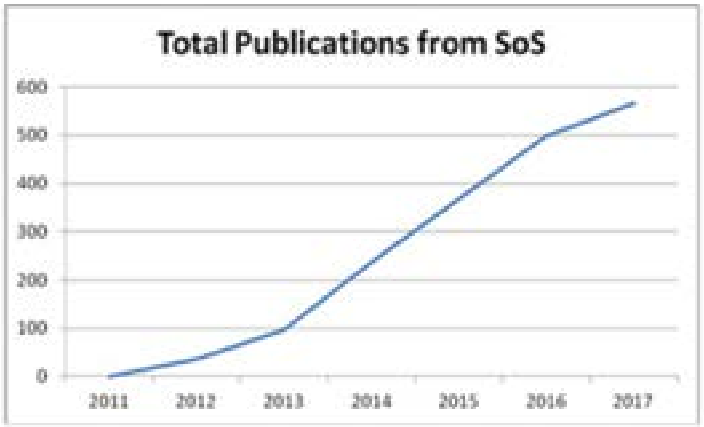

In 2018, the third generation of SoS Lablets will focus on twenty specific projects that address some of the most significant cybersecurity research challenges aligned against the five Hard Problems. The Lablets are Carnegie Mellon University, the International Computer Science Institute, North Carolina State University, the University of Illinois at Urbana-Champaign, the University of Kansas, and Vanderbilt University.

In 2018, the Science of Security and Privacy initiative will continue to engage with academia in foundational research as described above, promote rigorous scientific principles through the 6th Annual Best Scientific Cybersecurity Paper Competition and sponsorship of ISEF awards, and grow the community through its use of multiple outreach endeavors. Accomplishment of these goals will continue to raise awareness of the need for foundational cybersecurity science to mature the cybersecurity discipline and to underpin advances in cyberdefense.

PRIVACY

The Science of Security (SoS) initiative has changed its name to the Science of Security and Privacy initiative to reinforce the critical nature of privacy research. While privacy issues have long been considered in SoS research, privacy is significant enough that it deserves to be addressed separately. Privacy, which considers how authorized access to information is achieved, requires a trustworthy infrastructure to be assured, and a major tenet of providing trustworthy infrastructure is by developing a Science of Security. Only if the infrastructure can be trusted can privacy be assured. In 2017, the Science of Security and Privacy initiative continued to develop principles that are building the foundations for a Science of Privacy.

Privacy has been part of SoS since its inception and was prominent when the Science of Security Hard Problems were first developed. The five Hard Problems provided a clearly stated and well-scoped problem set to focus SoS research. Privacy is a major component of the Hard Problem of Policy Governed Secure Collaboration, since it focuses on authorized sharing of information and thus includes the research needed to share information. The two major components of research in this Hard Problem are the technical element of how to share information securely and the policy element of when and how to share information. Privacy consideration is an integral part of the second element, deciding on the when and the how.

There are four privacy projects under way at two of the SoS Lablets. The North Carolina State University (NCSU) Lablet has a project to build the first comprehensive encyclopedia and database of privacy incidents. Although there are databases on data breaches, privacy incidents cover a wider range of issues. In this research project, a privacy incident is defined as one of the following: an instance of accidental or unauthorized collection; use or exposure of sensitive information about an individual or an event that creates the perception that unauthorized collection occurred; use or exposure of sensitive information about an individual that may occur; and the belief that information is either being collected, used, or shared in digital form. Most privacy incidents that fall into a category identified above (including cyber-bullying/slander/stalking, revenge porn, social media oversharing, data reidentification and surveillance) do not involve a security breach. Therefore, such incidents are not represented in current security incident databases. This lack of a centralized resource leads to widely varying measurements of privacy incidents. NCSU's publicly accessible database, launched in 2017, will enable the privacy technology and policy communities to reach consensus around patterns in privacy incidents. It includes, for example, a company misusing data and data that hasn't been disclosed. It is available at https://sites.google.com/site/privacyincidentsdatabase/

A second project at NCSU is a study of norms in information sharing and how to reason about these norms to make data flow decisions. This research can aid in consistent and reality-based policy decisions. This year, researchers developed an approach that helps maintain alignment of commitments (as an exemplary social construct) between autonomous parties communicating over an asynchronous medium. Their research contributes a formal grounding to Sociotechnical Systems (STS) for cybersecurity and privacy that leads to new methods for architecting security and privacy into a software system.

Researchers at the Carnegie Mellon University (CMU) Lablet and its Sub-Lablet the University of Texas, San Antonio undertook work on a framework to support holistic analysis of systems for privacy. There are four interrelated scientific thrusts: 1) formal languages, based on natural language privacy policies, designed to express privacy principles and check data flow for compliance in a system's architecture; 2) information type ontology theory to enable algorithmic checks of whether meanings assigned to predicates in a formal privacy language are consistent among policy makers and developers; 3) static and dynamic analysis techniques to find privacy violations by identifying and tracing data collection, use, and transfer practices throughout source code; and 4) privacy risk metrics to reliably measure how users and data subjects perceive the risk of specific data practices, so that developers can focus effort on mitigating the most important privacy threats. Privacy differs from security because privacy is characteristically about the data subject, or the person about whom the data is descriptive. One of the goals of this work is to trigger a new line of design thinking that is not limited to access control and anonymization, and which can introduce consent and use-based restrictions imposed by architecture.

A second project at CMU developed an empirically validated framework to measure perceived privacy risk. The research found that while individuals can perceive increased risk with increased likelihood, the contribution to overall risk perception is sub-linear: there are greater perceived differences among the risks of sharing different information types than the differences due solely to increased likelihood of a privacy harm for a single information type. Moreover, the research shows that individuals are more willing to share information about what they do than they are willing to share information about who they are. This indicates that privacy risk may increase non-linearly when identifiable information is combined with sensitive information types. With respect to scalability, researchers are currently investigating techniques to scale the information type ontology, to investigate the effect of data aggregation, and to identify cost-effective ways to re-sample privacy risk measures from individuals. This framework can be applied to sharing of cybersecurity information to respond to cyber threats.

Privacy concerns continue to grow. There are multiple published concerns about the collection of sensitive information and the potential threat of the information being misused, whether by technology companies or government agencies. The Science of Security and Privacy initiative is seeking to address such concerns through a combination of security and privacy research. Just as has been the case for a science of security, increased scientific understanding of privacy is critically important.

Privacy research will become more integrated with Science of Security research in the future. Historically the Science of Security Lablets focused on the five Hard Problems; in 2018 two additional Lablets will be especially focused on Privacy and Cyber Physical Systems. The Science of Security Lablet focused on Privacy will strengthen the scientific foundations for a Science of Privacy. All Lablets, regardless of their focus, will be charged with addressing privacy concerns as part of their research. Progress on privacy research will be highlighted in the next annual report.

Section 1

Foundational Research

Engaging the Academic Community for Foundational Research

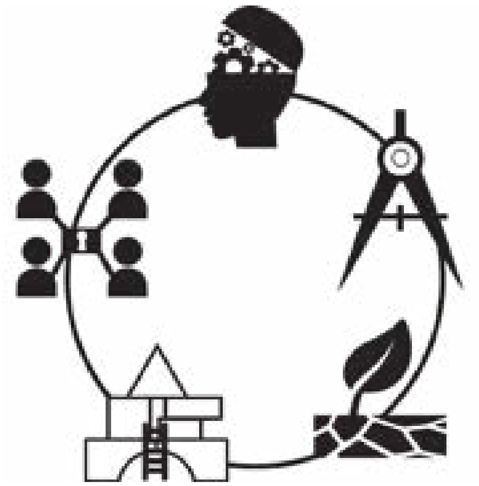

In 2017 the Science of Security initiative continued to focus its academic community engagement on the four Science of Security Lablets and their 25 Sub-Lablets. Since 2012, the Lablets and Sub-Lablets have published over 550 peer-reviewed papers and developed almost 60 tools. The papers have addressed multiple aspects of the five Hard Problems, and have been presented at conferences, symposia, and workshops around the world. The tools, which are displayed in the following section, address both design and analysis aspects of cybersecurity, and a number have been released to GitHub and to the public. The foundational research embodied by the papers and tools have contributed immeasurably 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 as a means of establishing challenging and critical research goals and to serve as the beginnings of a common language and a way to assess progress in foundational SoS research. The papers published over the past year provide tangible evidence of the impact that the Lablets' research has had on improving the Science of Security in the five Hard Problem focus areas.

Lablet foundational research is organized by projects which relate to one or more of the five Hard Problems. In 2017 there were 42 Lablet research projects that produced 69 publications. The Lablets also engage regularly in community outreach, participate in international conferences and workshops, and integrate Science of Security principles into their curricula. The Science of Security also sponsors the Science of SecUrity and Resilience for Cyber-Physical Systems (SURE) project which involves another four universities. SURE researchers have published 53 papers over the life of the SURE project, 21 of which were released this year. In addition to providing quarterly and annual reports on their activities, the four Lablets, along with Sublablets, collaborators, and NSA researchers, meet quarterly to present updates on their research projects and exchange information about issues related to Science of Security.

The progress on Hard Problems, Lablet activity, Lablet quarterly meetings, and the SURE project are detailed in this section.

Science of Security Lablets Progress on Hard Problems 2017

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 problems are: (1) Scalability and Composability; (2) Policy-Governed Secure Collaboration; (3) Security Metrics and Models; (4) Resilient Architectures; and (5) Human Behavior.

- 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.
- 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.
- Security Metrics and Models addresses the measurement of properties relevant to cybersecurity, and quantifying the degree to which a system satisfies those properties.
- 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.
- Human Behavior addresses how to handle the unpredictability of human actors in cybersecurity.

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. The updates below provide tangible evidence of the impact that the Lablets' research has had on improving the Science of Security in these five focus areas.

Scalability and Composability

The Hard Problem of Scalability and Composability 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.

Goal: 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.

Progress in 2017

In order to advance the study of Scalability and Composability, we need to provide practicable language support for least-privilege secure systems design. We designed, implemented, and achieved preliminary validation of a module system that is capability-safe, yet preserves most of the convenience of conventional module systems. Restricting the authority of untrusted code is difficult in today's systems because, by default, code has the same authority as the user running it. Object capabilities were a promising way to implement the principle of least authority, but previous approaches take away many conveniences provided by today's module systems. We developed a design that preserves those conveniences, paving the way for object capabilities as a practicable approach to software security. We demonstrated how to ensure key security and privacy properties of a program, using our module capabilities to reason about properties of the system as a consequence of the properties of individual modules. In a second, related, contribution, we addressed the problem of adding reflection to the language in a way that preserves the capability safety of the module system. Reflection is important for many applications that use dynamic extension and configuration mechanisms, but has been a source of security holes in Java, and no prior work has described how to make it capability safe. We used language design methods to provide a safe form of reflection in which the reflective mirror of an object provides exactly the same access as a typed capability would in the non-reflective, base system. This enables the construction of applications that are flexible and dynamic in structure, yet benefit from the compositional security guarantees of a capability system.

- Darya Melicher, Yangqingwei Shi, Alex Potanin, and Jonathan Aldrich, "A Capability-Based Module System for Authority Control," in *Proceedings, 31st International European Conference on Object-Oriented Programming (ECOOP)*, Barcelona, Spain, June 18 – 23, 2017. https://cps-vo.org/node/36376
- Esther Wang and Jonathan Aldrich, "Capability Safe Reflection for the Wyvern Language," in *Proceedings of the Workshop on Meta-Programming Techniques*

- *and Reflection (META)*, part of SPLASH, Amsterdam, Netherlands, October 30 – November 4, 2016. https://cps-vo.org/node/31395

We addressed the problem of compositional reasoning about properties of programs as they are being constructed. In this work, we developed a formal model of a program that is under construction, so that parts of it are implemented and typechecked, and parts of it are missing. Our formal model defines a type system that can check the parts of the system that are fully implemented, using appropriate typing assumptions for the parts of the system that are still being implemented. It also models programmer actions (e.g. in an integrated development environment) that implement the missing code, so that if the code being added conforms to the expected type, the overall system is guaranteed to typecheck.

- Cyrus Omar, Ian Voysey, Michael Hilton, Jonathan Aldrich, and Matthew Hammer, "Hazelnut: A Bidirectionally Typed Structure Editor Calculus," *Symposium on Principles of Programming Languages (POPL)*, Paris, France, January 15 – 21, 2017. https://cps-vo.org/node/34446
- Ian Voysey, Cyrus Omar, and Matthew Hammer, "Running Incomplete Programs," *Part of POPL, Off the Beaten Track Workshop (OBT)*, Paris, France, January 15 – 21, 2017. https://cps-vo.org/node/34448
- Cyrus Omar, Ian Voysey, Michael Hilton, Joshua Sunshine, Claire Le Goues, Jonathan Aldrich, and Matthew A. Hammer, "Toward Semantic Foundations for Program Editors," *2nd Summit on Advances in Programming Languages (SNAPL)*, Asilomar, CA, May 7 – 10, 2017. https://cps-vo.org/node/36379

We addressed the problem of security vulnerabilities that result from the interaction of multiple applications on a platform such as Android. Prior work on this problem was impractical due to the inability to scale to the numbers of applications that need to be considered. Our approach, encapsulated in the tool JITANA, applies program analysis at classloading time, rather than compilation time, allowing it to analyze multiple interacting applications on demand, including any dynamically generated code. Empirical studies of JITANA show that it is 4 to 30 times more efficient (depending on the scenario) than a state-of-the-art approach, and that it can effectively and efficiently analyze complex apps including Facebook, PokemonGo, and Pandora.

- Yutaka Tsutano, Shakthi Bachala, Witawas Srisa-an, Gregg Rothermel, and Jackson Dinh, "An Efficient, Robust, and Scalable Approach for Analyzing Interacting Android Apps," *39th International Conference on Software Engineering (ICSE)*, Buenos Aires, Argentina, May 20 – 28, 2017. https://cps-vo.org/node/36382

We developed and evaluated a compositional approach to "deep immutability" that addresses the frequent mismatch of programmer expectations when variables are declared to be constant or final: the reference itself might be guaranteed not to change, but the object referenced may nonetheless change. This "expectation gap" is the source of a range of escalation-of-privilege vulnerabilities, such as the Java 1.1 getSigners() method that shares an "immutable" reference to an array of signers that turns out to be mutable. Transitive immutability is a composition issue, since expectations must be expressed and analyzed across component boundaries. We applied type theory to design a system for verifying a semantically strong notion of immutability, then validated its usability (touching on the Human Behavior Hard Problem) with human-subjects studies demonstrating success rates that improved from a plain-Java control baseline of 0-40% (depending on task) to 70-90% success rates with our intervention. We also examined existing code bases to assess the extent of change required to model and analyze deep immutability, with results indicating very few minor changes were needed.

- Michael Coblenz, Whitney Nelsony, Jonathan Aldrich, Brad Myers, and Joshua Sunshine, "Glacier: Transitive Class Immutability for Java," *39th International Conference on Software Engineering ICSE)*, Buenos Aires, Argentina, May 20 – 28, 2017. https://cps-vo.org/node/36380

We addressed the problem of enforcing security practices in standardized software certification processes, such as Common Criteria. These certification approaches would ideally be compositional, allowing for certifications of systems to rely on certifications of their parts, and facilitating recertification after components are modified. We used the method of interviewing subject matter experts, and identified a number of tensions in the certification process, emphasizing that rapid recertification and composition issues are among the main concerns in today's certification. The effect of these issues is that old, vulnerable, (but certified) versions are often preferred to newer patched ones that are still going through the lengthy certification process. The studies confirm that a larger focus on rapid and compositional certification is critical to address security concerns for software systems, including those that are highly configurable.

- Gabriel Ferreira, "Software Certification in Practice: How Are Standards Being Applied?," in *Proceedings of the 2017 IEEE/ACM 39th International Conference on Software Engineering Companion (ICSE)*, Buenos Aires, Argentina, May 20 – 28, 2017. https://cps-vo.org/node/36437

We addressed the problem of security issues that arise from system evolution and configuration. While some versions and configurations of a system might be secure, changing a single configuration option, updating a single dependency, or combining multiple configuration options in unanticipated ways can expose new vulnerabilities. Heartbleed, for example, affected only users who activated the (default, but not required) heartbeat configuration option. In the Linux kernel there are 15K compile-time configuration settings, with the total number of configurations (a theoretical maximum of 2^15k) being intractably huge. We developed a new method to simultaneously analyze very large numbers of compile-time configurations of C code at once, without looking at each configuration individually, a kind of configuration-aware taint tracking. We analyzed which parts of the code might suffer from increased configuration complexity (e.g., many #ifdefs), and applied the ideas directly to examples such as the Linux kernel and OpenSSL.

We addressed the problem of how to predict which code fragments within a system such as the Linux kernel are most likely to contain security vulnerabilities. We applied network analysis techniques to revision data from software repositories to identify unusual commits that are outliers and merit security review. In an empirical evaluation, we found that our anomaly detection mechanism detected code commits that are actionable for developers to prioritize reviews, enabling the security review process to scale up more effectively.

- Rohit Goyal, Gabriel Ferreira, Christian Kästner, and James Herbsleb, "Identifying Unusual Commits on GitHub," *Journal of Software: Evolution and Process (JSEP)*, 2017. https://cps-vo.org/node/36439

We addressed the problem of determining and enforcing least-privilege architecture for Android, enabling a compositional approach to security. Our approach, embodied in the DELDroid tool, uses static analysis to identify the least-privilege level actually needed by an application, so this can be enforced at runtime, thus reducing security risks.

- Mahmoud Hammad, Hamid Bagheri, and Sam Malek, "DELDroid: Determination and Enforcement of Least-Privilege Architecture in Android," *IEEE International Conference on Software Architecture (ICSA 2017)*, Gothenburg, Sweden, April 3 – 7, 2017. https://cps-vo.org/node/36402

We addressed the problem of organizing existing research in the area of composable program analysis techniques for assessing the security of Android software. Applying a systematic literature review process revealed patterns, trends, and gaps in the existing literature, and underlined key challenges and opportunities that will shape the focus of future research efforts.

- Alireza Sadeghi, Hamid Bagheri, Joshua Garcia, and Sam Malek, "A Taxonomy and Qualitative Comparison of Program Analysis Techniques for Security Assessment of Android Software," *IEEE Transactions on Software Engineering*, Vol: 43(6), June 1, 2017. https://cps-vo.org/node/36405

We addressed the problem of efficiently and scalably classifying Android applications as benign or suspicious. Using a permissions analysis approach with pruning, we identified that only 22 out of 135 Android permissions are relevant for classification. Using a classifier based on these permissions, our tool was able to detect 93.62% of malware in our data set (which is comparable to a baseline approach), but with analysis times reduced by 4 to 32 times compared to the less scalable baseline.

We addressed the problem of effectively detecting data races and other concurrency faults, using a combination of dynamic analysis and virtualization to better control and observe system execution. The approach was able to reduce false negatives by 53%, and reported 77% of faults compared to 18% of faults detected by a stress testing baseline technique.

- Tingting Yu, Witty Srisa-an, and Gregg Rothermel, "SimExplorer: An Automated Framework to Support Testing for System-Level Race Conditions," *Journal of Software: Testing, Verification and Reliability*, Vol 27, 4-5. May 10, 2017. https://cps-vo.org/node/36387

Policy-Governed Secure Collaboration

The Hard Problem of Policy-Governed Secure Collaboration 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.

Goal: 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.

Progress in 2017

We investigated the specific problem of reducing errors in policy specification by reducing the complexity of policy expressions. Its specific method involved the development of a language that supports the tri-modular representation of firewall policies coupled with analysis of policies expressed in that language. The method enabled the specification of policies with fewer errors than otherwise possible. In particular, real-world security policies, such as firewall policies and SEAndroid policies, are often large and complex, and their specification is error-prone. Thus, errors in policies undermine sophisticated techniques for verifying and executing policies.

- Haining Chen, Ninghui Li, William Enck, Yousra Aafer, and Xiangyu Zhang, "Analysis of SEAndroid Policies: Combining MAC and DAC in Android," *Proceedings of the Annual Computer Security Application Conference (ACSAC)*, 2017. https://cps-vo.org/node/38108

We investigated the specific problem of how to extract knowledge from pre-deployment software artifacts (e.g., regulations) and post-deployment artifacts (e.g., breach reports) to improve security and privacy requirements. Its specific method was the analysis of text in regulations and breach reports to reason about how well existing regulations (represented as norms) would have prevented each specified breach (treated as a norm violation). Understanding whether the gap is in the norms or in how well they are applied can help improve the requirements to which software systems are built. This research helps reduce security vulnerabilities that are introduced through poor software engineering at the stages of requirements elicitation and development.

- Özgür Kafali, Jasmine Jones, Megan Petruso, Laurie Williams, and Munindar P. Singh, "How Good is a Security Policy against Real Breaches? A HIPAA Case Study," *International Conference on Software Engineering*, 2017. https://cps-vo.org/node/31257

We investigated the specific problem of instantiating a sociotechnical architecture based on norms and mechanisms so as to apply norms and policies consistently across authority domains. Its specific method is the development of an asynchronous communication protocol based on an information model that captures key and causal dependencies. This method achieved a limited form of consistency called *aligned*: aligned does not require complete agreement, which would be impossible to realize over asynchronous communication, but respects agreement along the direction of accountability: any norm (here commitment) that is active with its object (approximately, beneficiary) is active at the accountable party. This research contributes to a formal means to ground sociotechnical systems in which security and privacy are built in as norms.

- Thomas Christopher King, Akın Günay, Amit K. Chopra, and Munindar P. Singh, "Tosca: Operationalizing Commitments Over Information Protocols," *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, Melbourne, pages 256-264. 2017. https://cps-vo.org/node/36114

We investigated the specific problem of supporting reproducible empirical research into privacy attitudes and incidents. Its specific method is the development of the Privacy Incidents Database, (https://sites.google.com/site/privacyincidentsdatabase) which incorporates tools for reducing the human effort essential for maintaining the database. One tool suite helps identify privacy incidents from news media and social media reports (specifically, articles and tweets). Another tool suite helps visualize the chronology, entities, and location of privacy incidents to help users understand database entries. The database includes seed data based on news media articles relating to privacy incidents but is meant to grow through community contributions. The main benefit of the database would be in improving understanding of privacy risks and attitudes to help both software developers and members of the public.

- Karthik Sheshadri, Nirav Ajmeri, and Jessica Staddon, "No (Privacy) News is Good News: An Analysis of New York Times and Guardian Privacy News from 2010 to 2016," *Proceedings of 15th Annual Conference on Privacy, Security and Trust (PST)*, August 2017, Calgary, pages 1-12. https://cps-vo.org/node/38299

We investigated the specific problem of adoption of security tools during software development, including what incentives and other interventions may promote adoption. Its specific method was multiagent simulation of a simplified software development organization, including managers and developers, assumed to be rational players. The players are subject to differing incentives for timeliness, functionality, and security quality (vulnerability reduction) of the delivered software. The method helps understand the tradeoffs between the incentives that lead to software practices that promote or undermine security by complying with or violating norms and policies in software development. In a related effort, we are in the process of evaluating the predictions based on training, and evaluating a model based on a publicly available bug reports dataset.

Security Metrics and Models

The Hard Problem of Security Metrics and Models 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.

Goal: 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.

Progress in 2017

Given that one cannot sample and analyze everything, a natural question asks what one should sample, where one should sample it, and how often it needs to be sampled.

- John C. Mace, Nippun Thekkummal, Charles Morisset, and Aad van Moorsel, "ADaCS: A tool for Analysing Data Collection Strategies", *14th European Performance Engineering Workshop (EPEW 2017)*, Berlin, Germany, September 7-8, 2017. https://cps-vo.org/node/39034

One means of quantifying the effectiveness of an attack detection technique is to measure that performance against known attacks. To increase the diversity of attacks so one may use the output, we develop techniques for creating variations of known attacks and then measure detection effectiveness against the library of attacks founded in realism.

Cyber-physical systems create new problems in metrics definition and measurement. We explore a metric-based approach to synthesizing control laws in cyber-physical systems under the assumption that adversaries are attempting to disrupt the system.

Privacy (in the sense of not leaking sensitive information) is another problem made challenging in the context of cyber-physical systems. We have performed ground-breaking work on developing differential privacy mechanisms with an idea towards the tradeoffs between privacy and performance.

- Yu Wang, Zhenqi Huang, Sayan Mitra, and Geir Dullerud, "Differential Privacy in Linear Distributed Control Systems: Entropy Minimizing Mechanisms and Performance Tradeoffs", I*EEE Transactions on Control of Network Systems*, volume 4, issue 1, January 25, 2017. https://cps-vo.org/node/34798
- Zhenqi Huang, Yu Wang, Sayan Mitra, and Geir Dullerud. "Differential Privacy and Entropy in Distributed Feedback Systems: Minimizing Mechanisms and Performance Trade-offs," *IEEE Transactions on Network Control Systems* , volume 4, issue 1, March 2017. https://cps-vo.org/node/36593
- Yu Wang, Sayan Mitra, and Gier Dullerud, "Differential Privacy and Minimum Variance Unbiased Estimation in Multi-agent Control Systems," *20th World Congress of the International Federations of Automatic Control (IFAC 2017 World Congress)*, Toulouse, France, July 9-14, 2017. https://cps-vo.org/node/39064
- Yu Wang, Matthew Hale, Magnus Egerstedt, and Gier Dullerud, "Differentially Private Objective Functions in Distributed Cloud-based Optimization", *55th Conference on Decision and Control (CDC 2016)* , Las Vegas, NV, December 12-14, 2016. https://cps-vo.org/node/39035

Privacy in the context of messaging systems also depends on metrics describing that privacy.

- Giulia Fanti, Shaileshh Venkatakrishnan, and Pramod Viswanath, "Dandelion: Redesigning BitCoin Networking for Anonymity," *ACM SIGMETRICS* , Urbana, IL, June 5-9, 2017. https://cps-vo.org/node/36572

We investigated the specific problem of incorporating models of attacker behavior in the attacker's vulnerability discovery process. Its specific method involved the development of the attack surface "risky walk" metric based on the call graph of a system. This metric identifies the components that have high likelihood of containing a vulnerability based on a random walk on a system, beginning from a system's entry points (its attack surface) and traversing all of it, to simulate an attacker's vulnerability discovery process. This research thus helps identify vulnerabilities.

We investigated the specific problem of prioritizing security validation and verification efforts to those areas on the attack surface of a software system. Its specific method involved the development of metrics used to approximate the attack surface through an analysis of the artifacts (files, binaries) found on stack traces from system crash dumps. We call this method the risk-based attack surface approximation because the method identifies the artifacts that have high likelihood of containing a vulnerability that has been shown to be accessible through an event caused by a user/using system. This research thus helps identify vulnerabilities.

- Christopher Theisen, Brendan Murphy, Kim Herzig, and Laurie Williams, "Risk-Based Attack Surface Approximation: How Much Data is Enough?," *International Conference on Software Engineering (ICSE) Software Engineering in Practice (SEIP),* Buenos Aires, Argentina, pp. 273-282, 2017. https://cps-vo.org/node/37574

We investigated the specific problem of effectively configuring and sharing resources among collaborative intrusion detection systems (IDS). Its specific method consisted of a two-layer stochastic game-theoretic model, in which the reward function for each IDS depends upon the probability of attack detection and the corresponding benefit. In a case study of ten collaborative IDSs, we saw a performance improvement of 120%. This research addresses the hard problem by enabling efficient configuration and effective collaboration among IDS.

We investigated the specific problem of understanding the nature of intrusion detection and how it may be formalized through suitable metrics. Its specific method was to examine the consistency of accuracy, time, and space metrics pertinent to intrusion detection, as studied in the literature. This research establishes empirical validation metrics to capture evidence for the quality of a given intrusion detection technique.

We investigated the specific problem of using known behavior of a healthy system to proactively identify abnormalities and initiate countermeasures for mitigation and resolution. Its specific method concerned metrics of detector sensitivity and the resulting confidentiality, integrity, and availability preservation in the context of low intensity stealth attacks and zero-day attacks. It investigates the practicality of cloud-based and localized run-time back-to-back and acceptance testing for development of "soft" and re-configurable attack sensors, and the corresponding run-time mitigation models based on redundancy. It models a system as a workflow whose inputs and outputs are attack and/or policy violation sensors and are validated before data are allowed to enter or exit. This research helps derive conformance measures with respect to pre-established security properties.

One group of researchers explored techniques for assessing the security of systems in their deployment environments. They formalized several security metrics derived from field data, including the count of vulnerabilities exploited and the size of the attack surface actually exercised in real-world attacks, and showed that improvements in these metrics are often associated with the introduction of system security technologies (e.g. sandboxing), thus reflecting the effectiveness of these technologies in the field. Building on these initial results, they created a machine-learning model to forecast which vulnerabilities will be exploited in the wild, and showed that making accurate predictions required consideration of information diffusion, derived from vulnerability-related discourse on social media.

Resilient Architectures

The Hard Problem of Resilient Architectures 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.

Goal: Develop means to design and analyze system architectures that deliver required service in the face of compromised components.

Progress in 2017

We addressed the specific problem of understanding how to quantify the strength of resilience techniques with respect to a mission and creating the most appropriate resilient architecture considering both benefit and cost. Its specific method was based on new metrics for measuring the effectiveness of resilience techniques, such as isolation, diversity, and automated response, along with a formal framework to apply those metrics to synthesize resilient architectures based on mission requirements. This method increased resistance and detectability by using the isolation appropriate for networks, services, and flows. The method increased service and system diversity to make it costly for attackers to compromise a service. The method deceives and confuses the attacker through a moving target defense.

We addressed the specific problem of scaling Network Intrusion Prevention and Detection Systems (NIDS/NIPS) to growing and varying traffic demands and policy-enforcement responsibilities. Its specific method involved developing abstractions to simplify specification of resource-optimization and policy-enforcement goals in networks as well as methods for translating such specifications into near-optimal solutions fast enough to adapt to network changes, and deploying those solutions with a focus on Software Defined Networking (SDN) contexts. These capabilities simplify and accelerate the specification and deployment of novel resource-management and policy-enforcement capabilities in SDN networks.

We addressed the specific problem of enhancing security isolation, a fundamental building block for resilient architectures, to work against intelligent adversaries. Its specific method concerned identifying principles of policies, mechanisms, and metrics underlying adaptive and dynamic security isolation within hosts and distributed systems. It identified and explored primitives, such as vulnerability propagation pattern extraction and adaptive information flow control, to enable dynamic and adaptive security isolation within large-scale distributed systems. This research has enhanced understanding of building resilient architectures that withstand attacks, a major practical need for resilience.

- Rui Shu, Xiaohui Gu, William Enck, "A Study of Security Vulnerabilities on Docker Hub," *Proceedings of the ACM Conference on Data and Application Security and Privacy (CODASPY)*, 2017. https://cps-vo.org/node/34161

We addressed the specific problem of scalability of automated techniques for synthesizing resilient architectures. Its specific method concerned developing formal models of cyber and cyber-physical systems, and defining their resilience properties based on Satisfiability Modulo Theories. This research increases resilience by increasing isolation in various dimensions of the system design through model decomposition and refinement.

We addressed the specific problem of verifying that policy rules for resilience are conflict-free and that such rules can be executed simultaneously. Its specific method concerned defining a formal language and a controller to represent and implement reactive resilience policies. This research increases resilience by ensuring that the appropriate defensive course of action is initiated timely and accurately.

We addressed the specific problem of security fault elimination based on an analysis of discovery rates for security problems. Its specific method concerned the application of enhancements of the classical Software Reliability Engineering (SRE) methodology in combination with safety-oriented fault management processes. Our results indicate that generation of a repeatable automated test-strategy that explicitly covers the "top 25" security problems (epistemic errors) may help considerably by eliminating as many as 50% of the field observable security problems. This research can help improve resilience through a combination of security fault-elimination and run-time fault tolerance techniques.

- Donghoon Kim, Henry E. Schaffer, Mladen A. Vouk, "About PaaS Security," *International Journal of Cloud Computing*, 2017. https://cps-vo.org/node/13263

Network resiliency is a critical component of architectural resiliency. Software defined networking offers significant advantages in that one can craft a connection architecture that is resilient to faults, and develop schemes for reacting dynamically to changing network conditions so as to maintain network functionality.

Part of ensuring network resiliency is ensuring that SDN changes to the flow rules do not exceed pre-identified limits on congestion permitted on a per link basis.

- Xin Liu and Dong Jin, "ConVenus: Congestion Verification of Network Updates in Software-defined Networks", *Winter Simulation Conference (WSC 2016)*, Washington, DC, December 11-14, 2016. https://cps-vo.org/node/31574

One of the challenges encountered developing SDN strategies for resilience is that the algorithms for generating rules must simultaneously ensure correctness of the ruleset and optimize performance. As correctness is paramount, means of separating the correctness concerns from performance (and have the latter attributes attended to automatically and transparently) make design of SDN resilient networks more achievable.

- Santhosh Prabhu, Mo Dong, Tong Meng, Brighten Godfrey, and Matthew Caesar, "Let Me Rephrase That: Transparent Optimization in SDNs," *ACM Symposium on SDN Research (SOSR 2017)*, Santa Clara, CA, April 3-4, 2017. https://cps-vo.org/node/34796

The assurance of network resiliency depends on an ability to verify the properties of the network. By combining symbolic modeling of data plane and control plane with explicit state exploration, our tool Plankton performs a goal-directed search on a finite-state transition system that captures the behavior of the network as well as the various events that can influence it. In this way, Plankton can automatically find policy violations that can occur due to a sequence of network events, starting from the current state. Initial experiments have successfully predicted scenarios like BGP Wedgies.

- Santhosh Prabhu, Ali Kheradmand, Brighten Godfrey, and Matthew Caesar, "Predicting Network Futures with Plankton", *1st Asia-Pacific Workshop on Networking (APNet'17)*, Hong Kong, China, August 3-4, 2017. https://cps-vo.org/node/37477

An SDN architecture must incorporate various notions of accountability for achieving systemwide cyber resiliency goals. We analyze accountability based on a conceptual framework, and we identify how that analysis fits in with the SDN architecture's entities and processes.

- Benjamin Ujcich, Andrew Miller, Adam Bates, and William Sanders, "Towards an Accountable Software-Defined Networking Architecture", *3rd IEEE Conference on Network Softwarization (NetSoft 2017)*, Bologna, Italy, July 3-7, 2017. https://cps-vo.org/node/36600

Resilience implies response, and response implies observation and monitoring. Probability of detection increases as one uses multiple techniques for the anomaly detection. We developed a method for detecting low-activity lateral movement of attackers within a network using multiple detection schemes simultaneously.

- Atul Bohara, Mohammad Noureddine, Ahmed Fawaz, and William Sanders, "An Unsupervised Multi-Detector Approach for Identifying Malicious Lateral Movement," *36th IEEE International Symposium on Reliable Distributed Systems (SRDS 2017)*, Hong Kong, China, September 26-29, 2017. https://cps-vo.org/node/39033

We addressed the problem of anomaly detection, gathering evidence pointing to possible insider attacks based on the interaction of sequences of user interactions with resilient architectural behaviors. The research method involved a combination of dynamic network analysis techniques with behavioral models of software architecture. An evaluation simulating a large-scale web infrastructure showed that the approach could perfectly rank anomalous traces on an architecture with 240 components and 10% anomalous behaviors, and it could maintain better than 95% perfect ranking of anomalous traces with up to 40% of simulated traces as anomalies.

- Hemank Lamba, Thomas J. Glazier, Javier Camara, Bradley Schmerl, David Garlan and Jurgen Pfeffer, "Model-based Cluster Analysis for Identifying Suspicious Activity Sequences in Software," *3rd International Workshop on Security and Privacy Analytics (IWSPA)*, co-located with ACM CODASPY, Scottsdale, AZ, March 23 – 25, 2017. https://cps-vo.org/node/34450

Human Behavior

The Hard Problem of Human Behavior addresses how to handle the unpredictability and complexity of human actors in cybersecurity. These actors include malicious attackers, system users, and software/system developers.

Goal: Develop models of human behavior that enable the design, modeling, and analysis of systems with specified security properties.

Progress in 2017

We investigated the specific problem of understanding the psychosocial and cognitive characteristics of users in relation to their vulnerability to phishing given persuasive attributes of phishing messages. Its specific method was to empirically collect measures of individual user characteristics and message attributes and analyze the interaction of such variables and assess phishing vulnerability. This research informs our understanding of how users approach security-related decision making.

We investigated the specific problem of understanding user interface interventions that would help users resist phishing and other deceptive practices of attackers. Its specific method was based on the development of an information-processing model that emphasizes the actions users take to counteract phishing scams. It empirically established that procedural knowledge of how to identify the legitimacy of a webpage helps users and that such knowledge can be provided by embedding training within warnings. This research helps develop user interfaces that enable users to take security-enhancing decisions.

- Aiping Xiong, Robert W. Proctor, Ninghui Li, Weining Yang, "Is domain highlighting actually helpful in identifying phishing webpages?," *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 2017. https://cps-vo.org/node/34190

We investigated the specific problem of distinguishing between human users and bots, which might be used for automated security exploits. Its specific method involved applying established cognitive architectures to develop task-level and cognitive models of typing and of sequential pointer movements, and determining baseline behaviors for comparison in the context of security. This research advances system security by providing a detailed representation of security-relevant user behavior, by tying behaviors and performance to implicit or inferrable information, such as user goals and user decision points. We addressed the problem of understanding the way humans understand privacy and personal data security online. This research demonstrated that users vastly underestimate how much of their data is being collected and used, and that they are uncomfortable with common uses of this data. The results emphasize the importance of developing clear ways of explaining data collection and use, as well as nascent user demand for more control over their personal data.

- Jennifer Golbeck, "The Importance of Consent in User Comfort with Personalization," *International Conference on Social Informatic*s, 2017.
- Jennifer Golbeck, "I'll be Watching You: Policing the Line between Personalization and Privacy," in *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization*, ACM, 2017.
- Jennifer Golbeck. "User Privacy Concerns with Common Data Used in Recommender Systems," *International Conference on Social Informatics*, 2016.

We addressed the problem of how people learn about recommended security behaviors, and how they decide whether or not to follow these recommendations. In particular, our research examined how users' security beliefs, knowledge, and demographics correlate with their sources of security advice, as well as how these factors influence their security behaviors. The researchers used a near-census-representative survey sample and logistic regression models to identify factors that correlate with different potential sources of advice. They also used statistical hypothesis testing to understand the most important reasons why users accept and reject advice they are aware of. Researchers found that advice sources are correlated with income, education, and internet skill level. They also found that advice was rejected for inconvenience, for including too much marketing material, and because the user had not previously had a negative experience. On the other hand, users were determined to be more likely to accept advice about unfamiliar security concepts based on their trust in the source, rather than on the content of the advice. These results can be used to improve the delivery of security advice in order to increase the likelihood of uptake of important security behaviors.

We considered the problem of how to motivate users to adopt secure behaviors. In particular, the researchers examine the effectiveness of video "edutainment" to motivate users to apply software updates more quickly. This was done using the participatory design method to inform the design of an edutainment story about software updating, which was developed both as text and video. In an online experiment, users were exposed to either the developed edutainment materials, a control unrelated to software updating, or a control using traditional, non-story-based security advice; users then evaluated their behavioral intention for software updating longitudinally, using the well validated SEBIS instrument. Using statistical hypothesis testing, researchers found that edutainment produced a small but significant improvement in user intent to update software. These results suggest that edutainment-based approaches may have value as one component of a strategy for motivating important security behaviors.

By surveying Android users, researchers established how users' past experiences with software updating and users' psychometric traits differentiate those users who avoid application auto-updates from those who apply them, as well as users' preferences towards auto-updating their applications. Their findings reveal that users who avoid application auto-updates are more likely to have had past negative experiences with software updating, tend to take fewer risks, and display greater proactive security awareness. Users' perceived level of trust with mobile applications also determined how comfortable they are auto-updating these applications. These findings suggest ways that Android can improve the design of application update systems to encourage users to auto-update and keep their devices secure.

- Arunesh Mathur and Marshini Chetty, "Impact of User Characteristics on Attitudes Towards Automatic Mobile Application Updates," *SOUPS 2017.*

Researchers carried out a study of 30 users' software-updating practices, developed a low-fidelity prototype to address the issues identified in their formative work, and then evaluated the prototype with a study of 22 users. The findings suggest that updates interrupt users, that users lack sufficient information to decide whether or not to update, and that users vary in terms of how they want to be notified and provide consent for updates. Based on the study, the researchers made four recommendations to improve desktop updating interfaces, and outlined socio-technical considerations for software updating that will ultimately affect end-user security.

The results of previous research indicate that the presentation of deterring situational stimuli in an attacked computing environment shapes system trespassers' avoiding online behaviors during the progression of a system trespassing event. Nevertheless, none of these studies comprised an investigation of whether the effect of deterring cues influence system trespassers' activities on the system. Moreover, no prior research has been aimed at exploring whether the effect of deterring cues is consistent across different types of system trespassers. We examine whether the effect of situational deterring cues in an attacked computer system influenced the likelihood of system trespassers engaging in active online behaviors on an attacked system, and whether this effect varies based on different levels of administrative privileges taken by system trespassers. By using data from a randomized experiment, we find that a situational deterring cue reduced the probability of system trespassers with fewer privileges on the attacked computer system (non-administrative users) to enter activity commands. In contrast, the presence of these cues in the attacked system did not affect the probability of system trespassers with the highest level of privileges (administrative users) to enter these commands.

- Alexander Testa, David Maimon, Bertrand Sobesto, and Michel Cukier, "Illegal Roaming and File Manipulation on Target Computers," *Criminology & Public Policy 16*: 689–726, 2017.

Cloud platforms contain large amounts of user data, ranging from emails and files to location and search history. Applications can request access to the data stored in these platforms in order to provide additional services and functions to users. Currently, most cloud platforms provide all-or-nothing access to each class of data (for example, access to all email or no email) with no options in between. Researchers studied the feasibility of a "data-centric" access-control scheme in which data is organized, at the platform level, into groups that can be granularly shared or not shared with different applications. In an online study, they used machine learning to suggest access-control groups, called bubbles, for content extracted from participants' Gmail, Google Drive, and Google Calendar accounts. They then asked participants to rate the accuracy and appropriateness of the suggested bubbles. The results shed light on the potential capability of data-centric access control to support more fine-grained policy than is currently available in many cloud services, without a significant increase in policy-specification effort.

Humans circumvent security rules, and when that circumvention is recognized the question remains how to approach the humans so as to change that behavior.

- Ross Koppel, Vijay Kothari, Sean Smith, and Jim Blythe, "Beyond Pleading with or Restricting Users to Achieve Cyber Security Goals: Approaches to Understanding and Responding to Circumvention," *CRA CCC Sociotechnical Cybersecurity Workshop*, College Park, MD, December 12-13, 2016. https://cps-vo.org/node/31622

The way humans form models of how cyber systems in general, and security in particular, influences their behavior. Wrong models lead to bad decisions.

- Sean Smith, Vijay Kothari, Jim Blythe, and Ross Koppel, "Flawed Mental Models Lead to Bad Cyber Security Decisions: Let's Do a Better Job", *CRACCC Sociotechnical Cybersecurity Workshop*, College Park, MD, December 12-13, 2016. https://cps-vo.org/node/34177

Mobile applications frequently request sensitive data. While prior work has focused on analyzing sensitive-data uses originating from well-defined API calls in the system, the security and privacy implications of inputs requested via application user interfaces have been widely unexplored. Our goal is to understand the broad implications of such requests in terms of the type of sensitive data being requested by applications.

- Benjamin Andow, Akhil Acharya, Dengfeng Li, William Enck, Kapil Singh, and Tao Xie, "UiRef: Analysis of Sensitive User Inputs in Android Applications," *10th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec 2017)*, Boston, MA, July 18-20, 2017. https://cps-vo.org/node/36590

We are addressing the problem of enabling automatic static and dynamic checks for policy compliance on the basis of formalized models of underlying policies. Our approach includes four thrusts: formal languages designed to express privacy principles and enable checking of data flow for compliance in a system's architecture; information type ontology theory to enable algorithmic checks of whether meanings assigned to predicates in a formal privacy language are consistent among policy makers and developers; static and dynamic analysis techniques to find privacy violations by identifying and tracing data collection, use, and transfer practices throughout source code; and privacy risk metrics to reliably measure how users and data subjects perceive the risk of specific data practices, so that developers can focus effort on mitigating the most important privacy threats. Our goal is to facilitate communication between policy makers and developers by assuring formal consistency in terminology and data practices. This includes tracing high-level data types to low-level expressions in source code, including both platform API-managed data types and user-provided data.

- Jaspreet Bhatia, Travis Breaux, and Florian Schaub, "Mining Privacy Goals from Privacy Policies Using Hybridized Task Recomposition," *ACM Transactions on Software Engineering Methodology*, Article 22 (May 2016), 24 pages. Invited for presentation at ICSE 2017. https://cps-vo.org/node/36590
- Jaspreet Bhatia, Travis Breaux, Joel R. Reidenberg, and Thomas Norton, "A Theory of Vagueness and Privacy Risk Perception," in *Proceedings, Requirements Engineering Conference*, 2016. Nominated for best paper. https://cps-vo.org/node/31343
- Joel Reidenberg, Jaspreet Bhatia, Travis Breaux, and Thomas Norton, "Ambiguity in Privacy Policies and the Impact of Regulation," *Journal of Legal Studies*, Fordham Law Legal Studies Research Paper No. 2715164.

We are addressing the problem of how to gather useful data about security-related user behaviors when accessing websites. Our approach has been to develop a multi-purpose observational resource, the Security Behavior Observatory (SBO), to collect data from Windows home computers. The SBO collects a wide array of system, network, and browser data from a cohort of over 500 home Windows computer users. This resource has already led to some surprising results. Participants who are more engaged with computer security and maintenance did not necessarily have better security outcomes, in the sense of failures to patch, evidence of malware, etc. We discovered that these failures are not necessarily best addressed through automated security measures, which work well for some users but not for others who circumvent those measures. Using the Security Behavior Intentions Scale we developed as an instrument, we discovered that security-related intentions are not a good proxy for actual behavior. Finally, based on data from actual password use, we have discovered that partial password reuse (e.g., a common stem) is rampant, which suggests potential security risks that go beyond reuse of entire passwords.

- Sarah Pearman, Jeremy Thomas, Pardis Emami Naeini, Hana Habib, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, Serge Egelman, Alain Forget, "Let's go in for a closer look: Observing passwords in their natural habitat," *ACM CCS 2017*.

CARNEGIE MELLON
UNIVERSITY

William Scherlis

The Carnegie Mellon University (CMU) Science of Security (SoS) Lablet team is led by Principal Investigator (PI) William Scherlis, with co-PIs Travis Breaux and Jonathan Aldrich. The lead financial manager is Monika DeReno. The team is engaged in nine research projects, and has seven university subcontracts with about fifty faculty, postdoctoral, and PhD student researchers working at eight campuses.

At Carnegie Mellon, researchers are drawn from six different academic departments in three different colleges:

- In the School of Computer Science (SCS): the Institute for Software Research (ISR), the Human Computer Interaction Institute (HCII), and the Computer Science Department (CSD).
- In the Engineering College (the Carnegie Institute of Technology, (CIT)): the Electrical and Computer Engineering (ECE) department, and the Engineering and Public Policy (EPP) department.
- In the Heinz College: the School of Information Systems and Management.

Approximately sixteen PhD students participate in the program, drawn from the departments and institutions identified above. In addition, several MS and undergraduate students participate directly in Lablet projects. Many of the Lablet faculty and graduate students are affiliated with the CMU CyLab, the coordinating entity for cybersecurity research.

Subcontractors (and sub-PI's) include the following:

- Cornell (Dexter Kozen)
- University of California, Irvine (UCI) (Sam Malek)
- University of California, Berkeley (UCB) (Serge Egelman)
- University of Pittsburgh (UP) (Scott Beach)
- Wayne State University (WSU) (Marwan Abi-Antoun)
- University of Nebraska (UN) (Matt Dwyer, Witawas Srisa-an)
- University of Texas, San Antonio (UTSA) (Jianwei Niu)

The team also collaborates with researchers at Bosch, Google, and Microsoft.

The CMU Lablet addresses all five of the Science of Security Hard Problems, with emphasis on Scalability and Composability and Human Behavior. There is significant additional effort relating to the other three problems, Policy-Governed Secure Collaboration, Security Metrics and Models, and Resilient Architectures.

FUNDAMENTAL RESEARCH

Project: Frameworks, APIs, and Composable Security Models

PI: Jonathan Aldrich

Participating Sub-Lablets: University of California, Irvine; Wayne State University

Hard Problems: Scalability and Composability, Security Metrics and Models, Human Behavior

Developers of plugins for frameworks, such as apps in an app store, require ways to analyze their code to ensure that security properties hold in the context of a highly complex and rich software framework. One of the challenges in framework use by developers is the complexity of the rules of engagement at the framework APIs, particularly regarding protocols. An empirical study of software developers by this team shows specifically how developers have trouble understanding the complex protocols associated with framework APIs. This leads to problems with productivity, program correctness, and security. An intervention was tested that shows the value to developer quality—and productivity—of adding protocol guidance to developer tooling, with developers working twice as fast and eight times less likely to make a mistake in our intervention.

Project: Limiting Recertification in Highly Configurable Systems: Interactions and Isolation among Configuration Options
PI: Christian Kästner
Hard Problems: Scalability and Composability, Security Metrics and Models

Both evolution and configuration of a system increase the security challenge significantly. While some versions and configurations of a system might be secure, changing a single configuration option, updating a single dependency, or combining multiple configuration options in unanticipated ways can expose new vulnerabilities. Heartbleed, for example, affected only users who activated the (default, but not required) heartbeat configuration option. In the Linux kernel there are 15K compile-time configuration settings, with the total number of configurations (a theoretical maximum of 2^15k) being intractably huge. The team investigated techniques to simultaneously analyze very large numbers of compile-time configurations of C code at once, without looking at each configuration individually, a kind of configuration-aware taint tracking. They analyzed which parts of the code might suffer from increased configuration complexity (e.g., many #ifdefs), and applied the ideas directly to examples such as the Linux kernel and OpenSSL.

It turns out that network analysis techniques can be a predictor of which code fragments are more likely to contain security vulnerabilities in the Linux kernel. At the same time, we have investigated evolution issues by analyzing unusual commits that are outliers and merit security review. In an empirical evaluation, we found that our anomaly detection mechanism detected code commits that are actionable for developers to prioritize reviews. (The figure below illustrates the use of graph-theoretic network metrics to identify areas of variational code needing more scrutiny.)

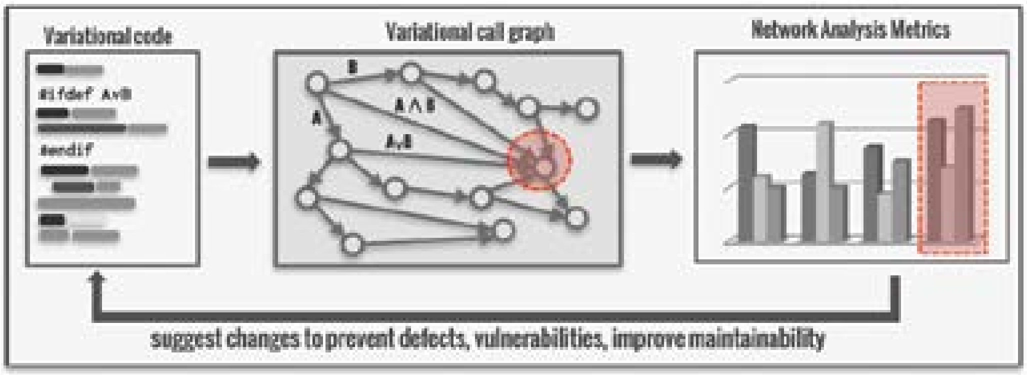

In a related study, the team analyzed how security practices are enforced in standardized software certification processes, such as Common Criteria. Interviews with Subject Matter Experts led to identification of tensions in the certification process, emphasizing that rapid recertification and composition issues are among the main concerns in today's certification. The effect of these issues is that old, vulnerable, but certified versions are often preferred to newer patched ones that are still going through the lengthy certification process. The studies confirm that a larger focus on rapid and compositional certification is critical to address security concerns for software systems, including those that are highly configurable.

Project: Security Reasoning for Distributed Systems with Uncertainties
PI: Andre Platzer
Participating Sub-Lablet: Cornell University
Hard Problems: Scalability and Composability, Resilient Architectures

Our research project investigates using artificial intelligence and optimization techniques for security problems related to systems with physical behavior. Many real-world devices (such as industrial plants and autonomous vehicles) use complicated control policies to guide the underlying dynamics of the system to desirable states. These control policies are, in general, difficult to analyze and prove free of security defects. Our research project is focused on improving capabilities of automatic synthesis and verification of the control of such systems, especially in the face of both benign uncertainty (e.g., sensor noise) and adversarial action. In particular, we are solving an open problem in this field: designing controllers using approximate techniques while retaining guarantees about the resulting behavior quality. As an alternative to exact approaches, we use approximate techniques that can be faster but need guarantees on their errors. We have made substantial progress on approximately solving a class of optimization problem related to both anomaly detection and security policy synthesis. Anomaly detection and security policy synthesis are essential capabilities for security in complicated systems where uncertainty is present and must be handled in a principled fashion. We have developed a robust implementation of an approximate solver for such problems. This includes solving a wealth of issues associated with numerical stability and basis generation. Our work includes developing a novel variant of Monte Carlo Tree Search (MCTS). This MCTS variant allows us to use our approximate solver to tackle continuous real-world policy synthesis problems that are difficult to solve using other methods because of the exponential runtime scaling of conventional methods. This project has focused on approximate optimization techniques with specific application to security planning and policy setting. These techniques involve solving an optimization problem within a restricted basis of functions. This technique, called Galerkin approximation, allows problems to be solved rapidly but with some loss in solution quality. We are focused on applying these methods to policy synthesis questions, but they can also be applied to anomaly detection problems. Many anomaly detection and policy synthesis questions can be cast as instances of a general problem called the Linear Complementarity Problem (LCP) (Zawadzki, Gordon, & Platzer, "A Projection Algorithm for Strictly Monotone Linear Complementarity Problems" 2013).

We are working on LCP instances based on Markov decision processes (MDPs). MDPs are a popular framework for making sequential decisions in uncertain environments. Synthesizing good security polices is an example of this kind of task. For example, we may want to automatically generate a good policy for when access control restrictions should be escalated in response to anomalous system behavior, while still allowing legitimate users to work productively. Our MDP approximation techniques will improve policy synthesis capability for many of these large security problems.

Our work ties into our earlier work on #E-SAT solving (Zawadzki, Platzer, & Gordon, "A Generalization of SAT and #SAT for Robust Policy Evaluation" 2013) which offered new capabilities to analyze the robustness of security policies under uncertainty. Together, these methods represent two important links in a chain of tools for generating and verifying policies in a number of different security applications that exhibit uncertainty. For example, one can first synthesize a policy using our MDP techniques, and then can analyze the policy's robustness to random failure using our #E-SAT methods.

Project: Compositional Security

PI: Anupam Datta
Hard Problems: Scalability and Composability, Policy-Governed Secure Collaboration

The goal of this project is to allow both sequential and parallel composition of arbitrary number of verified and unverified system program modules—and in a potentially adversarial environment (e.g., with untrusted system modules). The team developed a semantic model that supports compositionality of components in the presence of such adversaries. The UberSpark system enforces verifiably-secure object abstractions in systems software written in C99 and assembly. This is significant because the absence of explicit abstractions in low-level programming is a major source of vulnerabilities (and the subject of much binary analysis). UberSpark provides a verifiable overlay of abstractions, and can compositionally verify security properties of extensible commodity systems software (e.g., BIOS, OS, hypervisor) written in low-level languages. The UberSpark abstractions include interfaces, function-call semantics for implementations of interfaces, access control on interfaces, and forms of concurrency. It uses a combination of hardware mechanisms and light-weight static analysis. UberSpark provides a verifiable assembly language CASM which is compiled using the certifying CompCert compiler into executable binaries. The team has validated UberSpark on a performant hypervisor, demonstrating only minor performance overhead, and low verification costs. UberSpark is available at http://uberspark.org

Project: A Language and Framework for Development of Secure Mobile Applications

PI: Jonathan Aldrich
Hard Problems: Scalability and Composability, Human Behavior

Project lead Jonathan Aldrich, with co-lead Joshua Sunshine, developed language and framework concepts that support "building security in" for mobile and web applications, with specific benefits relating to both command injection, such as SQL injection attacks, and also least privilege, in how applications use system resources. The key idea is to embed domain-specific command languages, such as SQL, within the programming language—and to do this in a way such that it is more convenient to use the embedded (and secure) DSL than to use insecure string operations, which is typical in current web applications. A key feature of the approach is to do this in a general manner, enabling extension of a host language with new command DSLs, and to do this in a modular way. Past approaches were not adopted because they were un-modular, preventing separately-defined embedded DSLs to be used together.The team developed a novel composition mechanism, type-specific languages, that supports modular DSL embeddings by associating a unique DSL with appropriate types.

Project: Multi-Model Run Time Security Analysis

PI: David Garlan
Participating Sub-Lablets: University of California, Irvine; Wayne State University
Hard Problems: Scalability and Composability, Security Metrics and Models, Resilient Architectures

Project PI David Garlan, with co-leads Jonathan Aldrich, Bradley Schmerl, Javier Cámara, Sam Malek, and Marwan Abi-Antoun, combined techniques from dynamic network analysis and software architecture to improve security to create a model-based anomaly detection algorithm. The algorithm identifies evidence for insider attacks based on the interaction of sequences of user interactions with resilient architectural behaviors. The algorithm proved better at detecting suspicious activities than other approaches that do not consider architectural behavior, such as unigrams or bigrams, but which are traditionally used for anomaly detection.

One of the challenges faced by the team was measurement, i.e., how to evaluate the anomaly detection approach.

Experimental evaluation was hindered by a combination of a lack of real insider threat data and real platforms to run experiments on. This led to the implementation of an open source testbed platform whose design is based on knowledge of real attacks on enterprise software systems. With the simulated large-scale web infrastructure, the team showed 1) the approach could perfectly rank anomalous traces on an architecture with 240 components and 10% anomalous behaviors; 2) it could maintain better than 95% perfect ranking of anomalous traces with up to 40% of simulated traces as anomalies; and 3) while large architectures give better results, even with small architectures (those with 100 nodes), the team was able to rank 60% of the anomalies.This work is on-going with a goal of enabling simulation of real advanced persistent threats on the testbed, leading to both improved self-protecting resilience technologies to mitigate threats and also improved measurement techniques, based on testbed use.

Project: Race Vulnerability Study and Hybrid Race Detection
PI: Jonathan Aldrich
Participating Sub-Lablet: University of Nebraska
Hard Problems: Scalability and Composability, Human Behavior

Jonathan Aldrich and his team developed a compositional approach to "deep immutability" that addresses the frequent mismatch of programmer expectations when reference variables are declared immutable: the reference itself might be guaranteed not to change, but the object referenced may nonetheless change. This "expectation gap" is the source of a range of escalation-of-privilege vulnerabilities, (such as the Java 1.1 getSigners() method that shares an "immutable" reference to an array of signers that turns out to be mutable). Transitive immutability is a composition issue, since expectations must be expressed and analyzed across component boundaries.

Evaluation of this approach included not just mathematical analysis with respect to composition, but also empirical evaluations involving developers and code. As noted above, the underlying issue is the semantic expectations of developers. The team therefore assessed "developer usability" issues with human-subject studies demonstrating success rates that improved from a plain-Java control baseline of 0-40% (depending on task) to 70-90% success rates with our intervention. The team also examined existing code bases to assess the extent of change required to model and analyze deep immutability, with results indicating very few minor changes were needed. The results were reported in ICSE 2016 and ICSE 2017.

Project: Usable Formal Methods for the Design and Composition of Security and Privacy Policies
PI: Travis Breaux
Participating Sub-Lablet: University of Texas, San Antonio
Hard Problems: Scalability and Composability, Policy-Governed Secure Collaboration, Human Behavior

This team undertook work on a framework to support holistic analysis of systems for privacy. There are four interrelated scientific thrusts: 1) formal languages, based on natural language privacy policies, designed to express privacy principles and check data flow for compliance in a system's architecture; 2) information type ontology theory to enable algorithmic checks of whether meanings assigned to predicates in a formal privacy language are consistent among policy makers and developers; 3) static and dynamic analysis techniques to find privacy violations by identifying and tracing data collection, use, and transfer practices throughout source code; and 4) privacy risk metrics to reliably measure how users and data subjects perceive the risk of specific data practices, so that developers can focus effort on mitigating the most important privacy threats.These results are significant because it is otherwise too easy for policy makers and developers to "speak past" each other by hiding inconsistent data practices through ambiguous terminology. This includes tracing high-level data types to low-level expressions in source code, including both platform API-managed data types and user-provided data.Privacy differs from security, because privacy is characteristically about the data subject, or the person about whom the data is descriptive. One of the goals of this work is to trigger a new line of design thinking that is not limited to access control and anonymization, and which can introduce consent and use-based restrictions imposed by architecture.

Project: Understanding User Behavior when Security is a Secondary Task
PI: Lorrie Cranor
Participating Sub-Lablets: University of California, Berkeley; University of Pittsburgh
Hard Problems: Security Metrics and Models, Human Behavior

In the past three years, the team has developed a multi-purpose observational resource, the Security Behavior Observatory (SBO), to collect data from Windows home computers. The SBO collects a wide array of system, network, and browser data from over 500 home Windows computer users (who participate as human subjects).
This resource has already led to some surprising discoveries: 1) Participants who are more engaged with computer security and maintenance did not necessarily have better security outcomes, in the sense of failures to patch, evidence of malware, etc. 2) These deficiencies are not necessarily best addressed through automated security measures, which work well for some users but fail others who circumvent those measures. 3) The subcontractor partner developed and validated a Security Behavior Intentions Scale to quickly measure user security behavior. One conclusion is that end-user survey results show that intentions are not a good proxy for actual behavior. 4) More recently, the team has been collecting high quality password data (i.e., not from self-reports or lab studies, but from actual password use), and has shown that partial password reuse (e.g., a common stem) is rampant, which suggests potential security risks that go beyond reuse of entire passwords

PUBLICATIONS

Project: Frameworks, APIs, and Composable Security Models

- Esther Wang and Jonathan Aldrich, "Capability Safe Reflection for the Wyvern Language," in *Proceedings of the Workshop on Meta-Programming Techniques and Reflection (META)*, part of Systems, Programming, Languages and Applications: Software for Humanity (SPLASH), Amsterdam, Netherlands, October 30 – November 4, 2016. Presented by Esther Wang.
- Mahmoud Hammad, Hamid Bagheri, and Sam Malek, "DELDroid: Determination and Enforcement of Least-Privilege Architecture in Android," *IEEE International Conference on Software Architecture (ICSA 2017)*, Gothenburg, Sweden, April 3 – 7, 2017. Presented by Hamid Bagheri.
- Alireza Sadeghi, Hamid Bagheri, Joshua Garcia, and Sam Malek, "A Taxonomy and Qualitative Comparison of Program Analysis Techniques for Security Assessment of Android Software," Published online October 6, 2016: http://ieeexplore.ieee.org/document/7583740/ *IEEE Transactions on Software Engineering, Vol: 43(6), June 1, 2017*. DOI: 10.1109/TSE.2016.2615307
- Alireza Sadeghi, Naeem Esfahani, and Sam Malek, "Ensuring the Consistency of Adaptation through Inter- and Intra-Component Dependency Analysis," *Journal ACM Transactions on Software Engineering and Methodology (TOSEM). Vol 26(1), Article 2, July 12, 2017.* https://doi.org/10.1145/3063385

Project: Limiting Recertification in Highly Configurable Systems: Interactions and Isolation among Configuration Options

- Jafar Al-Kofahi, Tien Nguyen, and Christian Kästner, "Escaping AutoHell: A Vision for Automated Analysis and Migration of Autotools Build Systems," in *Proceedings of the 4th International Workshop on Release Engineering (RELENG),* New York, NY: ACM Press. Seattle, WA, November 18, 2016. Presented by Jafar Al-Kofahi.
- Meng Meng, Jens Meinicke, Chu-Pan Wong, Eric Walkingshaw, and Christian Kästner, "A Choice of Variational Stacks: Exploring Variational Data Structures," in *Proceedings of the 11st International Workshop on Variability Modelling of Software-Intensive Systems (VaMoS)*, Eindhoven, The Netherlands, February 1 – 3, 2017. Presented by Eric Walkingshaw.
- Flávio Medeiros, Márcio Ribeiro, Rohit Gheyi, Sven Apel, Christian Kästner, Bruno Ferreira, Luiz Carvalho, and Baldoino Fonseca, "Discipline Matters: Refactoring of Preprocessor Directives in the #ifdef Hell," *IEEE Transactions on Software Engineering (TSE), March 28, 2017*. DOI: 10.1109/TSE.2017.2688333. Presented by Flávio Medeiros.
- Gabriel Ferreira, "Software Certification in Practice: How Are Standards Being Applied?" in *Proceedings of the 2017 IEEE/ACM 39th International Conference on Software Engineering Companion (ICSE)*, Buenos Aires, Argentina, May 20 – 28, 2017. Presented by Gabriel Ferreira.
- Rohit Goyal, Gabriel Ferreira, Christian Kästner, and James Herbsleb, "Identifying Unusual Commits on GitHub," *Journal of Software: Evolution and Process (JSEP)*, 2017.

Project: A Language and Framework for Development of Secure Mobile Applications

- Jonathan Aldrich and Alex Potanin "Naturally Embedded DSLs," *Workshop on Domain-Specific Language Design and Implementation (DSLDI),* part of Systems, Programming, Languages and Applications: Software for Humanity (SPLASH), Amsterdam, Netherlands, October 31-Nov 1, 2016. Presented by Jonathan Aldrich.
- Ian Voysey, Cyrus Omar, and Matthew Hammer, "Running Incomplete Programs," part of Principles of Programming Language (POPL), *Off the Beaten Track Workshop (OBT),* Paris, France, January 15 – 21, 2017. Presented by Ian Voysey.
- Cyrus Omar, Ian Voysey, Michael Hilton, Joshua Sunshine, Claire Le Goues, Jonathan Aldrich, and Matthew Hammer, "Toward Semantic Foundations for Program Editors," *2nd Summit on Advances in Programming Languages (SNAPL)*, Asilomar, CA, May 7 – 10, 2017. DOI: 10.4230/LIPIcs.SNAPL.2017.23. Presented by Cyrus Omar.
- Darya Melicher, Yangqingwei Shi, Alex Potanin, and Jonathan Aldrich, "A Capability-Based Module System for Authority Control," in *Proceedings of the 31st International European Conference on Object-Oriented Programming (ECOOP)*, Barcelona, Spain, June 18 – 23, 2017. Presented by Darya Melicher.

Project: Multi-Model Run Time Security Analysis

- Hemank Lamba, Thomas Glazier, Javier Cámara, Bradley Schmerl, David Garlan, and Jürgen Pfeffer, "Model-based Cluster Analysis for Identifying Suspicious Activity Sequences in Software," *3rd International Workshop on Security and Privacy Analytics (IWSPA)*, co-located with ACM CODASPY. Scottsdale, AZ, March 23 – 25, 2017. Presented by Hemank Lamba.

Project: Race Vulnerability Study and Hybrid Race Detection

- Tingting Yu, Witty Srisa-an, and Gregg Rothermel, "SimExplorer: An Automated Framework to Support Testing for System-Level Race Conditions," *Journal of Software: Testing, Verification and Reliability. Vol 27, 4-5. May 10, 2017*. DOI: 10.1002/stvr.1634
- Junjie Qian, Hong Jiang, Witawas Srisa-an, Sharad Seth, Stan Skelton, and Joseph Moore, "Energy-efficient I/O Thread Schedulers for NVMe SSDs on NUMA," *17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, Madrid, Spain, May 14 – 17, 2017. Presented by Junjie Qian.

- Yutaka Tsutano, Shakthi Bachala, Witawas Srisa-an, Gregg Rothermel, and Jackson Dinh, “An Efficient, Robust, and Scalable Approach for Analyzing Interacting Android Apps,” *39th International Conference on Software Engineering (ICSE),* Buenos Aires, Argentina, May 20 – 28, 2017. Presented by Yutaka Tsutano.
- Michael Coblenz, Whitney Nelsony, Jonathan Aldrich, Brad Myers, and Joshua Sunshine, “Glacier: Transitive Class Immutability for Java,” *39th International Conference on Software Engineering (ICSE),* Buenos Aires, Argentina, May 20 – 28, 2017. Presented by Michael Coblenz.

Project: Usable Formal Methods for the Design and Composition of Security and Privacy Policies

- Jaspreet Bhatia, Travis Breaux, and Florian Schaub, “Mining Privacy Goals from Privacy Policies Using Hybridized Task Recomposition,” *ACM Transactions on Software Engineering Methodology*, Article 22 (May 2016), 24 pages. Invited for presentation at ICSE 2017.
- Jaspreet Bhatia, Travis D. Breaux, Joel R. Reidenberg, and Thomas B. Norton. “A Theory of Vagueness and Privacy Risk Perception”. *Proceedings of the Requirements Engineering Conference*, 2016. (Nominated for best paper).
- Joel R. Reidenberg, Jaspreet Bhatia, Travis D. Breaux, and Thomas B. Norton, “Ambiguity in Privacy Policies and the Impact of Regulation ,” *Journal of Legal Studies,* Fordham Law Legal Studies Research Paper No.2715164

Project: Understanding User Behavior when Security is a Secondary Task

- Joshua Tan, Lujo Bauer, Joseph Bonneau, Lorrie Cranor, Jeremy Thomas, and Blase Ur, “Can Unicorns Help Users Compare Crypto Key Fingerprints?” in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems,* Denver, CO, May 6 – 11, 2017. Presented by Joshua Tan.
- Sarah Pearman, Jeremy Thomas, Pardis Emami Naeini, Hana Habib, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, Serge Egelman, Alain Forget. Let’s go in for a closer look: Observing passwords in their natural habitat. *ACM CCS*, 2017.

EDUCATIONAL

Project: Multi-Model Run Time Security Analysis

We have engaged a team from the Master of Information Technology Strategy program for development of the exemplar testbed as part of their practicum project. We have worked with a Masters of Software Engineering student in an independent study for formal modeling of advanced persistent threats. We have engaged an undergraduate from Duke University to implement an instance of a simulated advanced persistent threat scenario on the exemplar testbed.

COMMUNITY OUTREACH

Project: Compositional Security

We have open-sourced and released the first academic prototype of UberSpark. The development and evolution of the framework continues at http://uberspark.org Our hope, and ultimate vision, for UberSpark is to foster the growth of the next generation of systems developers who can design and develop low-level verifiable and trustworthy systems with a focus on verifiability, development compatibility and performance.

Project: Multi-Model Run Time Security Analysis

In April 2017, David Garlan and Bradley Schmerl, with colleague Mary Shaw, received the CMU School of Computer Science Allen Newell Award for Research Excellence for their work on software architecture and adaptive systems. This is top annual award for research in the School of Computer Science.

Project: Usable Formal Methods for the Design and Composition of Security and Privacy Policies

The journal paper “Privacy Goal Mining through Hybridized Task Re-composition,” was invited for presentation at the 39th International Conference on Software Engineering (ICSE), 2017.

CMU and UTSA published a website for mobile app developers and regulators to help detect and repair potential privacy violations in mobile app source code. The website is located at http://polidroid.org/tutorials/developer

Project: Understanding User Behavior when Security is a Secondary Task

Work done by Cranor and colleagues was cited as providing the scientific rationale for several of the major changes in the NIST password guidelines.

North Carolina *State University*

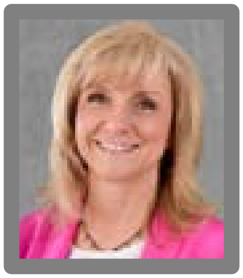

Laurie Williams

The North Carolina State University (NCSU) Science of Security (SoS) Lablet, led by Principal Investigator (PI) Laurie Williams, has embraced and helped build a foundation for the National Security Agency (NSA) 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. Continual evaluation of our research and community development efforts are key to our approach. NCSU Sub-Lablets include Purdue University, Rochester Institute of Technology (RIT), University of Alabama, University of North Carolina (UNC), University of North Carolina at Charlotte (UNCC), and University of Virginia (UVA). NCSU and its Sub-Lablets are working on sixteen research projects, two of which were added in 2017 ("High-Assurance Active Cyber Defense Policies for Auto-Resiliency" and "Formal Analysis of Breach Reports"). Project by project details of the research are listed under the Fundamental Research section of the report.

TEAMS

We have formed teams to conduct scientific research and evaluate progress on four of the five Hard Problems: Policy-Governed Secure Collaboration, Security Metrics and Models; Resilient Architectures, and Human Behavior. While we have no explicit team addressing Scalability and Composability, we address it as a secondary Hard Problem in several of our projects. Each Hard Problem team is composed of three or more projects researching complementary aspects of the Hard Problem. We also have additional teams for Research Methods, Community Development and Support, and for Evaluation.

Research Methods, Community Development and Support:

University of Alabama: Jeff Carver
North Carolina State University: Lindsey McGowen,
Jon Stallings, Laurie Williams, David Wright

Evaluation:

North Carolina State University: Lindsey McGowen,
Jon Stallings,David Wright
University of Alabama: Jeff Carver

FUNDAMENTAL RESEARCH

Project: Attack Surface and Defense-in-Depth Metrics
PIs: Andy Meneely, Laurie Williams
Participating Sub-Lablet: Rochester Institute of Technology
Hard Problems: Scalability and Composability, Security Metrics and Models, Resilient Architectures

We developed metrics and techniques to assess security risk by examining the attack surface using stack traces and call graphs. We conducted a systematic literature review on the notion of attack surface, and created a unified definition of attack surface. We developed metrics to estimate the attack surface and vulnerability discovery to better predict vulnerabilities. This research produces better measurements of security risk by providing actionable intelligence on problematic sections of source code. Our models outperform competing models in vulnerability prediction accuracy.

Project: Systematization of Knowledge from Intrusion Detection Models
PIs: Haiyu Dai, Andy Meneely
Participating Sub-Lablet: Rochester Institute of Technology
Hard Problems: Scalability and Composability, Security Metrics and Models, Resilient Architectures, Human Behavior

We examined how intrusion detection research today can be generalized by examining the validation criteria used in empirical studies of Intrusion Detection Systems (IDS). We conducted a systematic literature review on the metrics used in intrusion detection research. We found that IDS literature has some standardization of metrics, but typically lacks full examination of trade-offs such as memory vs speed. This study provides a foundation and guidance for empirical studies to be conducted on intrusion detection research. The metrics we uncovered and systematically aggregated will help provide a consistent picture of the quality of any IDS technique. We developed a distributed incentive mechanism for IDS collaboration that approximates socially optimal VCG auction-based scheme with significantly reduced complexity. To study collaborative security with privacy concerns, we developed a repeated two-layer single-leader multi-follower game and obtained the optimal collaboration and defending strategies based on privacy policies of collaborative entities. Our results help quantify the benefits of IDS collaboration, and the trade-off between collaboration utility and privacy.

Project: Vulnerability and Resilience Prediction Models
PIs: Mladen Vouk, Laurie Williams
Hard Problems: Scalability and Composability, Security Metrics and Models, Resilient Architectures

We conducted a survey of Software as a Service (SaaS) and Platform as a Service (PaaS) vulnerabilities and countermeasures. We showed that well-understood data-flow, security awareness and provenance information models can cost-effectively represent and manage resilience of cloud-based application networks, including prioritization of security-related validation and verification efforts based on vulnerability prediction. Depending on the security priorities (e.g., Confidentiality vs. Integrity vs. Availability), we construct a data-flow model of interacting components, “learn” its normal operational profile using a provenance-collecting engine, and identify vulnerable externally facing interfaces. This model then serves to position and implement appropriate countermeasures (e.g., access controls, reputation firewalls, and data integrity checks). The modeling and assessment platform (with GUIs) is based on an augmented version of Kepler, a scientific workflow modeling tool. This research supports recognition of attacks through metrics and models, and provides models to predict resilience to attacks in the field, thereby cutting across the Hard Problems of Scalability and Composability, Security Metrics and Models, and Resilient Architectures. Additionally, we produced vulnerability prediction models for Infrastructure as a Service (IaS) scripts which are often used in continuous deployment/DevOps environments.

Project: Warning of Phishing Attacks: Supporting Human Information Processing, Identifying Phishing Deception Indicators, and Reducing Vulnerability
PIs: Christopher Mayhorn, Emerson Murphy-Hill
Hard Problem: Human Behavior

We explored the content of phishing emails and conducted a temporal analysis to determine what attack vectors were used over time and how these tactics have evolved. Phishing email content was categorized according to the principles of persuasion identified by Cialdini (1987). Recent work investigated how personality attributes interacted with phishing email content to make individuals deferentially susceptible to phishing email attacks that used specific persuasion techniques. Over time, phishing email messages have increasingly come to rely on the following persuasion principles: commitment, liking, authority, and scarcity. Personality characteristics such as high extroversion were found to make people particularly susceptible to phishing attacks. By exploring the interaction between phishing message content and personality characteristics, this work contributes to solving the Hard Problem of Human Behavior by cataloging when and why security vulnerabilities exist in the phishing domain.

Project: A Human Information-Processing Analysis of Online Deception Detection
PIs: Robert Proctor, Ninghui Li
Participating Sub-Lablet: Purdue University
Hard Problem: Human Behavior

We conducted field experiments and online and laboratory studies investigating influences of warnings and training on responses to phishing attacks. We showed that training-embedded warnings are effective at improving users' ability to identify phishing webpages. Users who received both warnings and training did not fall for simulated phishing attacks in the natural context, whereas those who received one or the other did. Domain highlighting attracted users' attention, but they did not know how to apply that information. Thus, we established that users must not only be aware of online deception but also skilled enough to make informed decisions given the risk.

Project: Leveraging the Effects of Cognitive Function on Input Device Analytics to Improve Security
PIs: David Roberts, Robert St. Amant
Hard Problem: Human Behavior

We planned, designed, piloted, and evaluated studies on the efficacy of Human Subtlety Proofs (HSPs) for detecting the effects of cognition on input device usage. Additionally, we completed the model-based labeling of event data from previous studies, providing a detailed corpus of data with labels describing cognitive processes at the millisecond resolution. Most notably, we completed the validation of our study design through numerous iterations and refinement steps. Our results lay an important foundation for attacking the Hard Problem of Human Behavior by contributing the knowledge that scaffolds cognitive-model-based interpretation of input device usage. This novel interpretive ability could be leveraged in a wide-variety of ways to address the human factors that lead to security issues.

Project: Understanding Effects of Norms and Policies on the Robustness, Liveness, and Resilience of Systems
PIs: Emily Berglund, Jon Doyle, Munindar Singh
Hard Problems: Policy-Governed Secure Collaboration, Resilient Architectures

We developed a framework for modeling the adoption of security practices in software development and explored sanctioning mechanisms that may promote greater adoption of these practices among developers. The multiagent simulation framework incorporates developer and manager roles, where developers maximize task completion and compliance with security policies, and the manager enforces sanctions based on functionality and security of the project. The adoption of security practices emerges through the interaction of manager and developer agents in time-critical projects. Results indicate that group sanctioning for security practices yields better adoption of security practices, while individual sanctioning results in lower retention of developers. The model provides comparison of security adoption in developers with different preferences and provides guidance for managers to identify appropriate sanctioning mechanisms for increasing the adoption of security tools in software development.

Project: Formal Specification and Analysis of Security-Critical Norms and Policies
PIs: Jon Doyle, Munindar Singh, Rada Chirkova
Hard Problems: Scalability and Composability, Policy-Governed Secure Collaboration

Using our representation for Sociotechnical Systems (STS) in terms of norms (to characterize the social architecture) and mechanisms (to characterize the technical architecture), we developed an approach to realize interactions between autonomous parties relating to their commitments directed to each other. This approach places commitments and other norms over an information model that captures key and causal dependencies, enabling flexible realization in a decentralized setting. We developed an approach that helps maintain alignment of commitments (as an exemplary social construct) between autonomous parties communicating over an asynchronous medium. Our research contributes a formal grounding to STS for cybersecurity and privacy that leads to new methods for architecting security and privacy into a software system.

Project: Scientific Understanding of Policy Complexity
PIs: Ninghui Li, Robert Proctor, Emerson Murphy-Hill
Participating Sub-Lablet: Purdue University
Hard Problems: Policy-Governed Secure Collaboration, Human Behavior

We analyzed potential policy misconfigurations in SEAndroid policies, in part by comparing against underlying Unix/Linux policies. We analyzed how policy misconfigurations evolve over time. We found that several kinds of policy misconfiguration problems exist. One such example is composition privilege escalation, in which combined effects of multiple rules grant more access than the sum of its parts. Whereas such problems exist across multiple versions of Android, the situation is improving. Our research contributes to methods for comparing policies specified in different layers of a large, complicated, and widely used system.

Project: Privacy Incidents Database
PI: Jessica Staddon
Hard Problem: Policy-Governed Secure Collaboration

We launched the Privacy Incidents Database web site at https://sites.google.com/site/privacyincidentsdatabase We have completed the development of a privacy incidents classifier that achieves an overall accuracy (F1 measure) of more than 93% which is 12% better than keyword-based classifiers (i.e., identifying incidents by classifying articles containing "privacy" and other privacy related keywords, as privacy incident articles). We prototyped a classifier for tweets that are about privacy incidents. The privacy incidents classifier significantly reduces the amount of human review needed to identify news articles that are about privacy incidents (and so, are content for the database). We developed visualizations of the chronology, entities, and location of privacy incidents to help users understand database entries.

Project: Resilience Requirements, Design, and Testing
PIs: Kevin Sullivan, Mladen Vouk, Ehab Al-Shaer
Participating Sub-Lablets: University of North Carolina at Charlotte, University of Virginia
Hard Problems: Security Metrics and Models, Resilient Architectures

We developed a formal framework for automated Active Cyber Defense (ACD) systems that consist of a reactive control policy controller for defense action synthesis to create investigation and reconfiguration courses-of-actions, and a verifier to ensure the consistency, correctness and safety of automated ACD. Combining cyber automation and policy verification can enable adaptive cyber defense that mitigates risk in a timely manner, with minimum human intervention. Our work creates the foundation for verifiable ACD which is a key mechanism for cyber resiliency. We also analyzed resilience issues and countermeasures in cloud SaaS and PaaS layers, focusing on multi-tenancy, isolation, and virtualization. We considered how they relate to PaaS services and security trends, as well as concerns such as user and resource isolation, side-channel vulnerabilities in multi-tenant environments, and protection of sensitive data. We recommend a number of effective procedures which can help in providing security.

Project: Redundancy for Network Intrusion Prevention Systems (NIPS)
PI: Mike Reiter
Participating Sub-Lablet: University of North Carolina
Hard Problem: Resilient Architectures

We developed a framework to support specification and composition of network monitoring, policy-enforcement, and resource-management applications in Software Defined Networking (SDN) environments. We demonstrated information leakage that results from the reactive deployment of rules to SDN switches, and how this leakage might enable an attacker to conduct reconnaissance on the flows occurring in a network. Our SDN application framework simplifies the specification and composition of applications while achieving resource efficiency, responsiveness, and fairness. Our work on SDN leakage showed how an attacker can construct an approximately optimal probe to maximize its information gain. Our research contributes a new management and policy-enforcement capability that is accessible for network operators without advanced training in algorithm design or optimization.

Project: Smart Isolation in Large-Scale Production Computing Infrastructures
PIs: Xiaohui (Helen) Gu, William Enck
Hard Problem: Resilient Architectures

We studied how smart isolation can be practically incorporated into systems, discovering novel mechanisms (e.g., lazy polyinstantiation) and properties (e.g., vulnerability inheritance) that enable better design of resilient architectures. Our research created a taxonomy about existing security isolation techniques, identified security vulnerability in popular container-based systems, using information flow control as an adaptive policy for smart isolation. Our empirical study of the Docker Hub ecosystem identified pervasive concerns of unpatched software containing security vulnerabilities. This study further highlighted the potential for a security metric based on vulnerability inheritance within similar environments. Our investigation of policy specification produced the approach of Policy By Example (PyBE), which is based on the approach of Programming by Example (PBE) used for program synthesis.

Project: Automated Synthesis of Resilient Architectures
PI: Ehab Al-Shaer
Participating Sub-Lablet: University of North Carolina at Charlotte
Hard Problem: Resilient Architectures

We developed path and IP mutation for resilience to disguise or hide the network resource from external entities, so that in the reconnaissance phase the attacker gets incorrect information about the network. Cyber resilience depends on an ability to resist reconnaissance and Distributed Denial of Service (DDoS) attacks. We showed cyber mutation to be an effective technique for cyber resilience. It can provide 86% effectiveness (hiding critical resources) with affordable overhead.

Project: High-Assurance Active Cyber Defense Policies for Auto-Resiliency
PI: Ehab Al-Shaer
Participating Sub-Lablet: University of North Carolina at Charlotte
Hard Problem: Resilient Architectures

Active cyber defense through reactive policies is a key component in cyber resiliency. We developed metrics for measuring the safety of reactive policies for active cyber defense. Developing a scalable approach for verification of reactive policy is feasible if both static and dynamic analyses are employed.

Project: Formal Analysis of Breach Reports
PIs: Laurie Williams, Munindar Singh
Hard Problems: Policy-Governed Secure Collaboration, Security Metrics and Models

We developed a semantic reasoning framework that computes the coverage of security breaches by policies via comparison of individual policy clauses and breach descriptions, i.e., identifying the gaps in between. Our investigation of a subset of the breaches reported by the US Department of Health and Human Services (HHS) revealed the gaps between Health Insurance Portability and Accountability Act (HIPAA) and reported breaches, leading to a coverage of 65%. Additionally, our classification of the 1,577 HHS breaches shows that 44% of the breaches are accidental, not malicious, misuses. We found that HIPAA's gaps regarding accidental misuses are significantly larger than its gaps regarding malicious misuses. This result reinforces the fact that designing policies to address human factors is important and nontrivial. Preliminary findings from a study on Amazon Mechanical Turk indicate that crowdsourcing is an effective way of understanding normative relations between natural language artifacts such as security policies, regulations, and breach reports.

PUBLICATIONS

Project: Attack Surface and Defense-in-Depth Metrics

- Christopher Theisen, Brendan Murphy, Kim Herzig, and Laurie Williams, "Risk-Based Attack Surface Approximation: How Much Data is Enough?" *International Conference on Software Engineering (ICSE) Software Engineering in Practice (SEIP)* 2017, Buenos Aires, Argentina.

Project: Systematization of Knowledge from Intrusion Detection Models

- Richeng Jin, Xiaofan He, Huaiyu Dai, Rudra Dutta, and Peng Ning, "Towards Privacy-Aware Collaborative Security: A Game Theoretic Approach," *IEEE Symposium on Privacy-Aware Computing (PAC),* 2017.
- Xiaofan He, Mohammad Islam, Richeng Jin, and Huaiyu Dai, "Foresighted Deception in Dynamic Security Games," *IEEE International Conference on Communications (ICC)* 2017.
- Richeng Jin, Xiaofan He, and Huaiyu Dai, "On the Tradeoff between Privacy and Utility in Collaborative Intrusion Detection Networks - A Game Theoretical Approach," *Hot Topics in the Science of Security: Symposium and Bootcamp (HotSoS) 2017.*

Project: Vulnerability and Resilience Prediction Models

- Akond Rahman, Priysha Pradhan, Asif Partho, and Laurie Williams, "Predicting Android Application Security and Privacy Risk with Static Code Metrics," Short paper, *4th IEEE/ACM International Conference on Mobile Software Engineering and Systems*, Buenos Aires, Argentina, May 2017.
- Akond Rahman, Asif Partho, David Meder, and Laurie Williams, "Which Factors Influence Usage of Build Automation Tools?" *International Conference on Software Engineering (ICSE), 3rd International Workshop on Rapid Continuous Software Engineering (RCoSE)*, Buenos Aires, Argentina, May 2017.
- Donghoon Kim, Henry Schaffer, Mladen Vouk, "About PaaS Security," *International Journal of Cloud Computing*, 2017.

Project: Warning of Phishing Attacks: Supporting Human Information Processing, Identifying Phishing Deception Indicators, and Reducing Vulnerability

- Patrick Lawson and Christopher Mayhorn, "Interaction of Personality and Persuasion Tactics in Email Phishing Attacks," in *Proceedings of the Human Factors and Ergonomics Society 61st Annual Meeting*, Santa Monica, CA, Human Factors and Ergonomics Society.

Project: A Human Information-Processing Analysis of Online Deception Detection

- Aiping Xiong, Robert Proctor, Ninghui Li, and Weining Yang, "Is Domain Highlighting Actually Helpful in Identifying Phishing Webpages?" *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 2017.
- Weining Yang, Aiping Xiong, Jing Chen, Robert Proctor, and Ninghui Li, "Use of Phishing Training to Improve Security Warning Compliance: Evidence from a Field Experiment," *Hot Topics in the Science of Security: Symposium and Bootcamp (HotSoS), 2017.*
- Jing Chen, Robert Proctor, and Ninghui Li, "Human Trust in Automation in a Phishing Context," Talk presented at *46th Annual Meeting of the Society for Computers in Psychology (SCiP),* Boston, MA, November 2016.
- Aiping Xiong, Robert Proctor, Ninghui Li, and Weining Yang, "Use of Warnings for Instructing Users How to Detect Phishing Webpages," Talk presented at the *46th Annual Meeting of the Society for Computers in Psychology (SCiP),* Boston, MA, November 2016.

Project: Understanding Effects of Norms and Policies on the Robustness, Liveness, and Resilience of Systems

- Nirav Ajmeri, Hui Gu, Pradeep Murukannaiah, and Munindar Singh, "Arnor: Modeling Social Intelligence via Norms to Engineer Privacy-Aware Personal Agents," in *Proceedings of the 16th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, São Paulo, Brazil, May 2017, pages 1–9.

Project: Formal Specification and Analysis of Security-Critical Norms and Policies

- Thomas King, Akin Gunay, Amit Chopra, and Munindar Singh, "Tosca: Operationalizing Commitments Over Information Protocols," in *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, Melbourne, Australia, August 2017, pages 1-9.
- Munindar Singh and Amit Chopra, "The Internet of Things and Multiagent Systems: Decentralized Intelligence in Distributed Computing," in *Proceedings of the 37th IEEE International Conference on Distributed Computing Systems (ICDCS),* Blue Sky Thinking Track. Atlanta, GA, 2017, pages 1738–1747.
- Nirav Ajmeri, Chung-Wei Hang, Simon Parsons, and Munindar Singh, "Aragorn: Eliciting and Maintaining Secure Service Policies," *IEEE Computer 50(6)*, June 2017, pages 1–8.

- Özgür Kafali, Nirav Ajmeri, and Munindar Singh, "Kont: Computing Tradeoffs in Normative Multiagent Systems," in *Proceedings of the 31st Conference on Artificial Intelligence (AAAI)*, 2017.

Project: Privacy Incidents Database

- Pradeep Murukannaiah, Jessica Staddon, Heather Lipford, and Bart Kinijnenberg, "Is This a Privacy Incident? Using News Exemplars to Study End-user Perceptions of Privacy Incidents," *Usable Security Mini Conference (USEC)* 2017.
- Jessica Staddon, "Privacy Incidents Database: The Data Mining Challenges and Opportunities," *Cyber Security Practitioner*, November 2016.
- Karthik Sheshadri, Nirav Ajmeri, Jessica Staddon, "No (Privacy) News is Good News: An Analysis of New York Times and Guardian Privacy News from 2010 to 2016," *Proceedings of 15th Annual Conference on Privacy, Security and Trust (PST)*, August 2017, Calgary, pages 1-12.

Project: Redundancy for Network Intrusion Prevention Systems (NIPS)

- Sheng Liu, Michael Reiter, and Vyas Sekar, "Flow Reconnaissance via Timing Attacks on SDN Switches," *37th IEEE International Conference on Distributed Computing Systems,* 2017.

Project: Smart Isolation in Large-Scale Production Computing Infrastructures

- Rui Shu, Xiaohui Gu, and William Enck, "A Study of Security Vulnerabilities on Docker Hub," in *Proceedings of the ACM Conference on Data and Application Security and Privacy (CODASPY).*
- Haining Chen, Ninghui Li, William Enck, Yousra Aafer, and Xiangyu Zhang, "Analysis of SEAndroid Policies: Combining MAC and DAC in Android," *Proceedings of the Annual Computer Security Application Conference (ACSAC)*, 2017.

Project: Automated Synthesis of Resilient Architectures

- Mohammad Alsaleh and Ehab Al-Shaer, "Towards Automated Verification of Active Cyber Defense Strategies on Software Defined Networks," *ACM SafeConfig '16 Proceedings of the 2016 ACM Workshop on Automated Decision Making for Cyber Defense,* October 2016.

Project: Formal Analysis of Breach Reports

- Özgür Kafali, Jasmine Jones, Megan Petruso, Laurie Williams, and Munindar Singh, "How Good is a Security Policy against Real Breaches?: A HIPAA Case Study," in *Proceedings of the 39th International Conference on Software Engineering (ICSE),* 2017.
- Ricard López Fogués, Pradeep Murukannaiah, Jose Such, and Munindar Singh, "Understanding Sharing Policies in Multiparty Scenarios: Incorporating Context, Preferences, and Arguments into Decision Making," *ACM Transactions on Computer-Human Interaction (TOCHI),* 2017.

Research/Evaluation Team

- Morgan Burcham, Mahran Al-Zyoud, Jeffrey Carver, Mohammed Alsaleh, Hongying Du, Fida Gilani, Jun Jiang, Akond Rahman, Özgür Kafali, Ehab Al-Shaer, and Laurie Williams, "Characterizing Scientific Reporting in Security Literature: An Analysis of ACM CCS and IEEE S&P Papers," in *Proceedings of Hot Topics in the Science of Security: Symposium and Bootcamp (HotSoS) 2017*, pages 13-23.

EDUCATIONAL

Lablet funding has supported the training of multiple graduate students toward their MS and PhD degrees. The Lablet has driven these students toward principled, scientific approaches to their research. Additionally, students have contributed research efforts through independent study projects and collaborative activities with other non-Lablet research projects. The NCSU Lablet has also supported several Research Experiences for Undergraduates (REU) students who have made significant contributions. Students who participated in the Lablet have taken up research and faculty positions, including at RIT and William and Mary. Lablet research has influenced courses at multiple institutions as noted below.

Project: Attack Surface and Defense-in-Depth Metrics

- Our metrics research is shaping a course called Engineering Secure Software, a required course at RIT for all software engineering majors.

Project: Systematization of Knowledge from Intrusion Detection Models

- Based on the results from this work, we introduced cyber resiliency into graduate and undergraduate courses at UNCC.This work has also changed some of the ways that RIT teaches security and performance engineering in the Department of Software Engineering.
- Also at RIT, the discussion of trade-offs appears in a current course called Software Performance Engineering, and will help students discern the scientific literature on intrusion detection research.

Projects: Vulnerability and Resilience Prediction Models; Resilience Requirements, Design, and Testing

- Results from these projects have been incorporated into the security module of a cloud computing course at NCSU.

Project: Understanding Effects of Norms and Policies on the Robustness, Liveness, and Resilience of Systems

- We have used results from modeling the adoption of security practices in software development, and the sanctioning mechanisms that may promote greater adoption of these practices, among developers to guide new topics in our Social Computing course at NCSU.

Project: Formal Specification and Analysis of Security-Critical Norms and Policies

- We have used sociotechnical cybersecurity problems as exemplars for a combined graduate and undergraduate course in Social Computing at NCSU.

Project: Scientific Understanding of Policy Complexity

- We adopted problems and methods identified through our research on norms and breach reports as exemplars in an undergraduate course on Privacy.

Project: Smart Isolation in Large-Scale Production Computing Infrastructures

- We have added new class modules on cloud security in the graduate level classes the PIs regularly teach at NCSU.

Projects: Automated Synthesis of Resilient Architectures; High-Assurance Active Cyber Defense Policies for Auto-Resiliency

- Based on the results from work on these projects, we introduced reactive security policy into a graduate course at UNCC.

COMMUNITY OUTREACH

Lablet research has contributed to the development of datasets for research. Specifically,

- Our Privacy Incidents Database is an open research dataset, under development, of known privacy incidents.
- Our metrics research has contributed to the Vulnerability History Project, an open research dataset of vulnerabilities.

Lablet researchers gave numerous presentations at conferences and workshops emphasizing not only computer science but also other fields relevant to cybersecurity such as psychology and human factors. The following presentations promoted the Science of Security.

- “Privacy Incidents Database” at Indiana University’s Center for Applied Cybersecurity Research in October 2016.
- “Norms and analysis of breach reports” at the Carolina Privacy Officials Network (CPON) members and the International Association of Privacy Professionals (IAPP) meeting in April 2017.
- A keynote presentation which addressed Science of Security at the flagship software engineering conference, the Foundations of Software Engineering in September 2017.

Lablet researchers developed and extended collaborations with researchers outside of the four Lablets.

- We continued our international collaboration with researchers at Lancaster University in UK and at a number of institutions in the US.

Some Lablet research has garnered media attention.

- Our work on smart isolation was identified by The Morning Paper blog and on ACM’s official Twitter feed.

The Lablet hosted two workshops targeted at industry and government practitioners.

- We hosted our Industry Community Day in October, 2016. This event included 18 fast-paced student presentations, a poster session, and presentations from industry representatives to NCSU Lablet researchers. The goal of this annual event is to build collaborative relationships between our researchers and external organizations.

- We hosted our annual Healthcare IT Forum in April 2017. The forum included invited talks and moderated discussions between four industry influencers and NCSU Lablet researchers. We hosted a Summer Workshop in June 2017 for Lablet faculty, postdocs, and students, as well as guests from industry, academia, IT practitioners, and government, including the Laboratory for Analytic Sciences. The workshop program included keynote lectures, one each from industry and academia, as well as a panel of security professionals from government (HHS), nonprofits (RTI), the technology sector (Cisco), and the finance sector (Credit Suisse). The workshop included tutorials on scientific methodology as well as sessions on a methodological critique of recent security papers.

Additional engagement with industry and government:

- Lablet researchers had several discussions with, and gave presentations to, industry researchers and practitioners.
- Lablet researchers engaged with government researchers and practitioners, including those working for the NSA, the Federal Bureau of Investigation, and HHS. We also interacted with academic practitioners, such as in university and college-level IT organizations.

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

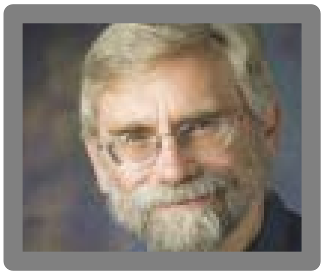

David Nichol

The University of Illinois at Urbana-Champaign (UIUC) Lablet is led by Principal Investigator (PI) David Nichol and lead researchers William Sanders and Sayan Mitra. In 2017 the UIUC Lablet continued to contribute to the development of security science by leveraging UIUC expertise in resiliency. 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 cyberattack. The Lablet's work also draws upon ideas from other areas of mathematics and engineering as well.

UIUC Sub-Lablets include Illinois Institute of Technology, Newcastle University, University of California at Berkeley, University of Southern California, University of Pennsylvania, Dartmouth College, and Rice University. UIUC, its Sub-Lablets, and collaborators are working on seven research projects, one of which was added in 2017 ("A Monitoring Fusion and Response Framework to Provide Cyber Resiliency"). Project by project details of the research are listed under the Fundamental Research section of the report.

The Science of Security has many attributes that range from use and development of scientific techniques in experimental security work, to modeling/mathematical foundations of systems where security and security properties are the object of the reasoning. UIUC contributes principally to the latter category with research that also supports the former category. We study how security properties are shaped by policy at different layers of the network stack, and use that to help define hypotheses that might be empirically tested. We are defining models of Cyber-Physical Systems that allow us to analyze how closely the system is allowed to skirt disaster, a measure of the system's resilience to disturbance. We are developing mathematical models of systems under attack, the attackers, and the defenders, to better understand how well the system is able to maintain required service levels through the attack, and to aid defensive decision-makers. We are applying sophisticated stochastic modeling techniques to describe vast volumes of data within which there are attacks; the models describe correlations between observations that might suggest attacks, and unobservable state that describes the attack. Finally, we are developing models of human behavior that seek to explain the how and why of humans circumventing security mechanisms. In short, the UIUC Science of Security research is exploring foundational mathematical modeling formalisms that quantitatively describe security attributes, and seek to predict those attributes as a function of context and environment.

FUNDAMENTAL RESEARCH

Project: A Hypothesis Testing Framework for Network Security
PI: Brighten Godfrey
Co-PIs: Matthew Caesar, David Nicol, William Sanders, Kevin Jin
Participating Sub-Lablet: Illinois Institute of Technology
Hard Problems: Scalability and Composability, Policy-Governed Secure Collaboration, Security Metrics and Models, Resilient Architectures

This project is developing the analysis methodology needed to support scientific reasoning about the security of networks, with a particular focus on information and data flow security. The core of this vision is Network Hypothesis Testing Methodology (NetHTM), a set of techniques for performing and integrating security analyses applied at different network layers, in different ways, to pose and rigorously answer quantitative hypotheses about the end-to-end security of a network. While our work touches on several Hard Problems, over the last year our key accomplishments focused on the Hard Problems of Scalability and Composability, Security Metrics and Models, and Resilient Architectures.

We have made progress on developing predictive security metrics with focus on predicting and verifying future behavior of networks including temporal properties. In real-world networks, correctness policies may be violated only through a particular combination of environment events and protocol actions, possibly in a non-deterministic sequence. However, tools in existence today are not capable of reasoning about all the possible network events, and all the subsequent execution paths that are enabled by those events.

We developed Plankton, a verification platform for identifying undesirable evolutions of networks. By combining symbolic modeling of data plane and control plane with explicit state exploration, Plankton performs a goal-directed search on a finite-state transition system that captures the behavior of the network as well as the various events that can influence it. In this way, Plankton can automatically find policy violations that can occur due to a sequence of network events, starting from the current state. An example use of the system would be verifying whether there exists a failure that could cause routes to change and circumvent a security control point or monitoring point, thus evading forensics. The system can prove whether such a dynamic event could occur, and if so, give an example.

We completed work on our project which ensures correct network virtualization. Current implementations lead to race conditions where, for example, a virtualized firewall, implemented behind the scenes at multiple physical locations, could erroneously block flows. We designed a system, COCONUT, which enables network elements to be automatically virtualized, i.e., implemented by a set of multiple physical elements which may be dynamic, while provably guaranteeing security properties.

We began work on two projects to automate improvements to software-defined networks, making them more robust and efficient. First, we developed an early implementation of NEAt (Network Error Auto-correct), a system that performs on-the-fly repair of updates that violate policies such as reachability, service chaining, and segmentation. NEAt sits between an SDN controller and the forwarding devices, intercepts updates proposed by SDN applications, and transforms a violating update into one that complies with the policy. Second, recognizing that commercial networks today have diverse security policies, we developed a framework that factors out the complexity of implementing security policies from the complexity of implementing performance optimizations, in the context of SDN controllers. Specifically, we are developing Oreo, a transparent performance enhancement layer for SDNs which guarantees that end-to-end reachability characteristics are preserved so security policies defined by the controller are not violated. Oreo performs these optimizations by using network modeling and verification mechanisms developed earlier in the NetHTM project.
To fully realize NetHTM, we need effective evaluation methodologies for large-scale and complex networked systems. We made advances in scalable evaluation methodology and platform using virtual-machine-based emulation and parallel simulation. We developed DSSNet (https://github.com/annonch/DSSnet), and utilized it to evaluate the SDN-based self-healing ability in critical energy systems and study the impact of various cyberattacks on network behavior.

We have also investigated resilient architectures for Industrial Control Systems (ICS) and, in particular, have taken infrastructure-level and application-level approaches to apply SDN technologies in ICS to make them more cyber secure and resilient. For example, we designed an SDN-based communication network architecture for microgrid operations and investigated multiple microgrid security applications, such as self-healing PMU, network verification, by leveraging the global visibility, direct networking controllability and programmability offered by SDN.

Project: Data-Driven-Model-Based Decision-Making
PIs: Bill Sanders, Masooda Bashir, David Nicol, Aad van Moorsel
Participating Sub-Lablet: Newcastle University, UK
Hard Problems: Security Metrics and Models, Human Behavior

Security analysis of complex systems is a challenging task, made especially difficult by the behavior of humans engaging the system. Rigorous mathematical models must be developed to capture the behavior of not only the autonomous aspects of these systems, but also the human participants.

We developed and continue to refine the HITOP modeling formalism to create models of human behaviors and decision-making. HITOP considers each human participant's opportunity, willingness, and capability to perform individual tasks throughout their typical daily routine. By studying HITOP models, we can better understand the critical role that good humans in the system impact the overall system performance and security. As with any model study, the quality of results is heavily dependent on the quality of input parameters. We have also developed a strategy for determining the data observation requirements necessary to ensure that useful, accurate quantitative metrics are produced.

Project: Data Driven Security Models and Analysis
PIs: Ravi Iyer, Zbigniew Kalbarczyk, Adam Slagell, Robin Sommer
Participating Sub-Lablet: University of California, Berkeley
Hard Problems: Security Metrics and Models, Resilient Architectures, Human Behavior

We have continued work on development of scientifically sound data-driven methods and tools with the goal of recognizing, mitigating, and containing attacks. The challenge is to capture and identify attackers' actions from the measurements, develop predictive models of attacker behavior before and during an attack, and thereby create a framework within which to reason about attacks, independent of the vulnerability exploited or the adopted attack pattern. This year we focused on the following:

- A novel application of Factor Graphs (probabilistic graphical models) to represent real-world multi-stage security incidents and develop methods for preemptive detection of attacks.
- Augmentation of the factor graph model based on analysis of recent credential stealing and infrastructure abuse attacks targeting Blue Waters, a petascale supercomputer hosted at the University of Illinois at Urbana-Champaign.
- Evaluation of detection capabilities of random factor functions (functions that return a random value when invoked in the factor graph evaluation) and factor functions defined based on the system knowledge and insights from the security experts. We found that detection performance of random factor functions has a substantially lower performance as compared to factor functions defined using system and domain experts.
- Introduction of mathematical underpinnings for automatic learning of factor functions. Specifically, we: 1) modelled multi-stage attacks; 2) classified factor functions by the number of input variables (univariate, bivariate, and multivariate) and the function body; and 3) developed a case study to demonstrate feasibility of automatic generation of factor functions.
- Indirect Cyberattacks: We investigated indirect cyberattacks against a large computing infrastructure through alteration of the CPS responsible for the cooling of the computing assets. We showed that a malicious user can attack a large computing infrastructure by compromising the environmental control systems in the facilities that host the compute nodes. We described real cases of failures due to problems in the cooling system of Blue Waters. We demonstrated, using real data, that the control systems that provide chilled water can be used as entry points by an attacker to indirectly compromise the computing functionality through the orchestration of clever alterations of sensing and control devices. In this way, the attacker does not leave any trace of his/her malicious activity on the nodes of the cluster. Failures of the cooling systems can trigger failure modes that can be recovered only after service interruption (including system reboot) and manual intervention.

Project: Science of Human Circumvention of Security
PIs: Tao Xie, Jim Blythe, Ross Koppel, Sean Smith
Participating Sub-Lablets: University of Southern California, University of Pennsylvania, Dartmouth College
Hard Problems: Scalability and Composability, Policy-Governed Secure Collaboration, Security Metrics and Models

We continue to study people's approaches to cybersecurity, and their use of authentication methods for accessing websites, their organization's databases, and the Internet. We focus especially on passwords as a prime method in the context of this trust (or suspicion or distrust). Use of passwords, adherence to password guidelines, and circumvention of password rules (e.g., sharing, writing them down on available files) are also excellent reflections of people's understanding, misunderstandings, and beliefs about personal and organizational efforts to protect individual and enterprise-level information. In addition, we are building and testing DASH agent models and beginning to test a mechanical Turk experiment/simulation to further examine users' use of passwords, workarounds, cyber trust, and strategies; measurements from the Turk experiment provide base calibration for the DASH model.

We have developed a new version of DASH in python that improves ease of development. We are also working with researchers at the University of Pennsylvania who have developed methods to learn agent behavior from observational data. To date, our results include constructing a semiotic framework for circumvention, validating our basic DASH model by reproducing behavior found in ground-truth human surveys, and duplication in our simulation of a version of "uncanny descent", in which making constraints on passwords more complex can decrease overall security. Last, we continue to administer two surveys: one on users' understanding of cybersecurity processes and their modes of circumvention; and one on security administrators' understanding of cybersecurity processes and their rationales for security policies and decisions. Also, to study people's trust in cybersecurity, especially mobile app security, we focus on collecting and analyzing UI text information faced by mobile app users to enable them to make informed decisions on mobile app security.

Project: Static-Dynamic Analysis of Security Metrics for Cyber-Physical System
PIs: Sayan Mitra, Geir Dullerud, Swarat Chaudhuri
Participating Sub-Lablet: Rice University
Hard Problems: Scalability and Composability, Security Metrics and Models

We have formulated the general problem of controller synthesis in the presence of resource constrained adversaries; namely, given an adversary of a certain class, parametrized according to the quantifiable resources available to them, we are creating a methodology to assess the worst-case potential impact and performance degradation of a control system from a threat of this class. We also consider the security and privacy implications for systems where the performance of the control system is optimized using data collected from human agents. Indeed, Cyber-Physical Systems (CPS) that contain physical agents have much more demanding and subtle security requirements than purely cyber-based systems, due to their significantly larger attack surface; that is, they can be attacked using all the approaches aimed at pure cyber-systems, but also in ways that are based on a deeper understanding of their physical dynamics and objectives, as well physically attacking agents and sensors themselves. Thus, new approaches for designing CPS need to take these potential attacks into account.

We have developed a sound and complete algorithm for solving this problem, for the special case of control systems with linear and monotonic dynamics and adversary resources characterized by their signal energy. The approach used to develop the algorithms brings together ideas from robust control and recent developments in syntax-guided program synthesis. Using our algorithms, we are able to synthesize controllers that are provably resilient to certain threat classes; in addition, we are also able to characterize the states of the systems in terms of their vulnerability levels.

We have considerably expanded our work on characterizing the trade-off between privacy and performance in CPS, particularly in cases where strategic preferences which govern dynamics are to be protected. In CPS, individual agents can interact through cloud-based distributed optimization, where each agent has an individualized objective function that is to be concealed to some level, while agents seek to minimize local objective functions and maintain global constraints. In this case, there are two challenges: First, the objective functions are used repeatedly in every iteration thus giving adversaries the possibility of a stream of measurements for estimating these objective functions, which once found accurately predict agent behavior. Secondly, the influence of perturbing objective functions in a traditional manner affects global system performance for all future times. During the past year we have completed significant work on this class of problems, and have developed analysis results on the propagation of perturbations on objective functions over time, and show how to derive upper bounds. With this, a noise-adding mechanism can be designed that randomizes the cloud-based distributed optimization algorithm to keep the individual objective function concealed. In addition, we have considered the trade-off between the secrecy of objective functions and the performance of the new cloud-based distributed optimization algorithms with noise. Such distributed optimization problems appear in various applications including energy systems, communications, signal processing, sensor networks, and machine learning, as well as in mobile robotics involving autonomous agents. In some of these applications it is common to have agents in a network directly share information with each other in the process of optimizing. However, in systems in which nefarious parties may be acting, when the agents share information with each other, a major concern is the concealing of their objectives and current situational states.

Our work on security and differential privacy constituted the very first published research in this area for CPS, and there is now an entire sub-community with the field of control theory working on these issues for CPS, and importantly, the security variants.

Project: Anonymous Messaging
PIs: Pramod Viswanath, Carl Gunter
Hard Problem: Scalability and Composability

This project is focused on the foundations of algorithms that broadcast information on networks efficiently and anonymously. We are particularly interested in applications to social networks and cryptocurrency networks.

Social networks: Anonymous social media platforms like Secret, Yik Yak, and Whisper have emerged as important tools for sharing ideas without the fear of judgment. Such anonymous platforms are also important in nations under authoritarian rule, where freedom of expression and the personal safety of message authors may depend on anonymity. Whether for fear of judgment or retribution, it is sometimes crucial to hide the identities of users who post sensitive messages. In this research, we consider a global adversary who wishes to identify the author of a message: it observes either a snapshot of the spread of a message at a certain time, sampled timestamp metadata, or both. Recent advances in rumor source detection show that existing messaging protocols are vulnerable against such an adversary. Our main technical contribution is the introduction of a novel messaging protocol, which we call adaptive diffusion, and we show that under the snapshot adversarial model, adaptive diffusion spreads content fast and achieves perfect obfuscation of the source when the underlying contact network is an infinite regular tree. That is, all users with the message are nearly equally likely to have been the origin of the message. When the contact network is an irregular tree, we characterize the probability of maximum likelihood detection by proving a concentration result over Galton-Watson trees. Experiments on a sampled Facebook network demonstrate that adaptive diffusion effectively hides the location of the source even when the graph is finite, irregular and has cycles. An Android implementation of our algorithm (titled WILDFIRE) is publicly available.

Cryptocurrency networks: Bitcoin and other cryptocurrencies have surged in popularity over the last decade. Although Bitcoin does not claim to provide anonymity for its users, it enjoys a public perception of being a "privacy-preserving" financial system. In reality, cryptocurrencies publish users' entire transaction histories in plaintext, albeit under a pseudonym; this is required for transaction validation. Therefore, if a user's pseudonym can be linked to their human identity, the privacy fallout can be significant. Recently, researchers have demonstrated deanonymization attacks that exploit weaknesses in the Bitcoin network's Peer-to-Peer (P2P) networking protocols. In particular, the P2P network currently forwards content in a structured way that allows observers to deanonymize users. In this work, we redesign the P2P network from first principles with the goal of providing strong, provable anonymity guarantees. We propose a simple networking policy called Dandelion, which achieves nearly-optimal anonymity guarantees at minimal cost to the network's utility. We also provide a practical implementation of Dandelion which is freely available on Github and is in the process of being implemented inside Bitcoin Core (which is the most popular version of Bitcoin).

Project: A Monitoring Fusion and Response Framework to Provide Cyber Resiliency

PI: Bill Sanders
Hard Problems: Policy-Governed Secure Collaboration, Resilient Architectures

Resilience has become a key strategy for protecting cyber systems. Although traditional cyber security protection mechanisms are an important component of an overall cybersecurity strategy, they are no longer sufficient for systems that must provide continuous service when under attack. Resiliency mechanisms offer a synergistic approach to securing systems based the realization that protection mechanisms are not perfect. We have developed a methodology for deploying a diverse set of monitors within a system, at different locations and at different levels in the system architecture, to serve as input to fusion and alert correlation algorithms whose goal is to detect attacks. We have also developed several fusion algorithms that could provide attack alerts to a set of response selection algorithms. We are exploring response selection algorithms that utilize game theory and control theory to find good or optimal response strategies. To evaluate our developed response algorithms, we have use a two-pronged strategy using real data and discrete event simulation. In particular, we intend to simulate attacker behavior with models learned using real attack data, where the attacker model is pitted against the response selection algorithms in a simulated environment.

PUBLICATIONS

Project: A Hypothesis Testing Framework for Network Security

- Christopher Hannon, Dong Jin, Chen Chen, Jianhui Wang, Cheol Won Lee, and Jong Cheol Moon, "Ultimate Forwarding Resilience in OpenFlow Networks," in the *ACM International Workshop on Security in Software Defined Networks & Network Function Virtualization (SDN-NFV Security 2017)*, Scottsdale, AZ, March 24, 2017.
- Santhosh Prabhu, Mo Dong, Tong Meng, Brighten Godfrey, and Matthew Caesar, "Let Me Rephrase That: Transparent Optimization in SDNs," *ACM Symposium on SDN Research (SOSR 2017)*, Santa Clara, CA, April 3-4, 2017.
- Soudeh Ghorbani and Brighten Godfrey, "COCONUT: Seamless Scale Out of Network Elements," *Twelfth European Conference on Computer Systems (EuroSys 2017)*, Belgrade, Serbia, April 23-26, 2017.
- Jiaqi Yan, Xin Liu, and Dong Jin, "Simulation of a Software-Defined Network as One Big Switch," *2017 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (ACM SIGSIM PADS)*, Singapore, May 24-26, 2017.
- Santhosh Prabhu, Ali Kheradmand, Brighten Godfrey, and Matthew Caesar, "Predicting Network Futures with Plankton," *1st Asia-Pacific Workshop on Networking (APNet'17)*, Hong Kong, China, August 3-4, 2017.
- Ning Liu, Adnan Haider, Dong Jin and Xian He Sun, "A Modeling and Simulation of Extreme-Scale Fat-Tree Networks for HPC Systems and Data Centers," *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, to appear.
- Dong Jin, Zhiyi Li, Christopher Hannon, Chen Chen, Jianhui Wang, Mohammad Shahidehpour, Cheol Won Lee and Jong Cheol Moon, "Towards a Resilient and Secure Microgrid Using Software-Defined Networking," *IEEE Transactions on Smart Grid, Special section on Smart Grid Cyber-Physical Security*, to appear.
- Xin Liu and Dong Jin, "ConVenus: Congestion Verification of NetworkUpdates in Software-defined Networks", *Winter Simulation Conference (WSC2016)*, Washington, DC, December 11-14, 2016.

Project: Data-Driven Model-Based Decision-Making

- John C. Mace, Nippun Thekkummal, Charles Morisset, and Aad van Moorsel, "ADaCS: A tool for Analysing Data Collection Strategies," *14th European Performance Engineering Workshop (EPEW 2017)*, Berlin, Germany, September 7-8, 2017.

Project: Science of Human Circumvention of Security

- Ross Koppel, Jim Blythe, Vijay Kothari, and Sean Smith, "Password Logbooks and What Their Amazon Reviews Reveal About Their Users' Motivations, Beliefs, and Behaviors," *2nd European Workshop on Usable Security (EuroUSEC 2017)*, Paris, France, April 29, 2017.
- Ross Koppel and Harold Thimbleby, "Lessons from the 100 Nation Ransomware Attack," May 14, 2017 The HealthCare Blog (THCB) http://thehealthcareblog.com/
- Haibing Zheng, Dengfeng Li, Xia Zeng, Beihai Liang, Wujie Zheng, Yuetang Deng, Wing Lam, Wei Yang, and Tao Xie, "Automated Test Input Generation for Android: Towards Getting There in an Industrial Case," *39th International Conference on Software Engineering (ICSE 2017)*, Software Engineering in Practice (SEIP), Buenos Aires, Argentina, May 20-28, 2017.
- Christopher Novak, Jim Blythe, Ross Koppel, Vijay Kothari, and Sean Smith, "Modeling Aggregate Security with User Agents that Employ Password Memorization Techniques," *Who Are You?! Adventures in Authentication (WAY 2017), workshop in conjunction with Symposium On Usable Privacy and Security (SOUPS 2017)*, July 12-14, 2017, Santa Clara, CA.
- Benjamin Andow, Akhil Acharya, Dengfeng Li, William Enck, Kapil Singh, and Tao Xie, "UiRef: Analysis of Sensitive User Inputs in Android Applications," *10th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec 2017)*, Boston, MA, July 18-20, 2017.
- Ross Koppel, Vijay Kothari, Sean W. Smith, and Jim Blythe, "Beyond Pleading With or Restricting Users to Achieve Cyber Security Goals: Approaches to Understanding and Responding to Circumvention", *CRA CCC Sociotechnical Cybersecurity Workshop*, College Park, MD, December 12-13, 2016.
- Sean W. Smith, Vijay Kothari, Jim Blythe, and Ross Koppel, "Flawed Mental Models Lead to Bad Cyber Security Decisions: Let's Do a Better Job", *CRA CCC Sociotechnical Cybersecurity Workshop*, College Park, MD, December 12-13, 2016.

Project: Static-Dynamic Analysis of Security Metrics for Cyber-Physical Systems

- Zhenqi Huang, Yu Wang, Sayan Mitra, and Geir Dullerud, "Differential Privacy and Entropy in Distributed Feedback Systems: Minimizing Mechanisms and Performance Trade-offs," *IEEE Transactions on Network Control Systems*, volume 4, issue 1, March 2017. (Jul 17)
- Hussein Sibaie and Sayan Mitra, "Optimal Data Rates for Estimation and Model Detection of Switched Dynamical Systems," *20th ACM International Conference on Hybrid Systems: Computation and Control in conjunction with CPS Week 2017*, Pittsburgh, PA, April 18-21, 2017.
- Joao Jansch Porto and Geir Dullerud, "Decentralized Control with Moving-Horizon Linear Switched Systems: Synthesis and Testbed Implementation," *American Control Conference 2017*, Seattle, WA, May 24-26, 2017.
- Yu Wang, Sayan Mitra, and Geir Dullerud, "Differential Privacy and Minimum-Variance Unbiased Estimation in Multi-Agent Control Systems," *IFAC World Congress*, Toulouse, France, July 9-14, 2017.
- Yu Wang, Zhenqi Huang, Sayan Mitra, and Geir Dullerud, "Differential Privacy in Linear Distributed Control Systems: Entropy Minimizing Mechanisms and Performance Tradeoffs", *IEEE Transactions on Control of Network Systems*, volume 4, issue 1, January 25, 2017.
- Yu Wang, Matthew Hale, Magnus Egerstedt, and Gier Dullerud, "Differentially Private Objective Functions in Distributed Cloud-based Optimization", *55th Conference on Decision and Control (CDC 2016)*, Las Vegas, NV, December 12-14, 2016.

Project: Anonymous Messaging

- Guilia Fanti, Shaileshh Venkatakrishnan and Pramod Viswanath, "Dandelion: Redesigning BitCoin Networking for Anonymity," *ACM Sigmetrics 2017*, Urbana, IL, June 5-9, 2017.

Project: A Monitoring Fusion and Response Framework to Provide Cyber Resiliency

- Carmen Cheh, Binbin Chen, William Temple, and William Sanders, "Data-Driven Model-Based Detection of Malicious Insiders via Physical Access Logs," *14th International Conference on Quantitative Evaluation of Systems (QEST 2017)*, Berlin, Germany, September 5-7, 2017.
- Atul Bohara, Mohammad Noureddine, Ahmed Fawaz, and William Sanders, "An Unsupervised Multi-Detector Approach for Identifying Malicious Lateral Movement," *36th IEEE International Symposium on Reliable Distributed Systems (SRDS 2017)*, Hong Kong, September 26-29, 2017.
- Benjamin Ujcich, Andrew Miller, Adam Bates, and William Sanders, "Towards an Accountable Software-Defined Networking Architecture," *3rd IEEE Conference on Network Softwarization (NetSoft 2017)*, Bologna, Italy, July 3-7, 2017.

EDUCATIONAL

A Hypothesis Testing Framework for Network Security

Brighten Godfrey covered network security in his graduate course, Advanced Computer Networking, including quantitative aspects of BGP security and formal verification of networks. These topics span lectures, reading, and student research projects developing new techniques for formal reasoning about networks.

This project team has been actively working on dissemination of knowledge through tutorials on network verification. Brighten Godfrey developed and presented a half-day tutorial at a workshop at Hebrew University. Santhosh Prabhu and Brighten Godfrey submitted a proposal to present an expanded tutorial at the IEEE/ACM International Conference on Software Engineering (ASE) in October 2017; this proposal was accepted.

Anonymous Messaging in Networks

Giulia Fanti and Pramod Viswanath gave a tutorial, "Information Limits on Finding and Hiding Message Sources on Networks: Social Media and Cryptocurrencies" at the IEEE International Symposium on Information Theory (ISIT) in Aachen, Germany on June 25, 2017.

Science of Human Circumvention of Security

Tao Xie attended the 2017 National Society of Black Engineers (NSBE) Convention during March 30-April 1, 2017, where he held discussions with a large number of participants (including his mentees) on various computer science problems including security problems.

UIUC SoS Lablet

We again held a SoS summer internship program with students from University of Arkansas, University of Maryland, and the University of Illinois at Urbana-Champaign. The interns worked on self-proposed research projects with a UIUC advisor. The internship program included educational programming in conjunction with other summer internship programs within the Illinois College of Engineering. This summer the interns presented to the NSA Information Assurance Research group, SoS Program Manager, and concluded the summer program with a poster session.

COMMUNITY OUTREACH

Project: A Hypothesis Testing Framework for Network Security

- Santhosh Prabhu: "Oreo: Transparent Optimization to Enable Flexible Policy Enforcement in Software Defined Networks," ITI Joint Trust and Security/Science of Security Seminar, October 2016.
- Kevin Jin: "Enabling a Cyber-Resilient and Secure Energy Infrastructure with Software-Defined Networking," Monthly UIUC/Information Assurance Research Presentation, January 2017.
- Christopher Hannon: "Ultimate Forwarding Resilience in OpenFlow Networks," technical presentation, ACM International Workshop on Security in Software Defined Networks & Network Function Virtualization (SDN-NFV Security 2017), March 2017.
- Christopher Hannon: "Securing the Smart Grid with Software Defined Networking," Monthly UIUC/Information Assurance Research Presentation, August 2017.

Project: Data Driven Security Models and Analysis

- Phuong Cao: "Automated Generation of Attack Signatures in Attack Graphs," Joint Trust and Security/Science of Security Seminar, November 2016.

Project: Science of Human Circumvention of Security

- Wing Lam, Dengfeng Li, and Wei Yang: "Towards Privacy-Preserving Mobile Utility Apps: A Balancing Act," Monthly UIUC/Information Assurance Research Presentation, February 2017.
- Tao Xie: "User Expectations in Mobile App Security," invited seminar, IEEE Rochester Section CS/CIS Joint Chapters/Department of Computing Security, Rochester Institute of Technology, March 2017.
- Jim Blythe, Ross Koppel, Sean Smith, Vijay Kothari, David Harmon, and Christopher Novak: "A Cross-Disciplinary Study of User Circumvention of Security," Monthly UIUC/Information Assurance Research Presentation, March 2017.

Project: Static-Dynamic Analysis of Security Metrics for Cyber-Physical Systems

- Sayan Mitra: “Auditing Algorithms,” Frontiers Seminar, Master of Technology Management Program, University of Illinois at Urbana-Champaign Business School, December 2016.
- George Pappas, Jerome Le Ny, Geir Dullerud, and Jorge Cortes: “Differential Privacy in Control and Network Systems,” invited tutorial, 55th Conference on Decision and Control, December 2016.

Project: Anonymous Messaging

- Giulia Fanti: “Anonymity in the Bitcoin Peer-to-Peer Network,” Lablet presentation, NSA SoS Quarterly Meeting, February 2017; ITI Joint Trust and Security/ Science of Security Seminar, February 2017; Illinois Bitcoin Meetup, Jump Labs at Research Park, University of Illinois of Illinois at Urbana-Champaign, February 2017; Security Seminar, Computer Science, University of California Berkeley, March 2017; ISL Colloquium, EE Department, Stanford University, March 2017; EE Department Seminar, University of Wisconsin, March 2017.
- Nitin Vaidya: “Privacy & Security in Machine Learning/ Optimization,” Lablet presentation, NSA SoS Quarterly Meeting, February 2017
- Shaileshh Venkatakrishnan: “Dandelion: Redesigning the Bitcoin Peer-to-Peer Network for Anonymity,” Security Seminar, MIT, March 2017.

Project: A Monitoring Fusion and Response Framework to Provide Cyber Resiliency

- Ahmed Fawaz: “PowerAlert: An Integrity Checker using Power Measurement,” Monthly UIUC/Information Assurance Research Presentation, November 2016.
- Atul Bohara: “A Framework for Detection and Containment of Lateral Movement-Based Attacks,” Monthly UIUC/Information Assurance Research Presentation, December 2016.
- Carmen Cheh: “Data-Driven Model-Based Detection of Malicious Insiders via Physical Access Logs,” Monthly UIUC/Information Assurance Research Presentation, January 2017.
- Uttam Thakore: “Prioritization of Cloud System Monitoring for Incident Response,” Monthly UIUC/Information Assurance Research Presentation, March 2017.
- Benjamin Ujcich: “Accountable SDNs for Cyber Resiliency,” Monthly UIUC/Information Assurance Research Presentation, March 2017.
- Mohammad Noureddine: “A Comprehensive Framework for DDoS Resiliency in the Cloud,” Monthly UIUC/ Information Assurance Research Presentation, April 2017.
- Ahmed Fawaz: “Lateral Movement Detection and Response,” Monthly UIUC/Information Assurance Research Presentation, June 2017.

Conference Panel and Committee Members

- Vijay Kothari, Panel Moderator; Ross Koppel, Panelist, “On Developing Authentication Solutions for Healthcare Settings,” SOUPS WAY Workshop, 2017
- Kevin Jin, Program CO-Chair, ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (PADS), 2017
- Brighten Godfrey, Program Committee Member, ACM SIGMETRICS, 2017
- Brighten Godfrey, Program Committee Member, ACM HoTNets, 2017
- Brighten Godfrey, Program Committee Member, ACM SIGCOMM, 2017: Workshop on Virtual Reality and Augmented Reality Network (VR/AR Network 2017)

UNIVERSITY OF MARYLAND

Jonathan Katz

The University of Maryland (UMD) Lablet, led by co-PIs Jonathan Katz and Michel Cukier, involves ten projects looking at different aspects of the five Hard Problems, with specific focus on Policy-Governed Secure Collaboration, Security Metrics and Models, and Human Behavior. The UMD Lablet consists of twenty faculty from both UMD and partner institutions. The fifteen UMD faculty are drawn from five different departments across campus, including Computer Science, Electrical and Computer Engineering, Information Studies, Criminology, and Reliability Engineering. The collaborators hail from the US Naval Academy (USNA), Virginia Polytechnic Institute and State University (Virginia Tech), University of Texas at Austin, Indiana University, and The George Washington University.

FUNDAMENTAL RESEARCH

Project: Understanding Developers' Reasoning about Privacy and Security
PIs: Michelle Mazurek, Charalampos Papamanthou, Mohit Tiwari
Participating Sub-Lablet: University of Texas at Austin
Hard Problem: Human Behavior

Techniques such as Information Flow Control can offer strong privacy guarantees but have failed to achieve traction among developers. This project has been developing an alternative scheme called Blox in which developers partition their apps based on functionality (analogous to a model-view-controller pattern) instead of using labels and information flow compilers. The team members are conducting studies to understand how users partition access to their content in cloud applications, to better understand the right abstractions that developers should use. A pilot study evaluated the usability of the Blox platform based on whether it is possible to accurately infer access-control domains using machine-learning techniques. The study demonstrated that the learning techniques are not yet sufficiently accurate, and may require additional supervision to gain accuracy. The team has been in the process of developing a measurement study to better understand data sharing in cloud platforms. Using the infrastructure platform developed for the first study, the team will collect usage data across multiple services and obtain mappings about how many items are shared with how many people, in what type of groups, and with what longevity. This will enable a characterization of the modern cloud-based access-control space, updating prior work that examined corporate and peer-to-peer networks.

Project: Measuring and Improving the Management of Today's PKI
PI: Dave Levin
Hard Problems: Security Metrics and Models, Human Behavior

This project focuses on metrics by means of large-scale measurements of the existing Public-Key Infrastructure (PKI) used in today's web. While online use of PKI is mostly automated, there is a surprising amount of human intervention in management tasks that are crucial to its proper operation. This project studies the following questions: Are administrators doing what users of the Web need them to do in order to ensure security?; and What can be done to help facilitate or automate these tasks? As part of this project, researchers are performing internet-wide measurements of how online certificates are actively being managed, including how quickly and thoroughly administrators revoke their certificates after a potential key compromise, and what role third-party hosting services play. In particular, they find that Content Distribution Networks (CDNs), which serve content for many of the most popular websites, have access to content providers' private keys, violating the fundamental assumption of PKI (i.e., that no one shares their private keys). We are performing the first widespread analyses of the extent to which websites are sharing their private keys, and exploring what impact this has on the management of the PKI and on users' privacy and security in general. The research group is also developing new systems that help improve clients' ability to stay up-to-date on certificate revocations. They are developing systems that leverage recent initiatives, such as Certificate Transparency, to more compactly represent revocation data. One such system, CRLite, uses a novel data structure that clients can query to determine whether or not a certificate is revoked; surprisingly, CRLite is able to do so at a cost of less than one byte per certificate. This shows that universal coverage of certificate revocations may at last be within reach.

Project: Trust, Recommendation Systems, and Collaboration
PIs: John Baras, Jennifer Goldbeck
Hard Problems: Scalability and Composability, Policy-Governed Secure Collaboration, Human Behavior

The overarching goal of the project is to develop a transformational framework for a science of trust, and assess its impact on local policies for collaboration in networked multi-agent systems. The framework takes human behavior into account from the start and is validated experimentally. Our work has led to novel results regarding the following: the evolution of opinions (or beliefs) over a social network modeled as a signed graph; new models and analytical methods for the investigation of consensus dynamics with both collaborative and non-collaborative node interactions; and new probabilistic models of multi-domain crowdsourcing tasks. The team has also formalized the problem of trust-aware task allocation in crowdsourcing and developed a principled way to solve it. The formulation models the workers' trustworthiness and the costs based on both the question and the worker group. In other work, the team has developed a new framework for modeling trust based on their recently developed foundational model for networked multi-agent systems in which they consider three interacting dynamic, directed graphs on the same underlying set of nodes: a social/agent network, which is relational; an information network, also relational; and a communication network that is physical. The links and nodes are annotated with dynamically changing "weights" representing trust metrics whose formal definition and mathematical representation can take one of several options, e.g., weights can be scalars, vectors, or even policies (i.e., rules). Within this new framework, the team is specifically focusing on the following fundamental problems: 1) theories and principles governing the spreading dynamics of trust and mistrust among members of a network; 2) design and analysis of recommendation systems, their dynamics and integrity; 3) development of a framework for understanding the composition of trust across various networks at the different layers of our basic model; and 4) analysis of the effects of trust on collaboration in networked multi-agent systems, using game-theoretic and economic principles.

Project: User-Centered Design for Security
PIs: Jennifer Goldbeck, Adam Aviv
Participating Sub-Lablet: United States Naval Academy
Hard Problems: Security Metrics and Models, Human Behavior

The overarching goal of the project was to better understand human behavior within security systems, and to use that knowledge to propose, design, and build better security systems. A system that is designed taking into account limitations on human memory, attention, and cognitive abilities will be easier to use and will thus lead people toward acting in secure ways. Conversely, systems that force users to carry out inherently difficult tasks lead people to circumvent security guidelines in order to get their tasks done efficiently. The team undertook several efforts in this space, in particular in understanding the security and usability of text and graphical passwords, as well as user opinions and expertise about security-related issues. In the password space, the team designed mechanisms to help people remember passwords more effectively, made progress in understanding mental models of privacy applied to mobile devices and how those models may affect the choice of mobile-authentication technique, and measured the strength of authentication systems to human attackers. Most recently, the latter project focused on shoulder-surfing attacks. Our work in this space has improved the understanding of human-password interaction and how authentication systems can be designed to be more usable and more secure.

The team also studied what users understand about how their personal information is used, how comfortable they are with that, and the role that consent and control play in their opinions. Numerous research projects have documented concerns that users have with data commonly used by recommender systems. We conducted several studies on this topic and found that users are uncomfortable with much data currently used in personalization technologies, that they don't know how to secure that data, and that consent is a critical component to their level of comfort.

Project: Reasoning about Protocols with Human Participants
PIs: Jonathan Katz, Poorvi Vora
Participating Sub-Lablet: George Washington University
Hard Problem: Human Behavior

The aim of this project is to study protocols (in particular, electronic-voting protocols) in which humans are explicitly modeled as participants. In the last year, the team has described security vulnerabilities in the remote voting system Helios, including one that allowed a dishonest voting terminal to change a voter's vote after it obtains the voter's credential. The team also proposed Apollo, a modified version of Helios, that addresses those vulnerabilities. With Apollo-lite, votes not authorized by the voter are detected by the public and prevented from being included in the tally. The full version of Apollo enables a voter to prove that her vote was changed. We also describe a very simple protocol for the voter to interact with any devices she employs to check on the voting system, to enable frequent and easy auditing of encryptions and checking of the bulletin board. Apollo uses some of the ideas of Remotegrity, and we are working on a common framework for definitions and proofs for Remotegrity and Apollo. As part of this project, Vora has also developed a taxonomy of voting systems using some new and some existing definitions, and applied the taxonomy to some of the more prominent voting systems in a survey paper.

Project: Empirical Models for Vulnerability Exploits
PI: Tudor Dimitras
Hard Problem: Security Metrics and Models

This project is exploring more informative metrics to quantify security of deployed systems. Work this past year had two thrusts: understanding and mitigating the misuse of cryptographic APIs, and characterizing the utility of hardware or virtualization indicators for detecting attacks against cloud-computing infrastructures.

In the first thrust, the team inferred five developer needs and showed that a good API would address those needs only partially. Building on this observation, the team proposed APIs that are semantically meaningful for developers, showed how the necessary interfaces can be implemented consistently on top of existing frameworks, and proposed build-management hooks for isolating security workarounds needed during the development and test phases. Through two case studies, the team showed that those APIs can be utilized to implement non-trivial client-server protocols and that they provide a better separation of concerns than existing frameworks. In the second thrust, the team investigated the information provided by hardware or virtualization indicators that could be utilized for detecting attacks against cloud-computing infrastructures. In those settings, the service providers are contractually prohibited from accessing the content of customer virtual machines, which makes it challenging for them to protect their infrastructures from malware infections. The research group extracted 614 features from traces generated by 3 tools, and analyzed them to extract the number, density, and distance between peaks in the signal. The group is currently investigating the extent to which these sub-semantic features are useful for detecting malicious activity in customer virtual machines. Early results are promising: on a ground truth with 529 malware samples and 529 benign programs, classifiers can be trained with accuracies above 99%.

Project: Human Behavior and Cyber Vulnerabilities
PI: V.S. Subrahmanian
Participating Sub-Lablet: Virginia Tech
Hard Problems: Security Metrics and Models, Human Behavior

This project had three thrusts over the past year. In the first thrust, a system called FeatureSmith was developed that performs automatic feature engineering by mining the security research literature. The effectiveness of machine-learning techniques primarily depends on the manual feature engineering process, which has traditionally been based on human knowledge and intuition. However, given attackers' efforts to evade detection and the growing volume of security reports and publications, human-driven feature engineering likely draws from only a fraction of the relevant knowledge. The team developed methods to engineer such features automatically, by mining natural-language documents such as research papers, industry reports, and hacker forums. As a proof of concept, the research group used this approach to train a classifier with automatically engineered features for detecting Android malware, and achieved performance comparable to that of a state-of-the-art malware detector that uses manually engineered features.

In the second thrust, the group finished analysis of their Spring 2016 study of mobile users' preferences towards autoupdates on their phones. This study, involving a survey of 550 Android users, found that those who do not autoupdate their applications tend to take fewer risks, are more security aware, and have had a previous negative experience with software updates. Users' preferences towards autoupdating were also found to be influenced by how satisfied they are with a mobile application, how important the application is to them, and how much they trust the application itself. Finally, results showed that users are more likely to want to autoupdate due to security updates rather than non-security-related updates. These findings led to several recommendations to improve notifications to encourage users to switch to autoupdates.

In the third thrust, we analyzed the performance of malware-family detection techniques. Android is the most widely used mobile OS today, and is also the biggest target for mobile malware. Given the increasing number of malware variants, detecting malware families is crucial for security analysts to reuse signatures of a known family to tackle new malware belonging to that family. In late 2016 we developed a thorough and systematic performance comparison of several traditional classification algorithms for the task of detecting Android malware families. We perform our evaluation on DREBIN, the largest public Android malware dataset with labeled families, on which we extract both static features from the code and dynamic features by executing malware samples in a controlled sandbox along with a network simulator. In particular, we defined a large set of features based on both static and dynamic code analysis, and showed that as long as the malware family contained at least 10 samples of malware variants in the training data, we could predict the families to which unlabeled samples belong with high accuracy, irrespective of the accuracy measures considered. Specific accuracy using microF score, macroF score, microAUC and macroAUC were 0.95, 0.89, 0.97, 0.93 respectively.

Project: Does the Presence of Honest Users Affect Intruders' Behavior?
PIs: Michel Cukier, David Maimon
Participating Sub-Lablet: University of Texas at Austin
Hard Problem: Human Behavior

In this project, Michel Cukier and David Maimon are applying criminological techniques to develop a better understanding of attacker behavior. One particular highlight of the past year is the examination of previously uninvestigated experimental data: an experiment that randomly assigned infiltrated target computers to have a certain type (administrative or non-administrative) and number (1 or 10) of users to appear on the system at the same time as the system trespasser. Using this data, the team examined whether the number and type of users present on a system reduced the seriousness and frequency of trespassing. Results indicated that the presence of an administrative user (regardless of the number of users) significantly reduced the number of system trespassing events. Additionally, with 10 users present, the presence of an administrative user significantly reduced the total amount of time attackers spent on the compromised system. Interestingly, comparing between conditions with different numbers of users, it was found that the number of users present on the system has no effect on the number of trespassing events or total time spent on the system. These findings together indicate that presence of an administrative user can produce a deterrent effect on system trespassers, while the number of users present on the system has no effect on system trespasser actions. Findings from these analyses were reported and presented during the Hot Topics in the Science of Security (HoTSoS) 2017 annual conference. Their poster "Is the Guardian Capable? A Routine Activity Theory Approach to Cyber Intrusion on Honeypot Systems" was named Best Poster at HotSoS 2017.

Project: Understanding How Users Process Security Advice
PI: Michelle Mazurek
Hard Problem: Human Behavior

This project addresses the Hard Problem of Human Behavior from the perspective of educational efforts. People encounter a tremendous amount of cybersecurity advice. It would be impossible to follow all the available advice, so people pick and choose which advice to follow and which to ignore in different circumstances. But the advice they pick is not always the most correct or useful. This project examines where people encounter security advice, how they evaluate its trustworthiness, and how they decide which advice to follow or reject. This year, we conducted a large-scale quantitative survey of how users learn about security advice. We also analyzed random-digit-dial national survey data touching on security advice and behaviors, finding a relationship between education level and advice sources. Based on results from last year's work, we hypothesized that edutainment could be an effective mechanism for motivating adoption of software updates. We conducted participatory design workshops to develop a storyline, which was professionally produced as a video. We used a longitudinal study to evaluate edutainment's effect on attitudes toward updates, finding a small but significant improvement for edutainment text relative to traditional security text. We also used participatory design to develop new approaches to motivating Two-Factor Authentication (2FA) adoption. We are conducting preliminary tests on the new design and hoping to test it in the field via a collaboration with a large software company.

Project: Trustworthy and Composable Software Systems with Contracts
PIs: David Van Horn, Jeffrey Foster, Michael Hicks, Sam Tobin-Hochstadt
Participating Sub-Lablet: Indiana University
Hard Problem: Scalability and Composability

As part of this project, researchers are investigating compositional-verification techniques using language-based mechanisms called contracts for specifying and enforcing program properties. Over the past 15 months, we have applied the technique to multi-language programs, security properties, and imperative languages. We have developed a foundational theory, extended it to these settings, and empirically evaluated the effectiveness of our prototype analysis tools. We made an earlier theoretical breakthrough showing how to generate counterexamples that witness contract violations. This is important for testing and debugging software that uses contracts. We have proved that our method is both "sound" and "relatively complete," which means the approach is powerful and capable of generating a large class of counterexamples. These results were first established in a purely functional setting, but over the past 15 months we have extended the result to higher-order imperative settings. We have worked on integrating this verification technique into a full-featured programming language and interactive development environment, and developed pedagogical tools for teaching verified software development. Currently, these tools are being used in an experimental section of the UMD introductory programming sequence.

PUBLICATIONS

Project: Measuring and Improving the Management of Today's PKI

- Taejoong Chung, Roland van Rijswijk-Deij, Balakrishnan Chandrasekaran, David Choffnes, Dave Levin,

- Bruce M. Maggs, Alan Mislove, and Christo Wilson, “Longitudinal, End-to-End View of the DNSSEC Ecosystem,” *USENIX Security 2017*. This paper received a distinguished paper award at the conference.
- James Larisch, David Choffnes, Dave Levin, Bruce M. Maggs, Alan Mislove, and Christo Wilson, “CRLite: A Scalable System for Pushing All TLS Revocations to All Browsers,” *IEEE Security & Privacy 2017*. This paper was awarded the IEEE Cybersecurity Award for Innovation.
- Taejoong Chung, Roland van Rijswijk-Deij, David Choffnes, Alan Mislove, Christo Wilson, Dave Levin, and Bruce M. Maggs, “Understanding the Role of Registrars in DNSSEC Deployment,” *ACM IMC 2017*.
- Taejoong Chung, Yabing Liu, David Choffnes, Dave Levin, Bruce M. Maggs, Alan Mislove, and Christo Wilson, “Measuring and Applying Invalid SSL Certificates: The Silent Majority,” *ACM IMC 2016*.
- Frank Cangialosi, Taejoong Chung, David Choffnes, Dave Levin, Bruce M. Maggs, Alan Mislove, and Christo Wilson, “Measurement and Analysis of Private Key Sharing in the HTTPS Ecosystem,” *ACM Conference on Computer and Communications Security 2016*.

Project: User-Centered Design for Security

- Jennifer Golbeck, “User Concerns with Personal Routers Used as Public Wi-fi Hotspots,” *IEEE UMECON*, 2017.
- Jennifer Golbeck, “The Importance of Consent in User Comfort with Personalization,” *International Conference on Social Informatics*, 2017.
- Jennifer Golbeck. “I’ll be Watching You: Policing the Line between Personalization and Privacy,” *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization*. ACM, 2017.
- Jennifer Golbeck. “User Privacy Concerns with Common Data Used in Recommender Systems,” *International Conference on Social Informatics*, 2016.

Project: Reasoning about Protocols with Human Participants

- Dawid Gawel, Maciej Kosarzecki, Poorvi Vora, Hua Wu, and Filip Zagorski, “Apollo - End-to-End Verifiable Internet Voting with Recovery from Vote Manipulation,” E-Vote-ID 2016.
- Josh Benaloh, Matthew Bernhard, Alex Halderman, Ronald Rivest, Peter Ryan, Philip Stark, Vanessa Teague, Poorvi Vora, and Dan Wallach, “Public evidence from secret ballots,” CoRR, abs/1707.08619, 2017. A shorter version of this has been accepted at E-Vote-ID 2017.

Project: Empirical Models for Vulnerability Exploits

- Soyumya Indela, Mukul Kulkarni, Kartik Nayak, and Tudor Dumitras, “Toward Semantic Cryptography APIs,” *IEEE Cybersecurity Development Conference 2016*.

Project: Human Behavior and Cyber Vulnerabilities

- Ziyun Zhu and Tudor Dumitras, “FeatureSmith: Automatically Engineering Features for Malware Detection by Mining the Security Literature,” *ACM Conference on Computer and Communications Security 2016*.
- Arunesh Mathur and Marshini Chetty, “Impact of User Characteristics on Attitudes Towards Automatic Mobile Application Updates, “ *SOUPS 2017*.

Project: Does the Presence of Honest Users Affect Intruders’ Behavior?

- Alexander Testa, David Maimon, Bertrand Sobesto, and Michel Cukier, “Illegal Roaming and File Manipulation on Target Computers,” *Criminology & Public Policy 16:* 689–726, 2017.

Project: Understanding How Users Process Security Advice

- Elissa Redmiles, Sean Kross, and Michelle Mazurek, “Where is the Digital Divide? Examining the Impact of Socioeconomics on Self-reported Security and Privacy Experiences, “ *ACM Conference on Human Factors in Computing Systems 2017*.

Project: Trustworthy and Composable Software Systems with Contracts

- David Darais, Nicholas Labich, Phúc C. Nguyên, and David Van Horn, “Abstracting Definitional Interpreters,” in *Proceedings of the ACM on Programming Languages*, 1(12), 2017.

EDUCATIONAL

Project: Measuring and Improving the Management of Today’s PKI

Dave Levin has incorporated the results from his research into both graduate and undergraduate courses on Computer and Network Security.

Project: Empirical Models for Vulnerability Exploits

In Fall 2017 Tudor Dumitras taught ENEE 657, a graduate computer security course that emphasizes empirical methods in security.

Project: Understanding How Users Process Security Advice

A class project in Michelle Mazurek's graduate class on human factors in security and privacy led to the work on edutainment.

Project: Trustworthy and Composable Software Systems with Contracts

PIs David Van Horn and Sam Tobin-Hochstadt lectured at the long-running Oregon Programming Languages Summer School.David Van Horn is currently teaching an experimental variant of the introductory programming sequence featuring prominent use of a design-by-contract programming methodology and automated verification and bug-finding tools recently developed as part of this effort.

General: Advanced Cybersecurity Program Experience for Students (ACES)

UMD Co-PI Michel Cukier also serves as Director of the UMD Advanced Cybersecurity Program Experience for Students (ACES). NSA provided mentors for the Fall 2016 and Spring 2017 ACES program. Each mentor has two students they work with, and they meet in person twice a semester. ACES leadership provides various topics of discussion so they can stay engaged.

COMMUNITY OUTREACH

Project: Measuring and Improving the Management of Today's PKI

Dave Levin presented research results to international collaborators at the University of Jordan, Princess Sumaya University for Technology, and the Hashemite University. Dave Levin and his colleagues have also been working with developers at Mozilla to explore incorporating the CRLite system into Firefox. If successful, this would result in all Firefox clients being fully up-to-date on all certificate revocations on a daily basis, a drastic improvement over today's status quo.

Project: User-Centered Design for Security

Jennifer Golbeck: "Foretold Futures from Digital Footprints: Artificial Intelligence, Behavior Prediction, and Privacy", Washington & Lee University Mudd Center for Ethics, February 2017 (keynote); University of Pittsburgh Big Data Science Colloquium, March 2017.
Jennifer Golbeck: ''Algorithmic Servants or Algorithmic Tyranny: Living with a Predicted Future" (keynote), University of Tennessee Social Media Week, February 2017.

Project: Reasoning about Protocols with Human Participants

Poorvi Vora served as a technical expert providing affidavits in support of Jill Stein's petition for a manual recount in the 2016 election in the states of Wisconsin and Michigan. She also served as an expert providing testimony to the Maryland Board of Elections on their proposed audits. She wrote an op-ed article in the Baltimore Sun, with Philip Stark, on why Maryland needs to manually examine paper ballots. In March 2017, she was awarded the Public Engagement Award for her work by the Election Verification Network, the premier network of election integrity experts.
In early 2017, Poorvi Vora wrote multiple articles and open letters to the Election Commission of India on the security of India's Electronic Voting Machines, which are under considerable public scrutiny after the most recent state elections in India.Doctoral student Hua Wu presented at E-Vote-ID 2016 and the DC-Area Privacy and Security seminar.

Project: Empirical Models for Vulnerability Exploits

In January 2017, Tudor Dumitras organized a second invitation-only workshop aimed at researchers interested in studying security empirically, using data-driven techniques. The 29 workshop participants came from 6 countries and represented organizations from academia, industry, and government. The discussion topics included understanding the motivations, capabilities, and limitations of real-world adversaries; putting theoretical assumptions to the test; accounting for the socio-economic incentives of attackers and for the properties of deployment environments; measuring and predicting security; secure data mining and machine learning techniques; automatically learning the semantics of security threats; clean-slate ideas, grounded in security measurements.

Project: Human Behavior and Cyber Vulnerabilities

Arunesh Mathur presented the work on secure software updates to girls ages 10-14 at the AspireIT camp at Princeton High School in Summer 2017. To encourage further research on natural-language processing for security, the data behind FeatureSmith has been released at http://featuresmith.org/ So far, six academic institutions have requested access.

Project: Understanding How Users Process Security Advice

Elissa Redmiles presented publications at the Conference on Communication and Computer Security (CCS 2016), CHI Conference on Human Factors in Computing Systems 2017, and at the Symposium on Usable Privacy and Security (SOUPS) 2017 WAY workshop. She presented a poster about the Edutainment project at Network and Distributed System Security Symposium (NDSS) 2017. She also presented this work at the Carnegie Mellon University CyLab Usable Privacy and Security seminar.

Elissa Redmiles wrote an article ("Why Installing Software Updates Makes Us WannaCry") for the academic news website The Conversation; the article was picked up by the Associated Press and Scientific American, among other publications.

Michelle Mazurek gave a presentation sponsored by Computing Research Association's Committee on the Status of Women in Computing Research (CRA-W) at Capital-Area Celebration of Women in Computing (CAPWIC) in February 2017. She also gave an invited talk on the topic of security behavior and advice at the 2017 Workshop on Technology and Consumer Protection (ConPro 2017).

General: Summer Camps

The Maryland Cybersecurity Center (MC2), which is home to the SoS UMD Lablet, and the National CyberWatch Center K-12 Division hosted three summer camps for students interested in the field of cybersecurity: Cyber Defense Training Camp, Intermediate CyberSTEM camp, and CyberSTEM camp. Each camp included hands-on activities designed to encourage students to gain confidence in their cybersecurity knowledge, with an emphasis on skill-building and interactive learning. NSA provided a guest speaker to each of the three summer camps.

Conference Panel and Committee Members

- Jonathan Katz served as program co-chair for the International Cryptology Conference (Crypto 2017).
- Adam Aviv served as Workshop and Tutorial Chair for SOUPS 2017.
- Adam Aviv and Michelle Mazurek served on the New Security Paradigms Workshop (NSPW) 2017 program committee.

Science of Security Quarterly Meetings

Winter 2017 Quarterly: North Carolina State University

The Winter 2017 Quarterly Science of Security (SoS) Lablet meeting was held at North Carolina State University (NCSU) on February 1 and 2, 2017, and hosted by Laurie Williams and Munindar Singh, NCSU Lablet Principal Investigators (PIs). The meeting focused on privacy and security with invited talks, Lablet presentations and the poster session all addressing the overarching theme.

Adam Tagert, NSA's Science of Security and Privacy Technical Director, provided a program update. He addressed selected activities at each Lablet including research, outreach, and education efforts. In addition to noting the NCSU Community Day held in October 2016, he highlighted the following NCSU research: a game theoretic model for Intrusion Detection Systems (IDS); how users do better against phishing with better tools; optimized reconnaissance of Software Defined Networks; and a new view of resilient architectures. Dr. Tagert covered the University of Illinois Urbana-Champaign (UIUC) summer intern program and a Bitcoin networking class as well as their research into modeling formalism for defenders, attackers, and users and open-sourcing DSSNet software. In addressing University of Maryland (UMD) Lablet activities, he noted that the Lablet held a workshop on Big Data and Machine Learning for cybersecurity and is working with the Office of Science and Technology Policy (OSTP) on how to get users to use 2-factor authentication. He covered two significant efforts at Carnegie Mellon University (CMU): collaboration with Sub-Lablet University of Texas, San Antonio (UTSA) on Polidroid, a website to help detect and repair potential privacy violations in mobile app source code; and expansion of Lablet work on UberSpark, the system they developed that enforces verifiably secure-object abstractions in systems software. He noted that in addition to the selected ongoing research, outreach, and educational activities covered above, the Lablets generated 114 publications since the last winter quarterly.

Tomas Vagoun of the National Coordination Office, Networking and Information Technology Research and Development program, addressed "Federal Privacy R&D Priorities". The modern definition of privacy expands the old definition from a right to be left alone to new concerns about large scale data collection, analysis and algorithmic decision-making. Privacy concerns are the effects of authorized Personally Identifiable Information (PII) processing. Dr. Vagoun discussed security and privacy concerns associated with PII, the National Privacy Research Strategy, the current federal priorities for privacy research, and privacy research efforts underway at a number of federal agencies. A National Science Foundation (NSF) survey identified developing approaches for remediation and recovery and reducing privacy risks of analytical algorithms as areas of major research gaps.

NSA's Dave Marcos presented "Researching the Science of Privacy". In his view, the Science of Privacy is a principled and methodological approach to privacy risk addressing the following research challenge questions: Can it be considered? Can a mathematical method be developed to evaluate privacy risk? How can a privacy accountability framework be built for Big Data? Can we apply current advances in engineering such as digital rights management, differential privacy, homomorphic encryption, and secure multi-party computation? How can the effectiveness of current privacy frameworks and associated controls be evaluated?

Jennifer Cowley of CMU's Software Engineering Institute/CERT division gave a presentation entitled "Why Can't I Put Down My Phone? The Paradox of Computing in Modern Work Environments" which discussed the unintended effects of the digital age, including the ability to think deeply. She challenged the community to think about whether implementing a new technology makes sense and whether it is good for people.

David Hoffman, Director of Security Policy and Global Privacy Officer at Intel, spoke on the relationship between privacy and security. In his presentation "It Takes Data to Protect Data," he said that security and privacy are neither tradeoffs nor a zero-sum game. Rather, the two should be thought of as needing to be in balance. It should be a process of adding to the other when one is increased. Risks are radically changing: new technologies have been created that allow a small group to inflict extreme harm on a large number of people using drones, germs, robots, and hackers, and the threat has become asymmetric. "Good cybersecurity is good for privacy," he concluded.

Bill Scherlis, Lablet PI at CMU, described a conference CMU hosted in the summer of 2016 on "Safety and Control for AI-based Systems". Artificial Intelligence (AI) is now embedded in critical infrastructure and has a big impact on security. We need assurance judgments about AI systems and for them to become reliable and trustworthy. AI safety is multidimensional and must be addressed in the mission context.

Lablet technical research presentations:

1. "Anonymity in the Bitcoin P2P Network"
Guilia Fanti (UIUC), presenter, with Shaileshh Venkatakrishnan (UIUC) and Pramod Viswanath (UIUC)
Research Project: Anonymous Messaging

This presentation addressed the reality that Bitcoin, like other cryptocurrencies, publish users' entire transaction histories in plaintext using a pseudonym, despite the public perception of it being a "privacy-preserving" financial system. Consequently, if a user's pseudonym can be linked to their human identity, the privacy fallout could be significant. The goal of the research being presented is to provide strong, provable anonymity guarantees, and the researchers propose a simple networking policy called Dandelion which achieves nearly-optimal anonymity guarantees at minimal cost to the network's utility.

2. "A Control-theoretic View of AI for Security"
Dave Roberts (NCSU)
Research Project: Leveraging the Effects of Cognitive Function on Input Device Analytics to Improve Security

The presenter noted that research in psychology, Human-Computer Interaction (HCI), and related fields tells us that human behavior can be predicted, influenced, and measured, and that those characteristics provide a security advantage given the right data, analytics, and models. He noted that one fundamental task for AI in security is to develop control-theoretic methods that enable systems to use analytics to reason about how users complete tasks and identify evidence of departures from normal behavior using control feedback to influence task conditions.

3. "Privacy Incidents, Privacy News and News about Incidents"
Jessica Staddon (NCSU)
Research Project: Privacy Incidents Database

The presenter addressed the vision and status of the privacy database, work in vetting the privacy database definitions and data and efforts to partially automate database maintenance. She also discussed privacy news, including trends in sentiment, entities, and keywords. The vision for the privacy database is that it becomes a Wikipedia for privacy incidents, and that it can provide visualization of incident trends and downloadable data sets.

4. "Discovering a Natural Language Semantics for Privacy"
Travis Breaux (CMU)
Research Project: Usable Formal Methods for the Design and Composition of Security Privacy Policies

This research focuses on developing semantic analysis techniques and tools to infer ontological relationships in natural-language privacy policies from information type variants and hypernymy patterns. The researchers are designing an empirical method to construct privacy ontologies from policies.

5. "Security and Privacy in Machine Learning"
Nitin Vaidya (UIUC)
Research Project: Anonymous Messaging

The speaker covered the motivation for this research as distributed machine learning, and the research problems as:
1) Privacy-preserving distributed optimization; 2) adversarial learning; and 3) robustness to adversarial samples. After covering each of these aspects in depth, the presenter

concluded that achieving privacy/security in learning is non-trivial, and while there has been some promising progress, there is much left to do.

6. "How Good is a Security Policy against Breaches?"
Özgür Kafali (NCSU), presenter, with Jasmine Jones, Megan Petruso, Laurie Williams, and Munindar Singh
Research Project: Formal Analysis of Breach reports

The goal of this research is to help analysts measure the gaps between security policies and reported breaches by developing a systematic process based on semantic reasoning. We propose SEMAVER, a framework for determining coverage of breaches by policies via comparison of individual policy clauses and breach descriptions. We represent a security policy as a set of norms. Norms (commitments, authorizations, and prohibitions) describe expected behaviors of users, and formalize who is accountable to whom and for what. A breach corresponds to a norm violation. We develop a semantic similarity metric for pairwise comparison between the norm that represents a policy clause and the norm that has been violated by a reported breach. We used the US Health Insurance Portability and Accountability Act (HIPAA) as a case study, and found that HIPAA's gaps regarding accidental misuses are significantly larger than its gaps regarding malicious misuses.

Posters

The poster sessions were held on each day of the meeting, and provided an opportunity for Quarterly attendees to discuss ongoing Lablet research. Some of the posters described recent publications while others addressed new findings and directions associated with Lablet fundamental research. Students who presented the posters engaged in spirited discussions with participants on opportunities to expand the research presented.

The following posters were presented. A link to the posters is available at https://cps-vo.org/node/27211

- Privacy-Aware Socially Intelligent Applications: Nirav Ajmeri, Munindar Singh
- Adoption of Security Analysis Tools in Software Development: Shams Al-Amin, Nirav Ajmeri, Munindar Singh, Jon Doyle, Emily Berglund
- UI Manipulation for Persistent, Subtle Security Proofs: Ignacio Dominguez, Nischal Shrestha, David Roberts, Robert St. Amant
- Identifying and Consolidating Security Requirements from Regulatory Documents: Sarah Elder, Hui Guo, Laurie Williams, Munindar Singh
- Resource-efficient Composition of Network Security Applications: Victor Heorhiadi, Michael Reiter, Vyas Sekar
- On the tradeoff between Utility and Privacy in Collaboration Intrusion Detection Networks: Richeng Jin, Huaiyu Dai
- How Good is a Security Policy against Breaches?: Özgür Kafali, Jasmine Jones, Megan Petruso, Laurie Williams, Munindar Singh
- Cloud Storage Service Organization Study: Yehuda Katz, Casen Hunger, Michelle Mazurek, Charalampos Papamanthou, Mohit Tiwari
- Flow Reconnaissance via Timing Attacks on SDN Switches: Sheng Liu, Michael Reiter
- Practical Secrecy Enforcement on Modern Commodity Platforms: Adwait Nadkarni, Benjamin Andow, William Enck
- Defect Analysis of Infrastructure as Code: Akond Rahman, Laurie Williams
- Leveraging Search Query Trends as a Privacy Preference Proxy: Shaown Sarker, Andrew McNamara, Jessica Staddon
- A Study of Security Vulnerabilities on Docker Hub: Rui Shu, Xiaohui (Helen) Gu, William Enck
- Attack Surface Definitions: Christopher Theisen, Nuthan Munaiah, Laurie Williams
- FeatureSmith: Automatically Engineering Features for Malware Detection by Mining the Security Literature: Ziyun Zhu, Tudor Dumitras
- An Analysis of Phishing Bait: Olga Zielinska, Patrick Lawson, Christopher Mayhorn

Spring 2017 Quarterly:

In lieu of a Spring 2017 Quarterly, the SoS Community met in April at HotSoS 2017 hosted by the University of Maryland (UMD) Lablet and the SURE project in Hanover, MD. See page 67 for full HotSoS 2017 details.

Summer 2017 Quarterly:
Carnegie Mellon University

The Summer 2017 Quarterly Science of Security (SoS) Lablet meeting was held at Carnegie Mellon University (CMU) on July 10 and 11 and hosted by Bill Scherlis, (PI) at the CMU Lablet. The Quarterly included a panel discussion entitled "Retrospective of the Science of Security Program: An NSA Perspective," as well as summaries by each of the Lablets on their current research, and the presentation of five technical papers.

This was the final quarterly meeting for the four Lablets that received awards under the Science of Security Broad Area Announcement (BAA) released in 2013: CMU, North Carolina State University (NCSU), University of Illinois at Urbana-Champaign (UIUC), and University of Maryland

(UMD). Under the most recent SoS initiative, 330 papers have been published and researchers have developed 59 tools focused on multiple aspects of security analysis and design. The Lablets have been described as small transdisciplinary labs for research to build a science of cybersecurity, and about one-fifth of the research has been done in cooperation with other institutions, including the 25 Sub-Lablets and 138 collaborators. While SoS Lablet research is built around the five Hard Problems (developed with the Lablets), their research also applies to other topics that fall under Information Assurance focus areas.

NSA's Science of Security and Privacy Technical Director, Adam Tagert, provided a program overview entitled "Laying the Cyber Foundation: The Science of Security and Privacy Initiative". He addressed the role of science in Science of Security and discussed the program's three pillars, presented the origins of the SoS program, and outlined each Lablets' current research projects. He also covered the SoS Lablet achievements described above and non-Lablet activities. In addition to the Lablets, the SoS initiative includes Vanderbilt's Science of SecUre and REsilient Cyber-Physical Systems (SURE) project, the Best Cybersecurity Paper competition and Intel's International Science and Engineering Fair (ISEF), where high school students are recognized for accomplishment in scientific cybersecurity research. He noted that there are close to 200 research organizations participating in the SoS initiative (Lablets, Sub-Lablets, SURE, collaborators, and Paper Competition submissions) in 30 countries on 5 continents.

The panel discussion featured five researchers who have participated in and monitored the program since its inception. Questions discussed included how each panelist defines Science of Security, whether there is or can be a standard definition, how and whether the five Hard Problems can be solved, and where the field goes from here. A robust interactive discussion both among the panelists and with the audience ensued. While there appeared to be consensus there is no standard definition of Science of Security, there was also a strong theme that there is significant value in adding rigorous methods and processes and probing for fundamental or basic principles. Some parts are mathematical and not empirical; the other side is the practical side. Most panelists described a progression in their thinking during the initiative, having become less wedded to the distinction between engineering and good empirical work from the mathematical. The initial concept focused on pure science, systemization of knowledge and clarification of ideas that can be a basis for advancement. The actual outcomes may not have provided complete "solutions" to the five Hard Problems, but there has been significant improvement both in terms of an ultimate solution and as benchmarks for the development of a Science of Security.

Summaries by the four PIs showed that Lablet research has been substantial. Michel Cukier (UMD) displayed a matrix of research by Hard Problem and provided updates on the ten projects UMD has under way. Eight of the ten projects focused on the Hard Problem of Human Behavior (either exclusively or along with another Hard Problem) while the next most prevalent Hard Problem addressed was Security Metrics and Models.

Bill Sanders and Sayan Mitra (UIUC) summarized the research into the six projects currently sponsored by SoS, and noted that they are adding another project. They reported that there have been 89 papers published under the SoS Lablet initiative and numerous outreach efforts, including their most recent summer intern program that concluded with a poster session.

Laurie Williams (NCSU) summarized the research on projects that address the Hard Problems, with five projects focused on Policy-Governed Secure Collaboration, three on Security Metrics and Modeling, five on Resilient Architectures, and three on Human Behavior. She also spoke about efforts in community building to include NCSU's annual industry day and annual summer workshop, their focus on co-authorship, and their project to analyze published papers.

Bill Scherlis (CMU) described his Lablet's goals as advancing scientific coherence through methods, validation, and productivity, and broadening the cybersecurity technical community via educational engagement and conferences such as HotSoS. Before summarizing the progress on CMU's nine projects, he presented CMU's three-part approach to SoS: 1) Address the most challenging Hard Problems areas, therefore focusing primarily on Scalability and Composability and Human Behavior; 2) Advance the process and methods by which science is done, including a specific focus on methods, synergies in the Lablet approach, and towards a common base for analysis and engineering; and 3) Engage with the broader research and technical community to address these goals.

Five technical research presentations were included in the meeting.

1. "Optimal State Estimation and Model Detection and Applications to Security and Privacy"
Sayan Mitra (UIUC), presenter, with Geir Dullerud (UIUC) and Swarat Chaudhuri (Rice)
Research Project: Static-Dynamic Analysis of Security Metrics for Cyber-Physical System

This project's objectives are to develop rigorous, model-based approaches for analyzing security metrics of large cyber-physical systems such as power systems, traffic control systems, and autonomous vehicles. In order to make the approaches scale to large models, the researchers are developing foundational results on compositional analysis. They also seek to formalize and characterize trade-offs between security/privacy on the one hand and performance and accuracy on the other. They investigated this problem from a perspective of topological entropy, introduced the notion of estimation entropy hRS, and showed it is impossible to monitor at bitrates below hRS. The upper bound on estimation entropy is hRS, ≤ L + a n. Detection algorithms with optimal bit-rate operate up to the entropy upper bound.

2. "Prioritizing Security Efforts with a Risk-Based Attack Surface Approximation (RASA)"
Chris Theisen (NCSU)
Research Project: Attack Surface and Defense-in-Depth Metrics

One prioritization technique is to identify the attack surface. Crashes are empirical evidence of data paths through software with flaws. Code that is covered by RASA are therefore more likely to have vulnerabilities, as there is evidence of flaws on RASA, and are more likely to be exploited, as they're on known traversable paths. Vulnerabilities are five times as likely to be in code that crashes than not. We are recovering the majority of vulnerabilities (94%). Future work will include comparing four vulnerability prediction models with RASA.

3. "Immutability for Integrity: Combining Language Theory and the Science of Usability in Glacier"
Jonathan Aldrich (CMU), presenter, with Michael Coblenz, Whitney Nelson, Brad Myers and Joshua Sunshine.
Research Project: Race Vulnerability Study and Hybrid Race Detection

Assessment of research systems, particularly Java, showed root problems: "we tend to build things that are too complex". The team built GLACIER: "Great Languages Allow Class Immutability Enforced Readily" to create simple, strong transitive immutability as an annotation system and checker for Java. They concluded, after testing, that GLACIER illustrates an effective approach to improving guidelines by using mathematical models to ensure correctness and power of tools and leveraging usability science to ensure benefit from that power in practice.

4. "Observing Passwords in Their Natural Habitat"
Lorrie Cranor (CMU), presenter and Sarah Pearman, Jeremy Thomas, Pardis Emami Naeini, Hana Habib, Lujo Bauer, Nicolas Christin, Serge Egelman, and Alain Forget
Research Project: Understanding User Behavior when Security is a Secondary Task

This study focuses on the problems of password reuse. The researchers ask how people manage all their passwords and collect empirical data to scientifically examine this problem. By using the Security Behavior Observatory to collect password data from home computer users, they were able to scientifically characterize user password behavior, moving beyond previous analyses that have been mostly anecdotal and based on speculation. After collecting data, they concluded that reuse is rampant, that users seem to cope with password demands through reuse strategies and use a mixture of reuse strategies; password managers may not be helping very much. Most of the issues are behavioral.

5. "Vulnerability Disclosure in the Age of Social Media: Exploiting Twitter for Predicting Real-World Exploits"
Octavian Suciu (UMD), presenter, with Carl Sabottke and Tudor Dumitraș
Research Project: Empirical Models for Vulnerability Exploits

This study seeks to predict exploits active in the wild. The growing number of vulnerabilities reported each year prompted the researchers to address whether Twitter analytics can be used for early detection. Conclusions and results included the design of a Twitter-based exploit detector that can be used for patching prioritization and risk assessment. Early detection of exploits active in the wild is possible, but performance depends on the quality of ground truth. Exploit detection under adversarial interference yields a security system without secrets.

SoS
Our Origin Story
CNCI Funding Begins Initiative

Identifying and Consolidating Security Requirements from Regulatory Documents
Different Approaches

Motivation: A Security Indicator for Cars

SYSTEM SCIENCE OF SECURITY AND RESILIENCE OF CPS
SURE
Key Ideas
Impact

Science of SecUre and REsilient Cyber-Physical Systems (SURE)

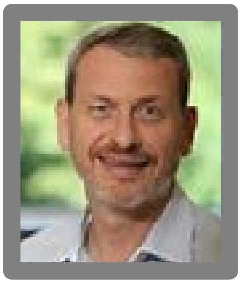

Xenofon Koutsoukos

The SURE project, under the direction of Principal Investigator (PI) Xenofon Koutsoukos of Vanderbilt University, draws together researchers from Vanderbilt, University of Hawaii, University of California, Berkeley, and the Massachusetts Institute of Technology to develop foundations and tools for designing, building, and assuring that Cyber-Physical Systems (CPS) can maintain essential system properties in the presence of adversaries. NSA's Information Assurance Research sponsorship of the SURE project began in 2014 and is aimed at improving scientific understanding of resiliency in CPS. The research addresses how to design systems that are resilient despite decentralization of resources and decision-making. The project has supported 8 graduate and 5 post-doctoral students and resulted in 53 publications in 40 different conferences and journals. In addition to advancing the science of CPS security, the project includes an effort to create a strong cadre of engineers dedicated to advancing security and resilience in CPS.

SURE RESEARCH THRUSTS

Hierarchical coordination and control

- Cyber risk analysis and incentive design: Goal is to develop regulations and strategies at the management level.
- Goal is resilient monitoring and control of the networked control system infrastructure.

Science of decentralized security

- Goal is to develop a framework that will enable reasoning about the security of all the integrated constituent CPS components.

Reliable and practical reasoning about secure computation and communication in networks

- Goal is to contribute a formal framework for reasoning about security in CPS.

Evaluation and experimentation

- Goal is to use modeling and simulation for the integration of cyber and physical platforms that directly interface with human decision-making.

Education and Outreach

- Goal is to educate the next generation of researchers in the field of security and resilience of CPS.

ANNUAL MEETING

In August 2017, SURE researchers meeting at Vanderbilt University briefed members of NSA's Information Assurance Research organization and the Air Force Research Lab on the outcomes of the three-year program.

Xenofon Koutsoukos, SURE PI, described the objectives of SURE as "developing a systematic body of knowledge with strong theoretical and empirical underpinnings to inform the engineering of secure and resilient CPS that can resist not only known, but also unanticipated attacks." He reported that among the most significant lessons-learned from the project are the following: attack modeling is hard; resilience depends on application and context; and controlled experiments present a design challenge of creating experiments with verisimilitude to the real world.

An important element of SURE's work has been the development and use of a customizable testbed (now a cloud-based platform) that allows simulations of such systems as transportation, water distribution and communications networks to develop resilience metrics and to apply specific methodologies to multiple systems. To date, researchers have modeled the smart grid, autonomous and human-guided automobiles, railway transport, and water distribution networks. They have also modeled cyber systems topology, wired/wireless distribution, protocols, and monitoring, and control systems. The lessons-learned changed how to perform research in CPS resilience and demonstrated that the testbed is scalable through its distributed heterogeneous simulation of CPS. While there are some attacks that cannot be replicated in the test bed, SURE researchers have been successful on focusing on hardware-in-the-loop, modeling side channel attacks, and algorithms for complexity attacks and their consequences. Modeling has shown SURE researchers how different systems can interact. The main research thrusts addressed via the testbed are cyber risk analysis and incentive design, resilient monitoring and control, and the science of decentralized security.

The live demonstration by Peter Volgyesi and Himanshu Neema showed actual testbed operations. The goal for the testbed is to produce repeatable experiments that are rapid and collaborative and have a CPS domain focus. New scenarios developed within the past year include automated simulation-based risk analysis and active resilience based on multistage games.

Following the testbed demonstration, Vanderbilt faculty and graduate students offered the following presentations:

Decentralized Control and Path Planning Games

Yevgeniy Vorobeychik

Decentralization is a fundamental problem in autonomous control. It is an important subject to study, as the world is moving towards greater decentralization, autonomy, and deregulation in CPS control decisions. In studying the problem, Nash equilibrium outcomes are typically similar to socially optimal in complex traffic light control; however, the gap can be large in special cases. Nash equilibriums are provably socially optimal in personalized configurations of Intrusion Detection Systems (IDS); however, they can cause both over- and under-investment in security settings. Over-investment may result in an “arms race”, and under-investment may occur when significant interdependencies exist. The outcomes can make the system highly susceptible to vindictive players (or cyber-attacks which attack one player to target another one), and can compromise safety in autonomous path planning.

Automated Simulation-Based Vulnerability Analysis Using Courses-Of-Action (COAS)

Himanshu Neema and Peter Volgyesi

Previously done on an ad hoc basis, experimental design is now focused on the development of a library of Courses of Action (COAs), a combinatorial approach with various subsets of COAs, and hierarchical grouping of them. COAs are sequence models created using actions that inject new information into a situation at given time points and produce outcomes matching a given pattern. They are Directed Acyclic Graph (DAG) by connecting atomic nodes with directed edges. The challenge is to create novel situations. The SURE testbed reuses the simulation integration toolkit based on two components: the SURE Testbed for CPS Security and resilience evaluation (CPSWT) and the Domain Independent Simulation Integration Toolkit (CPSWT). With these tools, links, routers, networks, and sensors, can be created to produce a communication network simulation. SURE researchers have been able to build a reusable cyberattack library. COAs are an intuitive means to enable system analysts to perform many scenario-driven experiments using the same integrated simulation.

Active Resilience Based on Multi-Stage Games

Aron Laszka

The challenge for Active Resilience Formulation is how to find an optimal defense strategy. This study approached the problem using a multi-stage game theoretic model that accounts for a strategic adversary and considers multiple defensive and adversarial actions using a case study in traffic simulation. Results show that in transportation networks, active resilience can significantly reduce the impact of ongoing cyber-attacks; in practice, defenders should be prepared to mitigate cyber-attacks by detecting and responding to them. Cyber-attacks in response to mitigation have low impact (active resilience can limit the impact of cyber-attacks even if adversaries remain active) and cyber-attacks have significant impact before mitigation, so it is difficult to anticipate and limit the impact of all possible attacks.

SURE Testbed Infrastructure and Technology Transition

Himanshu Neema

Over the last several years Vanderbilt has developed many open-source platform tools that can be “rewired” for reuse in many different application domains; the SURE Testbed is good example of one that was developed based on these core component tools. The core components and platforms have been successfully transitioned to many different applications. The technology transition also resulted in an open-source release of packaged tools on GitHub in the form of UCEF, which is already gathering a large community of users.

A Hardware-in-the Loop Testbed for Evaluating and Measuring Security and Resilience in CPS

Bradley Potteiger

SURE has developed two major attack scenarios on the test bed: a Moving Target Defense (MTD) and an autonomous car code injection attack. Simulation models may be based on simplifying assumptions or are incomplete. Comprehensive testing necessitates replicating the CPS using a test platform providing hardware and embedded control code similar to the real system. System dependent attacks (DDOS) and low-level code abstraction are hard to replicate in simulation environments. For the MTD experiment, CPS are vulnerable to code injection attacks through buffer overflow vulnerabilities; Instruction Set Randomization (ISR) mitigates this attack by altering machine code architecture. In this scenario, a malicious payload is injected into state controller through sensor spoofing: the payload opens a remote root shell to the attacker, allowing them to run a malicious fixed time state controller with a significant green time for the non-critical route. In the autonomous car attack scenario, a code injection attack using SOCKETCAN was used on an Udacity autonomous car simulator to manipulate various sensors including the GPS, camera, direction sensor, and state controller with the intention of causing a virtual car crash with "fatalities." The work produced strong results for both scenarios.

Foundations of CPS Resilience

Xenofon Koutsoukos with Aron Laszka, Waseem Abbas, and Yevgeniy Vorobeychik

Professor Koutsoukos described two new demonstrations developed within the past year that focused on more complex and nuanced automated simulation risk analysis and an active resilience based on multi-stage games. Three broad issues he presented were: 1) combining hardening and diversity to improve structural robustness of CPS networks; 2) integrating redundancy, diversity, and hardening for detection of cyber-physical attacks in water distribution systems; and 3) integrating diversity and hardening for resilient traffic control systems. He noted that one goal is to develop a model that allows principled investment in redundancy, diversity, and hardening for improving resilience in CPS. How to allocate these budgets for optimal deployment can be tested via algorithm. While the problem is computationally challenging, an efficient heuristic that works well in practice can be devised. This methodology and the metric algorithm presented have been successfully applied to water distribution, transportation systems, and power networks.

Software Vulnerability Analysis for Cybersecurity

Daniel Balasubramanian with Gabor Karsai (Vanderbilt), Corina Pasareanu (Carnegie Mellon University), and Tevfik Bultan (University of California, Santa Barbara)

According to the presenter "Once new defense technologies make vulnerabilities based on flawed implementations less common, adversaries will turn their attention to vulnerabilities inherent in the algorithms themselves, specifically, vulnerabilities stemming from space and time resource usage." This enables algorithmic complexity attacks and side-channel attacks. Algorithmic complexity attacks are denial-of-service attacks, where the adversary's goal is to deny service to the system's users, or to disable the system by choosing a worst-case input. Side-channel attacks recover secret inputs to programs from non-functional characteristics of computations, such as time consumed, number of memory accesses or size of output files. These vulnerabilities aren't "bugs" in the traditional sense. The goal of this research is to write algorithms/tools to detect whether a program contains one or the other. It is a hard problem, but building useful practical tools is possible.

CPS/IIoT Security

Gabor Karsai

The Cyber Physical Systems/Industrial Internet of Things presents at least three sets of challenges to researchers. The first set includes: identifying and specifying requirements in their physical context; determining the "security requirements language" and "security policy language" for CPS; defining and enforcing policies in their physical context; features in model generated code and systems; the "security architecture design" for CPS and its implications on non-functional properties; and what security requirements can be addressed (guaranteed) by the architecture of the system and by the model generated. The second set of challenges involves measurement and metrics including evaluation metrics, the quantitative metrics to evaluate alternative CPS designs and implementations with regard to resilience to security issues. Features of the third set of challenges include defensive mechanisms such as detection, prevention, and mitigation; security anomaly detection features; features available to prevent or mask the effect of security issues; and mitigation features available to contain the effects of security issues. Existing development paradigms and industry trends point towards platform-based security solutions. Model-driven component-based development relies on solid and robust execution platforms that could serve as the 'trusted computing base' for CPS. Interactions with and exposure to the physical environment are critical in CPS, and security breaches can have grave consequences that need to be analyzed and understood. There is a clear need for Platform-based Science of Security for CPS.

Adversarial Machine Learning in Cyber-Physical Systems: What it is and Why it Matters for CPS

Yevgeniy Vorobeychik

This work is fundamental research into adversarial classification using Adversarial Machine Learning (AML), specifically looking at the science of adversarial evasion modeling and evasion robust classification. It addresses questions related to key scientific questions for AML in CPS including the following: how to model attacks on ML in CPS particularly when there are physical consequences and how to validate such models; what are relevant attack vectors in CPS; and how to develop methods for making ML in CPS resilient to attacks. Little attention has been paid to adversarial manipulation of regression, which is crucial in control (end-to-end deep learning-based control, for example). The next issue will be how to manipulate vision systems to induce mistakes in control decisions.

The program agenda and links to slide presentations can be found online at: https://cps-vo.org/group/sos/sure/meetings

PUBLICATIONS

- Bo Li, Yevgeniy Vorobeychik, Muqun Li, and Bradley Malin, "An Iterative Classification Scheme for Sanitizing Large-Scale Datasets," *IEEE Transactions on Knowledge and Data Engineering,* 29(3):698-711, 2017.
- Bradley Potteiger, William Emfinger, Himanshu Neema, Xenofon Koutsoukos, CheeYee Tang, and Keith Stouffer, "Evaluating the Effects of Cyber-Attacks on Cyber Physical Systems using a Hardware-in-the-Loop Simulation Testbed," *National Symposium on Resilient Critical Infrastructure, Resilience Week 2017*, Wilmington, DE, September 18-22 2017. (Best paper award).
- Aron Laszka, Yevgeniy Vorobeychik, and Xenofon Koutsoukos, "A Game-Theoretic Approach for Integrity Assurance in Resource-Bounded Systems," *International Journal of Information Security*, January 31, 2017.
- Waseem Abbas, Aron Laszka, and Xenofon Koutsoukos, "Graph-Theoretic Approach for Increasing Participation in Social Sensing," *The 2nd International Workshop on Social Sensing (SocialSens 2017)*, Pittsburgh, PA, April 21, 2017. (In conjunction with CPS Week 2017).
- Saqib Hasan, Amin Ghafouri, Abhishek Dubey, Gabor Karsai, and Xenofon Koutsoukos, "Heuristics-Based Approach for Identifying Critical N - k Contingencies in Power Systems," *National Symposium on Resilient Critical Infrastructure, Resilience Week 2017*, Wilmington, DE, September 18-22 2017.
- Waseem Abbas and Xenofon Koutsoukos, "Improving Network Connectivity Using Trusted Nodes and Edges," *The 2017 American Control Conference (ACC 2017),* Seattle, WA, May 24-26, 2017.
- Bo Li, Kevin Roundy, Chris Gates and Yevgeniy Vorobeychik, "Large-scale Identification of Malicious Singleton Files," *Conference on Data and Applications Security (CODASPY)*, 2017.
- Nika Haghtalab, Aron Laszka, Ariel Procaccia, Yevgeniy Vorobeychik, and Xenofon Koutsoukos, "Monitoring Stealthy Diffusions," *Knowledge and Information Systems (KAIS)*, 52(3), 657-685, September 2017.
- Andrew Smith, Jian Lou, and Yevgeniy Vorobeychik. "Multidefender Security Games," *IEEE Intelligent Systems*, February 13, 2017.
- Swetasudha Panda and Yevgeniy Vorobeychik, "Near-optimal Interdiction of Factored MDPs," *Conference on Uncertainty in Artificial Intelligence*, 2017.
- Amin Ghafouri, Aron Laszka, Abhishek Dubey, and Xenofon Koutsoukos, "Optimal Detection of Faulty Traffic Sensors Used in Route Planning," *2nd International Workshop on Science of Smart City/Operations and Platforms Engineering (SCOPE),* Pittsburgh, PA, April 21, 2017. (In conjunction with CPS Week 2017).
- Ayan Mukhopadhyay, Yevgeniy Vorobeychik, Gautam Biswas, and Abhishek Dubey, "Prioritized Allocation of Emergency Responders Based on a Continuous-time Incident Prediction Model*," International Conference on Autonomous Agents and Multiagent Systems, (AAMAS 2017),* 2017.
- Heath LeBlanc and Xenofon Koutsoukos, "Resilient First-Order Consensus and Weakly Stable, Higher Order Synchronization of Continuous-Time Networked Multi-Agent Systems," *IEEE Transactions on Control of Network Systems*, Published online: 4 April 2017.
- Waseem Abbas, Lina Perelman, Saurabh Amin, and Xenofon Koutsoukos, "Resilient Sensor Placement for Fault Localization in Water Distribution Networks," *ACM/IEEE 8th International Conference on Cyber-Physical Systems (ICCPS 2017)*, Pittsburgh, PA, April 18 - 21, 2017.
- Chang Liu, Bo Li, Yevgeniy Vorobeychik, and Alina Oprea, "Robust Linear Regression Against Training Data Poisoning," *Proceedings of the ACM Workshop on AI and Security,* (AISec '17), Dallas TX, 91-102, November 3, 2017. (Best paper award finalist).
- Aron Laszka, Waseem Abbas, and Xenofon Koutsoukos, "Scheduling Battery-Powered Sensor Networks for Minimizing Detection Delays," *IEEE Communications Letters*, Published online: Jan 2017.
- Waseem Abbas, Aron Laszka, Yevgeniy Vorobeychik, and Xenofon Koutsoukos, "Scheduling Resource-Bounded Monitoring Devices for Event Detection and Isolation in Networks," *IEEE Transactions on Network Science and Engineering*, Published online: 31 July 2017.
- Jiarui Gan, Bo An, Yevgeniy Vorobeychik, and Brian Gauch, "Security Games On a Plane," *Association for the Advancement of Artificial Intelligence (AAAI),* 2017.
- Xenofon Koutsoukos, Gabor Karsai, Aron Laszka, Himanshu Neema, Bradley Potteiger, Peter Volgyesi, Yevgeniy Vorobeychik, and Janos Sztipanovits, "SURE: A Modeling and Simulation Integration Platform for Evaluation of SecUre and REsilient Cyber-Physical Systems," *Proceedings of the IEEE*, Published online: 15 August 2017.
- Aron Laszka, Waseem Abbas, Yevgeniy Vorobeychik, and Xenofon Koutsoukos, "Synergistic Security for Smart Water Networks: Redundancy, Diversity, and Hardening," *3rd International Workshop on Cyber-Physical Systems for Smart Water Networks (CySWater 2017)*, Pittsburgh, PA, April 21, 2017. (In conjunction with CPS Week 2017).
- Andrew Smith, Jackson Mayo, Vivian Kammler, Robert Armstrong and Yevgeniy Vorobeychik, "Using Computational Game Theory To Guide Verification and Security in Hardware Designs," *IEEE International Symposium on Hardware Security and Trust (HOST),* 2017.

SECTION 2

Promoting Rigorous Scientific Principles

In addition to sponsoring fundamental research at the four Lablets and through the Science of SecUre and REsilient Cyber-Physical Systems (SURE) project described in Section 1 of this report, the Science of Security (SoS) initiative uses two other means to promote rigorous scientific principles: the 2017 Annual Best Scientific Cybersecurity Paper Competition and sponsorship of awards at the Intel International Science and Engineering Fair (ISEF). There were 38 nominations in the Paper Competition, and the panel of Distinguished Experts selected one winner. Representatives from the SoS initiative narrowed the 1800 Intel ISEF submissions to 62 that were relevant to SoS, held 20 interviews, and awarded a First Place, two Second Places, and three Honorable Mentions. Details on the Paper Completion and the Intel ISEF can be found in the following pages.

BEST SCIENTIFIC CYBERSECURITY PAPER COMPETITION

In order to encourage the development of the scientific foundations of cybersecurity, the National Security Agency (NSA) established the Annual Best Scientific Cybersecurity Paper Competition in 2012. NSA invites nominations of papers that show an outstanding contribution to cybersecurity science and that come from any field of cybersecurity research. In order to be eligible to be nominated for the 2017 competition, papers had to have been published in 2016 in peer-reviewed journals, magazines, or technical conferences. Nominations describe the scientific contribution of the paper and explain why this paper merits the award. Nominators may not be an author or co-author of the nominated paper. In 2017 there were 38 submissions, bringing the total number of submissions to more than 200 during the five years of the competition. 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 following individuals (* first time reviewers) served as distinguished experts for the 5th 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 Angela Sasse, University College, London
- Professor Stefan Savage, University of California, San Diego*
- Professor Paul Van Oorschot, University of Carleton*
- Mr. Phil Venables, Goldman Sachs
- Professor David Wagner, University of California, Berkeley

The winning paper of the 5th Competition was

"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, and Christian Stransky.

These researchers are at the Center for IT-Security, Privacy, Accountability (CISPA), Saarland University, in Germany and at the University of Maryland, College Park, in the United States. The paper was presented at the 2016 IEEE Symposium on Security and Privacy.

Their research helps answer the question of why software developers are writing programs that have security vulnerabilities. The authors present scientific evidence that confirms anecdotal stories in the programming community. Specifically, the researchers investigated how different information sources available to the developer influence the developer's abilities to quickly program and to program securely. They studied 54 developers (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; 2) Stack Overflow only; 3) official Android documentation only; and 4) books only. 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. However, they are rarely used; in the study, one free choice participant used a book.

This paper was selected for excelling at multiple attributes of high quality scientific work and reporting. First, the authors developed a laboratory study to control factors so they could accurately measure the information source variable and help determine the root cause of software vulnerabilities. These choices were based on their preliminary research on Android App developers where they determined the best variable to measure. The research also included work to validate the results, and the researchers examined the limitations of their study. The paper did a thorough job explaining the research method which helps other researchers duplicate and build upon this work. The paper also has some actionable scientific-based advice on developing better materials for developers to write more secure programs.

This paper adds scientific knowledge to our understanding of how developers rely on information sources and the impact of the introduction of insecure software code.

The 5th Annual Best Scientific Cybersecurity Paper Competition Awards Ceremony was hosted by the Research Directorate at NSA on October 27, 2017.

Dr. Adam Tagert, the Science of Security Technical Director, gave the welcoming remarks. Following Dr. Tagert, Dr. Deborah Frinke, NSA's Director of Research, emphasized the value of research at NSA, noting that the Director of Research has a seat at NSA's Board of Directors. She confirmed that papers such as the winning paper "influence the outside world" in part by demonstrating how science can be used as a "common language and rigor to approach problems". The competition also serves to gauge the maturity of security research, and she acknowledged seeing improvement over the years.

Authors Yasemin Acar and Christian Stransky gave a brief presentation on their research which was inspired by a common problem. 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 research explores 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. After describing their methodology of subjecting Android developers to various security-relevant tasks and varying their choices of resources (Stack Overflow, official documentation, books, and free choice), they reviewed their findings on the impacts to both functional correctness, and security correctness. They concluded that project managers should "take developers offline and give them a book," and added that while professionals tended to produce functional code more reliably, they were no better at security.

Dr. Carl Landwehr then moderated a Question and Answer panel discussion with the awardees. When asked what kinds of blind alleys they might have gone down, they noted that they didn't expect how difficult it would be to recruit enough Android developers, and to get their development system to run on different systems, with different restrictions. The discussion led to conclusions about documentation, how it needs more troubleshooting to be an effective, preferred resource, and that the team is working with Google to improve their documentation. The team was then asked about the generalizability of their findings. They responded that there needs to be a change in mindset, so that security needs to be treated as a common goal, and that "documentation matters" even though nobody likes to write it. The team noted that it might be necessary to treat developers like end-users, in that most don't know enough about security. In response to Dr. Frincke's question about other human behavior issues, the team noted that people don't search for optimal solutions and that they take security advice from odd sources which discourages deeper learning.

Dr. George Coker, Chief of Information Assurance Research, gave the closing remarks, noting that the scientific approach to Science of Security is advancing as demonstrated by the research quality improving year over year, and announcing that the 6th annual competition begins in December.

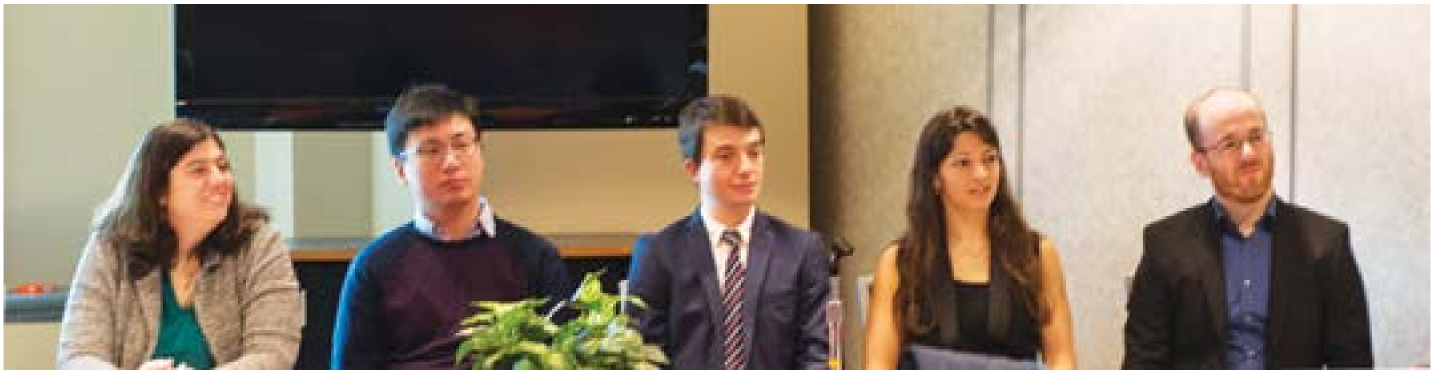

Intel International Science and Engineering Fair

For the third consecutive year, the NSA Research Directorate Science of Security initiative sponsored special awards at the Intel International Science and Engineering Fair (Intel ISEF) held May 14-19 2017. The nearly 70-year old ISEF, held annually, is the world's largest international pre-college science competition. Each year, approximately 1,800 high school students from more than 75 countries, regions, and territories are awarded the opportunity to showcase their independent research and compete for approximately $4 million in prizes.

The Science of Security initiative first sponsored an award at the Intel ISEF in 2015. The award was created to encourage high school students to pursue scientific research in cybersecurity and related fields. The award is open to high school students who are competing in the ISEF Finals. To be able to attend, the finalists had to have won at least one recognized science fair. In 2017, a new award was added to recognize outstanding mathematics contributions. The Science of Security and Mathematics Special Awards presented at ISEF promote study and research in science areas that can assure and protect cyberspace interactions in an increasingly interconnected world. NSA RD awarded a total of $6,000 to ISEF participants this year.

The Research Directorate's Science of Security initiative is dedicated to developing a scientific discipline focused on cyber security. The Research Directorate uses ISEF to spark student interest in research protecting, “cyberspace interactions in an increasingly interconnected world.” At Intel ISEF 2017, members of RD engaged with students, parents, and other STEM enthusiasts, to educate participants on the career opportunities that exist within the Agency and the cyber security profession as a whole.

NSA representatives worked with other Intelligence Community representatives to staff the Office of the Director of National Intelligence (ODNI) booth. The three NSA SoS judges and two NSA Mathematics judges, assisted by researchers from the Veterans Administration and the National Institute of Standards and Technology (NIST), reviewed all of the 1800 submitted projects to determine their relevance to SoS and Mathematics.

In addition to reviewing the posters that described the projects, the judges read the supplemental material provided with the posters (such as log books, reports, and other documentation) in order to determine how well the topic fit in the Science of Security domain. Key aspects considered were whether the submitters were doing scientific work and increasing understanding and whether it was related to security. Three of the posters they reviewed were submitted by prior winners of NSA SoS Intel ISEF awards. Of the 1800 submissions, the SoS judges selected 62 projects for further review and interviewed 20 students. They also recognized another 25 projects with SoS Appreciation cards.

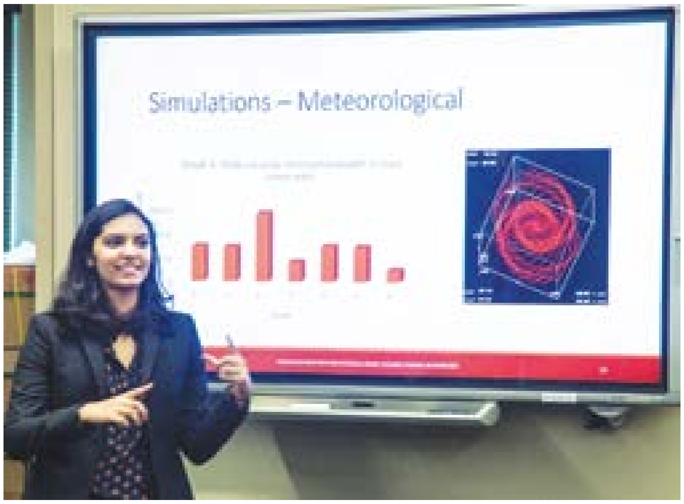

Rucha Joshi, *above,* from Westwood High School in Austin, Texas, took home the $3000 First Place Science of Security award for her project, “Power-efficient, Delay and Spatial Error Tolerant, Dynamic 3D Network Analysis.”

Joshi designed and coded an algorithm to determine optimal low-power communication solutions in evolving networks. The algorithm can adapt to changing environments, constantly calculating the most effective pathway for information to get from point A to point B. Joshi was one of the runner-up awardees for the SoS award in 2016.

Other SoS awardees were as follows:

- 2nd Place in Science of Security - Holly Jackson - $1,000 Prize

Unlocking History: An Algorithm to Virtually Unfold 3D Computed Tomography Scans of Unopened Historical Documents
Notre Dame High School, San Jose, CA

- 2nd Place in Science of Security - Mihir Patel and Nikhil Sardana - $1,000

Automating Identification of Terrorist Recruitment on Social Media Platforms
Thomas Jefferson High School for Science and Technology in Alexandria, VA.

- Honorable Mention in Science of Security - Nicky Wojtania, *above.*

Cellulose Nanocrystals for Security Applications: Embedding Non-Optical Signatures Provided by Nanoparticles into Cellulose Nanocrystal Chiral Nematic Films
Plano West Senior High School, Plano TX

- Honorable Mention in Science of Security - Mary Catherine Lorio

Quantum Eraser with Applications to Optical Quantum Information Processing of Polarization-Encoded Qubits
St. Joseph’s Academy, Baton Rouge, LA

- Honorable Mention in Science of Security - Michael Litt

MATCHLESS: A Linear Algebraic Approach to Duplicate File Identification
Orange High School, Pepper Pike, OH

Two of the SoS awardees subsequently applied for the NSA summer intern program.

Shobhita Sundaram, from Greenwich High School in Greenwich, Connecticut, received the $1000 First Place Mathematics award for her project, “Detection of Premalignant Pancreatic Cancer via Computational Analysis of Serum Proteomic Profiles.” Sundaram built a model that helps identify significant pancreatic cancer biomarker proteins quickly, enabling earlier diagnosis. Her software-based program examines the Mass Spectral data of a patient's blood and predicts with greater than 80% accuracy whether that person would develop pancreatic cancer in the future based on the presence of inter-related protein fragments contained within the patient's blood analysis.

Other Mathematics awardees were as follows:

- Honorable Mention in Mathematics - Carson Cato

Optimizing the Search for
Mersenne Primes
Arkansas School for Mathematics, Sciences and the Arts, Hot
Springs, AR

- Honorable Mention in Mathematics - Tassilo Schwarz

Done Defense System: Detection,
Tracking, Classification and Targeting of Flight Objects in 3D and
Real Time
Johannes-Heidenhain-Gymnasium Traunreut, Traunreut, Bayern
Germany

The Research Directorate recognized several of the SoS Intel ISEF awardees in a ceremony in July and the Mathematics winner in September. SoS winners Rucha Joshi, Mihir Patel, Nikhil Sardana, and Nicky Wojtania visited NSA and presented their research to NSA researchers. They also met with RD leadership and toured selected NSA mission offices.

Adam Tagert, Science of Security Technical Lead, explained, "These awards we hope inspire the greater science community to think about cybersecurity in their work and advancing cyber security in their projects. The cyber domain is constantly evolving. Therefore, it is imperative for current STEM professionals to seek future talent. NSA RD is committed to this search. We look to grow a community of researchers doing scientifically backed research and support rigorous security research measures, and we advance these two goals with our participation at ISEF."

SECTION 3

Growing the Science of Security

The Science of Security initiative's premier evernt, the 2017 Hot Topics in the Science of Security: Symposium and Bootcamp (HotSoS 2017) attracted participants from academia, industry, and government for the two-day workshop held in Hanover, MD. Sponsored by NSA in cooperation with the Association for Computer Machinery (ACM), HotSoS 2017 was hosted by the University of Maryland Lablet and the Vanderbilt University Science of SecUre and REsilient Cyber-Physical Systems (SURE) project. While HoTSoS 2017 again focused on the SoS Hard Problems, organizers also solicitated papers having specific applications to the topics of privacy and the security of Cyber-Physical Systems (CPS). The Science of Security Virtual Organization (SoS-VO) grew to over 1300 members in 2017 and continued to provide a centralized location for cybersecurity research, events, and news. By leveraging Lablet and SURE researchers, the Paper Competition, the Intel ISEF awards, HotSoS and the SoS-VO, the SoS community grew to encompass close to 200 institutions and thousands of researchers around the world in 2017. Details on HotSoS, the SoS-VO, and other outreach efforts are found in the following section.

sos-vo.org

NSA's Science of Security and Privacy Research Network

498 publications from171 Institutions (4 Lablets, 25 Sub-Lablets, 4 SURE, 138 Collaborators)

(superscript numbers represent the number of authors from that institution)

Air Force Research Laboratory (AFRL) - *Kirtland, NM & Rome, NY, USA*
Amazon Web Services (AMAZON) - *Washington, D.C., USA*
Argonne National Laboratory (ANL) - *Lemont, IL, USA*
Ariel University (ARIEL) - *Ariel, Israel*
AT&T Labs Research (AT&T) - *Florham Park & Bedminster, NJ, USA and University of Warwick, United Kingdom*
Auburn University (AUBURN) - *Auburn, AL, USA*
Bar-Ilan University (BIU) - *Tel Aviv, Israel*
Bard College (BARD) - *Annandale-On-Hudson, NY, USA*
Beihang University (BUAA) - *Beijing, China*
Boise State University (BSU) - *Boise, ID, USA*
Bosch Research and Technology Center (BOSCH) - *Pittsburgh, PA, USA*
Brooklyn College (BC) - *Brooklyn, NY, USA*
California Polytechnic State University (CAL POLY) - *San Luis Obispo, CA USA*
Carnegie Mellon University (CMU) - *Pittsburgh, PA, USA*
CERT Coordination Center (CERT/CC) - *Pittsburgh, PA, USA*
Chinese Academy of Sciences (CAS) - *Beijing, China*
Clarkson University (CCT) - *Potsdam, NY, USA*
Concordia University (CU CANADA) - *Montreal, Quebec, Canada*
Cornell University (CU) - *Ithaca, NY, USA*
Dartmouth College (DC) - *Hanover, NH, USA*
Dartmouth-Hitchcock Medical Center (DHMC) - *Lebanon, NH, USA*
DataVisor (DATAVISOR) - *Mountain View, CA, USA*
École Polytechnique Fédérale de Lausanne (EPFL) - *Lausanne, Switzerland*
Eindhoven University of Technology (TU/e) - *Eindhoven, Netherlands*
ETH Zürich (ETH ZÜRICH) - *Zürich, Switzerland*
Fordham University School of Law (FORDHAM) - *New York, NY, USA*
Fraunhofer Institute of Experimental Software Engineering (FRAUNHOFER IESE) - *Kaiserslautern, Germany*
Fujitsu Laboratories (FUJITSU) - *Sunnyvale, CA, USA*
George Mason University (GMU) - *Fairfax, VA, USA*
George Washington University (GWU) - *Washington, D.C., USA*
Georgia Institute of Technology (GT) - *Atlanta, GA, USA*
Goethe University (JWGU) - *Frankfurt, Germany*
Google Headquarters (GOOGLE) - *Mountain View, CA, USA*
Hampton University (HU) - *Hampton, VA, USA*
Hewlett Packard Laboratories (HP) - *Palo Alto, CA, USA*
IBM, T.J. Watson Research Center (IBM) - *Yorktown Heights, NY, USA*
Idaho National Laboratory (INL) - *Idaho Falls, ID, USA*
Illinois Institute of Technology (IIT) - *Chicago, IL, USA*
IMDEA Software Institute (IMDEA) - *Madrid, Spain, USA*
Indiana University (IU) - *Bloomington, IN, USA*
Institute for Molecular Manufacturing (IMM) - *Palo Alto, CA, USA*
Intel Labs (INTEL) - *Hillsboro, OR, USA*
International Computer Science Institute at UC Berkeley (ICSI) - *Berkeley, CA, USA*
Iowa State University (ISU) - *Ames, IA, USA*
Isfahan University of Technology (IUT) - *Isfahan, Iran*
Johns Hopkins University (JHU) - *Baltimore, MD, USA*
King's College (KINGS) - *Wilkes-Barre, PA, USA*
KTH Royal Institute of Technology (KTH) - *Stockholm, Sweden*
Lancaster University (LU) - *Lancaster, England, United Kingdom*
Lawrence Livermore National Laboratory (LLNL) - *Livermore, CA, USA*
Lewis & Clark College (L&C) - *Portland, OR, USA*
Massachusetts Institute of Technology (MIT) - *Cambridge, MA, USA*
Microsoft (MS) - *Redmond, WA, USA*
Microsoft Research (MSR) - *Redmond, WA, USA*
Nanyang Technological University (NTU) - *Western Water Catchment, Singapore*
National Center for Supercomputing Applications (NCSA) - *Urbana, IL, USA*
National Institute of Standards and Technology (NIST) - *Gaithersburg, MD, USA*
National Security Agency (NSA) - *Fort Meade, MD, USA*
National Technical University of Athens (NTUA) - *Athens, Greece*
National University of Sciences and Technology (NUST) - *Islamabad, Pakistan*
Navy Research Laboratory (NRL) - *Washington, D.C., USA*
NEC Laboratories America (NEC LABS) - *Princeton, NJ, USA*
NetFlix (NETFLIX) - *Los Gatos, CA, USA*
New York University, Abu Dhabi (NYUAD) - *Abu Dhabi, United Arab Emirates*
Newcastle University (NEWCASTLE UK) - *Newcastle upon Tyne, England, United Kingdom*
North Carolina State University (NCSU) - *Raleigh, NC, USA*
Northeastern University (NEU) - *Boston, MA, USA*
Northwestern University (NWU) - *Evanston, IL, USA*
Oregon Health and Science University (OHSU) - *Portland, OR, USA*
Oregon State University (OSU) - *Corvallis, OR, USA*
Otto Von Guericke University Madgeburg (OVGU) - *Magdeburg, Germany*
Peking University (PKU) - *Beijing, China*
Pennsylvania State University (PSU) - *State College, PA, USA*
Polytechnique Montreal (POLYMTL) - *Montreal, QC, Canada*
Princeton University (PRINCETON) - *Princeton, NJ, USA*
Providence College (PC) - *Providence, RI, USA*
Purdue University (PU) - *Lafayette, IN, USA*
Queen's University (QUEEN'S) - *Kingston, ON, Canada*
Rensselaer Polytechnic Institute (RPI) - *Troy, NY, USA*
Rice University (RICE) - *Houston, TX, USA*
Rochester Institute of Technology (RIT) - *Rochester, NY, USA*
Royal Holloway, University of London (ROYAL HOLLOWAY UNIVERSITY) - *London, England, United Kingdom*
RSA Laboratories (RSA) - *Cambridge, MA, USA*
Ruhr University Bochum (RUB) - *Bochum, Germany*
Rutgers University (RU) - *Camden, NJ, USA*
Samsung Research America-Dallas (SRA) - *Dallas, TX, USA*
Sandia National Laboratories (SANDIA) - *Albuquerque, NM, USA*
SentiMetrix (SENTIMETRIX) - *Bethesda, MD, USA*
Siemens Corporation (SIEMENS) - *Princeton, NJ, USA*
Simula Research Laboratory (SRL) - *Fornebu, Norway*
Singapore Management University (SMU) - *Bras Basah, Singapore*
Software Engineering Institute (SEI) - *Pittsburgh, PA, USA*
Stanford University (SU) - *Stanford, CA, USA*
Swansea University (SWANSEA) - *Swansea, Wales, United Kingdom*
Swarthmore College (SC) - *Swarthmore, PA, USA*
Symantec Research Labs (SYMANTEC) - *Culver City, CA, USA & Sophia Antipolis, France*
Technical University Braunschweig (TU BRAUNSCHWEIG) - *Braunschweig, Germany*
Technischen Universität Darmstadt (TU DARMSTADT) - *Darmstadt, Germany*
Temple University (TEMPLE) - *Philadelphia, PA, USA*
Tennessee Technical University (TTU) - *Cookeville, TN, USA*
Texas A&M University (TAMU) - *College Station, TX, USA*
The Australian National University (ANU) - *Canberra, Australia*
The University of Tokyo (TODAI) - *Tokyo, Japan*
Towson University (TU) - *Baltimore, MD, USA*
Tsinghua University (THU) - *Beijing, China*
United States Military Academy (USMA) - *Westpoint, NY, USA*
United States Naval Academy (USNA) - *Annapolis, MD, USA*
Universidad Politécnica de Madrid (UPM) - *Madrid, Spain*
Universidade Federal de Alagoas (UFAL) - *Maceió, Brazil*
Universidade Federal de Campina Grande (UFCG) - *Campina Grande, Brazil*
Universidade Nova de Lisboa (NOVA) - *Lisbon, Portugal*
Università della Svizzera italiana (USI) - *Lugano, Switzerland*
Universität des Saarlandes (UDS) - *Saarbrücken, Germany*
Universitat Politècnica de València (UPV) - *Valencia, Spain*
Université catholique de Louvain (UCL) - *Louvain-la-Neuve, Belgium*
Universiteit Twente (UT) - *Enschede, Netherlands*
University at Albany-SUNY (UALBANY) - *Albany, NY, USA*
University of Alabama (UA) - *Tuscaloosa, AL, USA*
University of California, Berkeley (UCB) - *Berkeley, CA, USA*
University of California, Irvine (UCI) - *Irvine, CA, USA*
University of California, San Diego (UCSD) - *San Diego, CA, USA*
University of California, Santa Barbara (UCSB) - *Santa Barbara, CA, USA*
University of Chinese Academy of Sciences (UCAS) - *Beijing, China*
University of Colorado (CU BOULDER) - *Boulder, CO, USA*
University of Hawai'i at Mānoa (UH MĀNOA) - *Mānoa, HI, USA*
University of Idaho (U IDAHO) - *Moscow, ID, USA*
University of Illinois Urbana-Champaign (UIUC) - *Urbana-Champaign, IL, USA*
University of Jinan (UJN) - *Jinan, China*
University of Kaiserslautern (TU KL) - *Kaiserslautern, Germany*
University of Leeds (LEEDS) - *Leeds, England, United Kingdom*
University of Lisbon (ULISBOA) - *Lisbon, Portugal*
University of Maryland (UMD) - *College Park, MD, USA*
University of Maryland, Baltimore County (UMBC) - *Baltimore, MD, USA*
University of Melbourne (UNIMELB) - *Melbourne, Australia*
University of Michigan (UMich) - *Ann Arbor, MI, USA*
University of Nebraska at Lincoln (UNL) - *Lincoln, NB, USA*
University of New Mexico, Albuquerque (UNM) - *Albuquerque, NM, USA*
University of North Carolina at Chapel Hill (UNC-CH) - *Chapel Hill, NC, USA*
University of North Carolina at Charlotte (UNCC) - *Charlotte, NC, USA*
University of Nottingham (UON) - *Nottingham, United Kingdom*
University of Oslo (UiO) - *Oslo, Norway*
University of Oxford (OXFORD) - *Oxford, England, United Kingdom*
University of Passau (PASSAU) - *Passau, Germany*
University of Pennsylvania (PENN) - *Philadelphia, PA, USA*
University of Pittsburgh (PITT) - *Pittsburgh, PA, USA*
University of Porto (UP) - *Porto, Portugal*
University of São Paulo (USP) - *São Paulo, Brazil*
University of Southern California (USC) - *Los Angeles, CA, USA*
University of Surrey (UNIS) - *Guildford, United Kingdom*
University of Texas at Arlington (UTA) - *Arlington, TX, USA*
University of Texas at Austin (UT AUSTIN) - *Austin, TX, USA*
University of Texas at Dallas (UT DALLAS) - *Dallas, TX, USA*
University of Texas at San Antonio (UTSA) - *San Antonio, TX, USA*
University of Verona (UNIVR) - *Verona, Italy*
University of Virginia (UVA) - *Charlottesville, VA, USA*
University of Warwick (UWAR) - *Coventry, England, United Kingdom*
University of Washington (UW) - *Seattle, WA, USA*
University of Waterloo (UOFW) - *Waterloo, ON, Canada*
University of Wisconsin (UW-MADISON) - *Madison, WI, USA*
University Politehnica of Bucharest (UPB) - *Bucharest, Romania*
US Army Research Lab (ARL) - *Adelphi, MD, USA*
Vanderbilt University (VU) - *Nashville, TN, USA*
Verisign Labs (Verisign) - *Reston, VA, USA*
Viatec Research (VIATEC) - *Raleigh, NC, USA*
Victoria University of Wellington (VUW) - *Wellington, New Zealand*
Virginia Polytechnic Institute and State University (VT) - *Blacksburg, VA, USA*
Wandoujia Lab (WANDOUJIA) - *Beijing, China*
Wayne State University (WSU) - *Detroit, MI, USA*
Western University (UWO) - *London, Ontario, Canada*
Wroclaw University of Technology (WRUT) - *Wroclaw, Poland*
Zhejiang University (ZJU) - *Zhejiang, China*

704 Authors

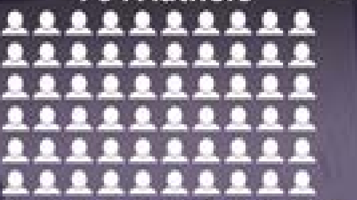

27 Countries

- Abu Dhabi
- Australia
- Belgium
- Brazil
- Canada
- China
- England
- France
- Germany
- Greece
- Iran
- Israel
- Italy
- Japan
- Netherlands
- New Zealand
- Norway
- Pakistan
- Poland
- Portugal
- Romania
- Singapore
- Spain
- Sweden
- Switzerland
- United States
- Wales

5 Continents

http://sos-vo.org/map

HotSoS

The Hot Topics in the Science of Security: Symposium and Bootcamp (HotSoS) is an annual event which is sponsored by the National Security Agency (NSA) in cooperation with the Association for Computer Machinery (ACM). HotSoS is focused on addressing the fundamental problems of cybersecurity in a principled manner and brings together researchers from diverse disciplines to promote advancement of work related to the Science of Security. As in previous years, HoTSoS 2017 created a forum for dialogue centered upon the development and advancement of research, specifically focusing on the Science of Security Hard Problems. For the 2017 event, HoTSoS also solicitated papers having specific applications to privacy, broadly construed, and the security of Cyber-Physical Systems (CPS).

Hosted by the University of Maryland Lablet and the Vanderbilt University SURE Project (Science of SecUre and REsilient Cyber-Physical Systems), HotSoS 2017 was held April 4-5 in Hanover, MD. There were 135 attendees, almost half of whom were from the government. This was the fourth HotSoS, bringing researchers together to interact and to see presentations demonstrating rigorous scientific approaches to identify, prevent and remove cyber threats. A major continuing focus of HotSoS is the advancement of scientific methods, including data gathering and analysis, experimental methods, and mathematical models for modeling and reasoning. Outside speakers addressed Science of Security from the perspectives of Cyber-Physical Systems, privacy, differential privacy and data analysis. A panel discussion, papers, tutorials, and poster sessions rounded out the agenda. Details on each of the agenda elements are provided below.

HotSoS 2017 proceedings have been published by ACM and are available online at the ACM Digital Library at https://dl.acm.org/citation.cfm?id=3055305&picked=prox.

The HotSoS 2017 Program Chair was Jonathan Katz, Principal Investigator (PI) of the University of Maryland Lablet. Co-chairs were Michel Cukier, University of Maryland Lablet Co-PI, and Xenofon Koutsoukos, PI for the SURE project at Vanderbilt University.

HotSoS Program Committee Members:

Adam Aviv, United States Naval Academy
Travis Breaux, Carnegie Mellon University
Alvaro Cardenas, University of Texas, Dallas
Will Enck, North Carolina State University
Chris Gates, Symantec
Limin Jia, Carnegie Mellon University
Michael Maass, Bosch Research
Sam Malek, George Mason University
Michelle Mazurek, University of Maryland
Sayan Mitra, University of Illinois at Ubrana-Champaign
Charles Morisset, Newcastle University
Bill Sanders, University of Illinois at Urbana-Champaign
Bill Scherlis, Carnegie Mellon University
Jessica Staddon, North Carolina State University
Adam Taggert, National Security Agency
Eugene Vorobeychik, Vanderbilt University
Shouhuai Xu, University of Texas, San Antonio

Welcome

Dr. Deborah Frincke, Director of NSA's Research Directorate, introduced the event. She noted that "NSA continues a strong commitment to SoS and the connections and collaborations that have been growing in recent years." She cited the value of the work being done in SoS as having an impact. "The work is showing up in various places, underpinnings of security are improving, and we are seeing the impact of using the principles found in SoS," she added. "The need is both for dramatic discovery applying basic principles, then transition to practice and baking those discoveries and principles in."

Research Papers

Nine refereed papers were presented on research studies related to CPS properties, scientific reporting quality, optimization of security investments, building a privacy incident database, tradeoffs between privacy and utility, phishing training, variations in attack encounters, uncertainty in network security analysis, and security practice adherence in software development.

1. Leveraging Unique CPS Properties to Design Better Privacy-Enhancing Algorithms

Jairo Giraldo, Alvaro A. Cardenas, and Murat Kantarcioglu, University of Texas at Dallas

Cyber-Physical Systems have unique properties that can be used to design new privacy-enhancing technologies that minimize the negative impact to the utility of CPS. In this paper the authors show two examples of these properties. The first example looks at how differential privacy degrades CPS performance due to the large noise addition, and how the inherent noise of CPS can be leveraged to reduce the additional noise added by differential privacy algorithms, and thereby minimize the negative impact on the system utility and safety. The second example looks at the ability to sample at sensor readings on demand, and how this flexibility can be used to design adaptive sensor sampling algorithms that hide sensitive information without the need to add noise.

2. Characterizing Scientific Reporting in Security Literature: An analysis of ACM CCS and IEEE S&P Papers

Morgan Burcham, Mahran Al-Zyoud, and Jeffrey C. Carver, University of Alabama; Ehab Al-Shaer, Mohammed Alsaleh, and Fida Gilani, UNC Charlotte; Jun Jiang, UNC Chapel Hill; Akond Rahman, Hongying Du, Özgür Kafali, and Laurie Williams, NC State University

Scientific advancement is fueled by solid fundamental research, followed by replication, meta-analysis, and theory building. To support such advancement, researchers and government agencies have been working towards a "science of security". As in other sciences, security science requires high-quality fundamental research addressing important problems, and reporting approaches that capture the information necessary for replication, meta-analysis, and theory building. The goal of this paper is to aid security researchers in establishing a baseline of the state of scientific reporting in security through an analysis of indicators of scientific research as reported in the 2015 ACM Computer and Communications Security (CCS) and 2016 IEEE Security & Privacy (S&P) proceedings. To conduct this analysis, we employed a series of rubrics to analyze the completeness of information reported in papers relative to the type of evaluation used (e.g. empirical study, proof, discussion). Our findings indicated that some important information is often missing from papers, including explicit documentation of research objectives and the threats to validity. Our findings show a relatively small number of replications reported in the literature. We hope that this initial analysis will serve as a baseline against which we can measure the advancement of the science of security.

3. Optimal Security Investments in a Prevention and Detection Game

Carlos Barreto, Alvaro A. Cardenas, and Alain Bensoussan, University of Texas at Dallas

Most security defenses can be breached by motivated adversaries. According to the authors, in addition to attack-prevention technologies, firms investing in cybersecurity for their information technology infrastructure need to consider attack-detection and restoration tools to detect intruders and restore their system to a safe condition. Attackers face similar investment alternatives: they need to invest resources to finding vulnerabilities in a protected system, and once the protection has been broken, to invest in the infrastructure necessary to exploit these attacks and maintain stealthy persistence in the compromised infrastructure. The authors model these dual considerations as a dynamic programming problem between attackers and defenders and then study the Nash equilibrium of the game. Their goal is to find models and alternatives that can contribute to understanding optimal security investments in prevention and detection against advanced rational adversaries.

4. Learning a Privacy Incidents Database

Pradeep K. Murukannaiah, Rochester Institute of Technology; Chinmaya Dabral, Karthik Sheshadri, Esha Sharma, and Jessica Staddon, NC State University

A repository of privacy incidents is essential for understanding the attributes of products and policies that lead to privacy incidents. The authors describe a vision for a novel privacy incidents database and progress toward building a prototype. Key challenges in gathering such a database include bootstrapping and sustainability. Their work proposes a semi-automated framework that can recognize privacy incidents and related information from various online sources such as news, blogs, and social media. The crux of the framework is an incident classifier that identifies whether a piece of text in natural language is related to a privacy incident or not. They curate a dataset consisting of 1324 news articles of which 543 articles are about one or more privacy incidents and train the incident classifier on this dataset, considering a variety of feature engineering, feature selection, and classification techniques. This incident classifier yields an F1 measure of 93.1%, which is about 12% higher than the keyword search-based baselines they adopted.

5. On the Tradeoff between Privacy and Utility in Collaborative Intrusion Detection Systems-A Game Theoretical Approach

Richeng Jin and Huaiyu Dai; NC State University, Xiaofan He, Lamar University

Intrusion Detection Systems (IDS) are crucial security mechanisms widely deployed for critical network protection. However, conventional IDS become incompetent due to the rapid growth in network size and the sophistication of large scale attacks. To mitigate this problem, Collaborative IDS (CIDS) have been proposed in literature. In CIDS, a number of IDS exchange their intrusion alerts and other relevant data so as to achieve better intrusion detection performance. Nevertheless, the required information exchange may result in privacy leakage, especially when these IDS belong to different self-interested organizations. In order to obtain a quantitative understanding of the fundamental tradeoff between the intrusion detection accuracy and the organizations' privacy, a repeated two-layer single-leader multi-follower game is proposed in this work. Based on this game-theoretic analysis, the authors are able to derive the expected behaviors of both the attacker and the IDS and obtain the utility-privacy tradeoff curve. In addition, the existence of Nash Equilibrium (NE) is proved and an asynchronous dynamic update algorithm is proposed to compute the optimal collaboration strategies of IDS. Finally, simulation results are shown to validate the analysis.

6. Use of Phishing Training to Improve Security Warning Compliance: Evidence from a Field Experiment

Weining Yang, Aiping Xiong, Robert W. Proctor, and Ninghui Li, Purdue University; Jing Chen, New Mexico State University

The current approach to protect users from phishing attacks is to display a warning when the webpage is considered suspicious. The authors hypothesize that users are capable of making correct informed decisions when the warning also conveys the reasons why it is displayed. They chose to use traffic rankings of domains, which can be easily described to users, as a warning trigger and evaluated the effect of the phishing warning message and phishing training. The evaluation was conducted in a field experiment. They found that knowledge gained from the training enhances the effectiveness of phishing warnings, as the number of participants being phished was reduced. However, the knowledge by itself was not sufficient to provide phishing protection. They suggest that integrating training in the warning interface, involving traffic ranking in phishing detection, and explaining why warnings are generated will improve current phishing defense.

7. Global Variation in Attack Encounters and Hosting

Ghita Mezzour, International University of Rabat; Kathleen M. Carley and L. Richard Carley, Carnegie Mellon University

Countries vary greatly in the extent to which their computers encounter and host attacks. Empirically identifying factors behind such variation can provide a sound basis for policies to reduce attacks worldwide. However, the primary current approach to identify these factors consists of expert opinions with limited empirical validation. In this work, the authors empirically test hypotheses regarding social and technological factors behind such international variation. Using Symantec's Intrusion Prevention System (IPS) telemetry data collected from around 10 million Symantec customers worldwide, they found web attacks and fake applications are most prominent in Western Europe and North America. The results indicate a relationship between countries' wealth and technological sophistication and attack exposure, indicating that attackers probably target developed countries to maximize their profits. Moreover, Eastern Europe hosts disproportionate quantities of attacks. Their statistical analysis reveals a relationship between attack hosting and the combined effect of widespread corruption and computing resources. Surprisingly, China is not among the top 10 attack hosting countries and Africa hosts the smallest quantities of attacks. This work, according to the authors, has important policy implications.

8. An Approach to Incorporating Uncertainty in Network Security Analysis

Hoang Hai Nguyen, Kartik Palani, and David M. Nicol, University of Illinois at Urbana-Champaign

Attack graphs used in network security analysis are analyzed to determine sequences of exploits that lead to successful acquisition of privileges or data at critical assets. An attack graph edge corresponds to a vulnerability, tacitly assuming a connection exists and tacitly assuming the vulnerability is known to exist. This work explored use of uncertain graphs to extend the paradigm to include lack of certainty in connection and/or existence of a vulnerability. The authors extend the standard notion of uncertain graph (where the existence of each edge is probabilistically independent), as significant correlations on edge existence probabilities exist in practice, owing to common underlying causes for dis-connectivity and/or presence of vulnerabilities. That extension describes each edge probability as a Boolean expression of independent indicator random variables. The paper shows: 1) that this formalism is maximally descriptive in the sense that it can describe any joint probability distribution function of edge existence; 2) that when these Boolean expressions are monotone one can easily perform uncertainty analysis of edge probabilities; and 3) these results are used to model a partial attack graph of the Stuxnet worm and a small enterprise network and to answer important security-related questions in a probabilistic manner.

9. Surveying Security Practice Adherence in Software Development

Patrick Morrison and Laurie Williams, North Carolina State University

Software development teams are increasingly incorporating security practices into their software development processes. Little empirical evidence exists on the costs and benefits associated with the application of such security practices. Balancing the tradeoff between the costs in time, effort, and complexity of applying security practices, and the benefit of an appropriate level of security in delivered software requires measuring security practice benefits and costs. The goal of this research is to support researcher investigations of software development security practice adherence by building and validating a set of security practices and adherence measures through literature review and survey data analysis. The authors extracted 16 software development security practices from a review of the literature and established a set of adherence measures based on technology acceptance theory. They built a survey around the 13 most common practices and their adherence measures and surveyed 11 security-focused open source projects to collect empirical data to test their thesis about practice adherence. In the collected survey data, each of the 13 security practices identified was used daily by at least one survey participant; tracking vulnerabilities and applying secure coding standards are the practices most often applied daily. In the data, ease of use, effectiveness, and training, measured via Likert items, did not always show the expected theoretical relationship with practice use. In the data, training is positively correlated with practice use, while effectiveness and ease of use vary in their correlations with practice use on a practice-by-practice basis.

Keynotes

1. Security of Cyber-Physical Systems, Challenges and Approaches

Insup Lee

Insup Lee, University of Pennsylvania Professor of Computer and Information Science, spoke about the ways CPS are used to monitor and control real-world systems. He noted that securing CPS introduces additional challenges since the attack surface is increased compared to conventional systems. In addition to the cyber intrusions that apply to all computer systems, attacks on CPS can be through interference to the physical environment of CPS. A number of such attacks have emerged which suggest that conventional cyber-only security approaches will not be effective. His talk discussed approaches to making CPS resilient to cyber-physical attacks by exploiting spatial and temporal redundancy as well as dynamics of the underlying physical system.

2. Navigating Privacy Issues in a Data-Driven World

Jules Polonetsky

Jules Polonetsky, Chief Executive Officer of the Future of Privacy Forum, a non-profit organization, addressed the stresses on privacy created by technological advances. Online tracking for analytics and advertising has extended to mobile devices, interactive television and smart home devices. Social media sharing has achieved near ubiquity, with services integrating location, facial recognition, and live video sharing. Motor vehicles have become data collectors, and drones allow our public spaces to be more easily monitored. Big data strains against fair information practices of consent, limited purpose, and data minimization. Algorithmic decision-making and machine learning wreak havoc with efforts to provide transparency. Artificial Intelligence may leave us unsure as to who will even be accountable for data-driven determinations. These advances could create opportunities for progress.

3. Differential Privacy and Data Analysis

Aaron Roth

Aaron Roth, Associate Professor of Computer and Information Science at the University of Pennsylvania, gave a “friendly introduction” to differential privacy, which he described as “a rigorous methodology for analyzing data to provide provable privacy guarantees that has recently been widely deployed in several settings.” He specifically focused on the rich relationship between differential privacy and machine learning, including both the ability to do machine learning subject to differential privacy, and tools arising from differential privacy that can be used to make learning more reliable and robust even when privacy is not a concern.

Panel Discussion

Shifting the Balance in the Attack-Defend Cycle

Carnegie Mellon University Lablet PI Bill Scherlis moderated a lively panel discussion about the Defense Advanced Research Project Agency (DARPA) Cyber Grand Challenge (CGC) which took place August 4, 2016. Seven computers developed by teams of hackers played the world’s first all-machine game of “Capture the Flag”. The goal of the CGC was to accelerate the development of advanced, autonomous systems that can detect, evaluate, and patch software vulnerabilities in computers and networked electronic devices before adversaries have a chance to exploit them.

Panelists David Melski (GrammaTech), Charles Nelson (US Cyber Command), Sam Septembre (US Navy), and Yan Shoshitaishvili (University of California, Santa Barbara) discussed their participation in and lessons learned from the challenge. Panelists suggested that CGC revealed what is currently lacking in achieving cyber autonomy. Moving forward, when full autonomy is achieved, there will still be a need to address the human factor as human intelligence can and should be injected into the system. A video of the CGC final round is available at https://www.youtube.com/watch?v=n0kn4mDXY6I.

Posters

Twenty-three posters were presented, and one was awarded Best Poster. The evaluators looked to recognize cybersecurity research posters with scientific rigor, clarity of presentation, and global impact. The Best Poster award is to encourage scientists across multiple disciplines to address the fundamental problems of security in a principled manner. This year’s winning poster is an exemplar of the research needed to improve both the confidence we gain from scientific results and also the capacity and efficiency through which we address increasingly technical problems.

The poster “Is the Guardian Capable? A Routine Activity Theory Approach to Cyber Intrusion on Honeypot Systems” by Michael Becker, Michel Cukier, and David Maimon, a team of researchers from the University of Maryland, was named Best Poster.

1. Is the Guardian Capable? A Routine Activity Theory Approach to Cyber Intrusion on Honeypot Systems

Michael Becker, Michel Cukier, David Maimon

Research on the criminological side of system trespassing (i.e. unlawfully gaining access to a computer system) is relatively rare and has yet to examine the effect of the presence of other users on the system during the trespassing event (i.e. the time of communication between the trespasser’s system and the infiltrated system). This poster begins to analyze this relationship drawing on the concept of capable guardianship under Routine Activity Theory. Data were collected from a randomized control trial of target computers deployed on the Internet network of a large U.S. university. This poster examined whether the number (one or multiple) and type (administrative or non-administrative) of computer users present on a system reduced recidivism by trespassers on targeted systems. Results indicate that neither the number nor type of condition produced a deterrent effect through the role of increased capable guardianship on the target system.

2. A Value Model for Implementing Cyber Metrics and Best Practices

Natalie Scala, Emil Manuel, Rachel Fredman, Jasmin Farahani, Paul Goethals

Research in the five Hard Problems has led to innovative and interdisciplinary advances in cybersecurity. However, a disconnect exists in transferring the research into implementable industry solutions. For example, as of October 2016, 110 papers in the Hard Problem of Security Metrics and Models had been indexed by the Science of Security. An organization looking to improve its cybersecurity posture may be overwhelmed by the sheer volume of options; organizations who lack cyber expertise may shy away from implementing metrics altogether by not knowing where to start. This research promotes strengthening an organization’s security posture by formulating a value model to identify the preferred metrics and best practices for defending against cyber-attacks or intrusions. These practices may differ by organization, based on demographics and history such as size of firm and prior experiences of cyber-attacks and/or breaches. To identify the preferred metrics and best practices, we employ Multiple Objective Decision Analysis (MODA), which is grounded in utility theory, to evaluate a set of candidate metrics and best practices against desired attributes in cyber defense. This framework can be applied to any organization, customized by data applicable to that firm. We illustrate the model for the general supply chain. We identify six attributes that are valued in cybersecurity (data integrity, end to end security, cloud security, security policies, intrusion and threat detection, patch sets and hot fixes) and utilize NSA’s Information Assurance Directorate (IAD) Top Ten Information Assurance Mitigation Strategies as candidate metrics and best practices. We employ the combined standard of Parnell, et al. for model data collection, interviewing subject matter experts and researching policy documents. Preliminary results identify controlling administrative privileges, limiting workstation-to-workstation communication, and using web domain name system recognition as the top three preferred metrics and best practices for supply chain organizations. Continuing research will evaluate the sensitivity of these results as well as index model attributes for other industries.

3. Advanced Metrics for Risk-Based Attack Surface Approximation

Christopher Thiesen, Laurie Williams

Despite a growing number of threats, the software engineering community still faces a critical deficit of trained security professionals for defending against cyber-attacks. To combat this shortage, efficient prioritization of the effort of security professionals is needed. To address this issue, we present Risk Based Attack Surface Approximation (RASA), which uses crash dump stack traces to approximate the attack surface of a system. The poster describes several RASA metrics that could help security effort prioritization. These metrics include temporal metrics (how the attack surface changes over time), shape metrics (how code artifacts are interconnected), and depth metrics (how far code artifacts are from the entry and exit points of a system)

4. An Instruction Set Randomization (ISR) Framework for Developing Secure and Resilient CPS

Bradley Potteiger, Zhenkai Zhang, Xenofon Koutsoukos

A number of successful attacks against CPS have demonstrated that security and resilience of CPS is a critical problem, and that new methods and techniques are required to build high confidence systems. Among these attacks, code injection attacks are cited as a common attack vector taking advantage of software input processing vulnerabilities in order to run malicious code in a control process. Instruction set randomization is a very effective technique to mitigate against code injection attacks, randomizing the instruction set architecture of machine code for the purpose of making originally valid injected attack code invalid and un-executable. When the injected attack code is executed, the Central Processing Unit (CPU) will recognize the invalid instructions and the process will crash. However, system crashing is unacceptable in safety critical systems and continuous, reliable functionality has to be maintained. This poster illustrates a runtime framework that utilizes ISR techniques to mitigate against code injection attacks as well as integrating a control manager process to prevent the possibility of system crashing by reliably switching to a safe controller alternative. A traffic case study scenario is utilized to demonstrate the implementation of the control framework.

5. Analysis of Two Parallel Surveys on Cybersecurity: Users & Security

Jim Blythe, Ross Koppel, Vijay Kothari, Sean Smith

This poster described a pilot study comparing cybersecurity professionals' views and expectations of authentication workarounds with those of users. We examined the differential perceptions of cybersecurity professionals and of general users about access rules and passwords. We found that understanding the perceived reasonableness of cybersecurity policies offers opportunities for improvement, even if one finds users to be naive or misinformed. Only by understanding users' perceptions can we hope to better inform them and to respond to their needs. While both general users and cybersecurity professionals expressed dissatisfaction with access rules and passwords, their perceptions manifest many misunderstandings, and thus approaches, to improving security. A well-informed cybersecurity professional who understands the perceptions of general users will be in a better position to address users' concerns, to establish user trust, and to educate the user by dispelling user misperceptions and legitimizing existing (or new and better) security measures. Since this was only a pilot study, sample sizes were very small; generalizing to larger populations is unwise. We are, however, planning on expanding the research to larger samples and differing populations. The poster also addressed the next steps in this study.

6. Cyber Knowledge is Here, but Not Evenly Distributed

Susan Campbell, Meredith Hughes, Scott Jackson, Valerie Karuzis, Sunhee Kim, Alison Tseng

Current selection systems in cybersecurity tend to pick candidates based on existing knowledge of cybersecurity concepts and procedures. As part of a larger study designed to create an aptitude test for cyber operations, we surveyed a sample of students and community members to determine their cybersecurity knowledge, cognitive abilities, and demographic characteristics. We found that the best prediction for cybersecurity knowledge came from a background in computer science, identifying as a gamer, and general academic ability. The top 25% of candidates on cybersecurity knowledge were not as diverse as the top 25% of candidates on ability measures, however. This finding shows that selection based only on cybersecurity knowledge may systematically overlook talented candidates.

7. Exploring Defect Categories for Infrastructure as Code

Akond Rahman, Laurie Williams

Infrastructure as Code (IaC) technology refers to the process of automated and reproducible management of infrastructure that includes operating systems and software dependencies. IaC technologies, such as Ansible and Puppet, aim to reduce deployment errors and deployment overhead through automation. Practitioners have observed multiple benefits of using IaC technologies including rapid delivery of software changes and fewer configuration errors. Yet, IaC scripts can be complex in nature containing hundreds of lines of code, exposing IaC scripts to defects that are potentially difficult to debug. For example, a small change in infrastructure configuration incorrectly enabled Microsoft Azure blob storage, causing a large-scale system outage for Microsoft Azure. Systematic investigation of what categories of defects occur can help in understanding the nature of IaC defects and in taking appropriate actions to mitigate those defects. The goal of this poster is to aid software practitioners in understanding IaC-related defects through categorization of defects that occur in IaC scripts.

8. Factors for Differentiating Human from Automated Attacks

Masooda Bashir, Kelly Greeling, Alex Withers

Many Intrusion Detection Systems and Intrusion Protection Systems utilize behavior-based methodology, which seeks to identify a baseline of normal use that is then compared against real-time and non-real-time events in an effort to locate malicious activity. However, the rise of automated attacks has created a great deal of noise for security personnel to wade through to identify malicious behavior. If a human-based attack is significantly different than an automated attack, it would be extremely useful for security personnel to have a way to separate the behavior of an automated cyberattack tool from that of a human actor, as this would allow them to create separate tools to deal with each. This poster is an exploratory study into whether it is viable to use event time-difference and event pattern-occurrence as factors in behavior-based Intrusion Detection Systems for identifying the difference between human and automated program behavior.

9. FARM: A Toolkit for Finding the Appropriate Level of Realism for Modeling

Jim Blythe, Ross Koppel, Vijay Kothari, Christopher Novak, Sean Smith

In previous work we have developed the DASH agent platform aimed at simulating human behavior in order to understand and mitigate the negative consequences of human security decision-making for security protocols. The DASH architecture, shown in the poster, includes a dual-process approach with modules for rational and instinctive behavior, helping account for differing human responses to situations depending on factors such as level of attention, fatigue, stress, or effect. It includes a reactive planning module that allows the agent to exhibit goal-directed behavior that is responsive to changes in the environment and supports different mental models that humans may use to reason about a situation.

10. Flawed Mental Models Lead to Bad Cyber Security Decisions: Let’s Do a Better Job

Jim Blythe, Ross Koppel, Vijay Kothari, Sean Smith

Conventional computer security wisdom implicitly assumes models about humans and human organizations, and these models then translate into practices that conventional wisdom blesses as good. Unfortunately, fieldwork (by us and many, many others) shows that these models are not necessarily true, and that the practices resulting from them do not necessarily make things better, and in fact can make things worse. If we blindly apply conventional wisdom without validating the assumptions upon which it is based, we don’t see the security gains that we might expect. This gives rise to uncanny descents, scenarios where we turn up security knobs with the expectation that aggregate security will improve, but we instead observe that things get worse. We posit that these problems all result from the same underlying cause: flawed models of the interaction of humans and technology. Security policies, mechanisms, and recommendations are designed according to a human-conceived model of security, whether designed directly by the policy designer (e.g., by following tradition) or indirectly by the utilization of risk assessment or other security tools that are created by humans. In previous work, we’ve characterized these causes as mismorphisms, “mappings that fail to preserve structure,” especially mismatches between the security practitioner’s mental model, the user’s mental model, the model arising from system data, and the reality. This situation gives rise to a grand challenge: how do we unravel this problem? Flawed models lead to bad decisions. We need a way to make better decisions. This poster explores this problem, and presents both our current work as well as where we plan to go next. A solution would likely have several components: effective ways to talk about aggregate security in practice, effective ways to discover and correct flaws in mental models, and effective ways to make better security decisions despite such flaws.

11. Formulating a Method for Using Search Query Trends as a Measure of Mass-User Interest

Andrew McNamara, Shaown Sarker, Jessica Staddon

It is difficult to estimate the mass-user interest in a specific topic over time often due to the sensitivity of the self-reported data to bias (such as privacy-related surveys), which is a part of the privacy paradox. In 2015, Pew Research Center published a report discussing how the on-line activities and habits of American adults have changed since Edward Snowden began leaking documents about government intelligence programs. Their survey identified how people changed their on-line behavior as they became aware of surveillance. These findings indicate that revelations of privacy-invasive programs make people change their on-line behaviors to better protect their sensitive data.

Privacy-related decision-making like this has been actively studied for some time with much of the research focusing on the “privacy paradox” in which behavioral decisions seem at odds with reported preferences. In our work, we propose a novel methodology to measure the population’s interest over time by creating a robust index based on on-line searches and leveraging of the search trends of the queries in this index. We verify our index-based process using the Pew Research Center report on the Snowden revelations.

12. Leading the Convoy: What Happens When They Know What They’re Doing?

Christopher Mayhorn, Carl Pearson, Allaire Welk

In an increasingly technological environment, security-critical and safety-critical situations often give a system operator information from both human and automated sources simultaneously. This study was conducted to understand the personal characteristics and situational variables that affect reliance on human or automated decision aids, when the decision aids are in conflict. Based on the experimental protocol used by Lyons & Stokes, this study compares two samples, military and civilian. The civilian sample (N = 126, mean age = 19) was taken from an undergraduate population at a large Southeastern university. The military sample (N = 29, mean age = 38) was taken from a group of special operations members enrolled in an education program at a U.S. Army base. The experimental protocol used the Convoy Leader software platform, tasking participants to choose an optimal convoy route based on provided decision aid information. Participants marked their route decision, which indicated reliance on the automated tool recommendation, the human decision aid recommendation, or their own reasoning (the route not recommended by a decision aid). Participants also reported their perceived risk of the situation and perceived workload (NASA TLX) of the overall task. Results indicated that there were significant differences among civilian and military samples on reliance measures overall. A significantly higher proportion of the military sample relied upon themselves (on the non-recommended route) compared to the undergraduate sample. Non-significantly higher proportions of the civilian sample, compared to the military sample, relied on the automated tool and the human decision aid. Results also showed differences in situational perceptions. The military sample perceived significantly less risk in the situation than the civilian sample. The military sample also experienced significantly less workload than the civilian sample. The included measures indicate that the military sample experienced the situation differently than the civilian sample and came to different decisions. The background of participants in experimental simulations is vitally important, as it appears different life experiences and specialized training can drastically shift situational perceptions as well as situational outcomes in this context. This should be considered when designing experiments, especially in the population that conclusions are meant to generalize to. Designing decision support systems for both civilian and military populations should consider that the same situation will likely be perceived differently, and the same given information may not lead to the same decision output.

13. Learning Factor Graphs for Preempting Multi-Stage Attacks in Cloud Infrastructure

Phuong Cao, Ravishankar Iyer, Zbigniew Kalbarczyk

We discuss methods for: 1) learning parameters of multi-variate factor functions that capture relations among the events representing behavior of both a user and an attacker; and 2) construction of factor graphs to reason about an attack state with the purpose of preemptively detecting malicious activities. Our work is driven by real attacks reported in the wild. In the context of this analysis, we focus on multi-stage attacks that can be represented by five distinct stages: initial compromise, host hopping, escalation of privilege, maintaining presence, and delivery of payloads. User and attacker activities are represented by events, which are derived from the log files collected at runtime by security monitoring tools, such as intrusion detection systems, network flows, and system logs.

14. Multi-agent System for Detecting False Data Injection Attacks Against the Power Grid

Esther Amullen, Hui Lin, Zbigniew Kalbarczyk

In the power grid, control decisions and subsequent actions that directly impact the operation of the power grid are made based on estimation data obtained from the state estimator. A class of cyber-attacks called False Data Injection (FDI) attacks that target measurement data used for state estimation in the power grid are currently under study by the research community. These attacks modify sensor readings obtained from measuring equipment with the aim of misleading the control center into taking ill-advised response action. It has been shown that an attacker with knowledge of the network topology can craft an attack that bypasses existing bad data detection schemes (largely based on residual generation) employed in the power grid. We propose a multi-agent system that detects FDI attacks targeting state estimation in power grids. The multi-agent system is composed of software-implemented agents created for each substation. The software-based agents partition the power network into virtual sub-grids comprising adjacent substations connected by transmission lines and perform state estimation at each of these virtual sub-grids. Off-the-shelf computing and communication infrastructure is deployed in substations providing the software-based agents a means to communicate with each other. The agents facilitate exchange of meter readings among substations that are included in each sub-grid. We demonstrate that the information exchanged among substations, even untrusted, enables agents to cooperatively detect disparities between local state variables at the substation and global state variables computed by the state estimator. In the absence of FDI attacks, state estimation results at each sub-grid are identical to state estimation results for the whole grid.

However, in the presence of FDI attacks, compromised measurements can bypass bad data detection during state estimation for the whole grid. The virtual sub-grids and the main power grid have different topologies such that a successful FDI attack would have to evade bad data detection for the main grid and also for each of the virtual sub-grids making it extremely difficult for a false data injection attack to evade detection in a smart grid equipped with the multi-agent system proposed. To evaluate the proposed strategy for FDI detection, we conduct experiments with the IEEE 9-bus, 14-bus and 30- bus power system benchmarks using MATPOWER, an open source MATLAB toolbox. For each system, we generate 1000 attack cases that can bypass BDD schemes during state estimation for the whole grid. Then, we deploy agents that construct sub-grids at each substation in the power system and use state estimation at each sub-grid to analyze all FDI cases. In our experiments, we can detect all FDI attack cases with at least one agent.

15. No (Privacy) News is Good News: An Analysis of Privacy News in the U. S. and U. K. from 2010-2016

Nirav Ajmeri, Karthik Sheshadri, Jessica Staddon

News is a popular and influential source of privacy information and consequently, an important information source to analyze towards understanding privacy-related policy, product development and user perceptions. We provide the first large-scale text mining of privacy news using nearly 1700 articles from the New York Times and the Guardian over the years 2010-2016. Results of four independently-trained sentiment classifiers show that New York Times privacy news is predominantly negative in sentiment and more negative than randomly sampled news and some issues of global concern, such as the Syrian refugee crisis and the 2015 Paris terrorist attacks. We also provide a frequency analysis of entities involved (e.g. companies) over our entire data set and within certain privacy categories. Our analysis demonstrates that news text can be mined to facilitate both the tracking of important privacy events and efforts to build expressive taxonomies of privacy.

16. Obsidian: A Safer Blockchain Programming Language

Jonathan Aldrich, Michael Coblenz, Tyler Etzel, Eli Kanal, Brad Myers, Mark Sherman, Joshua Sunshine

Blockchain platforms, such as Ethereum, promise to facilitate transactions on a decentralized computing platform among parties that have not established trust. Recognition of the unique challenges of blockchain programming has inspired developers to create domain-specific languages, such as Solidity, for programming blockchain systems. Unfortunately, bugs in Solidity programs have recently been exploited to steal money. We propose a new programming language, Obsidian, that makes it easier for programmers to write correct programs.

17. On the Disconnect between CVSS Scores and Vulnerability Bounties

Andrew Meneely, Nuthan Munaiah

The Common Vulnerability Scoring System (CVSS) has been a standard for vulnerability severity assessment. The U.S. National Vulnerability Database (NVD) mandates CVSS base scoring of all vulnerabilities that it curates. The CVSS metrics, however, go unused when software vendors are looking to reward ethical hackers with monetary benefits ("bounties") for responsibly disclosing vulnerabilities. In this poster, we present the findings from an empirical evaluation of the relationship (or lack thereof) between the two conceptually related metrics: CVSS score and vulnerability bounty. We collected the CVSS scores and bounties for 851 vulnerabilities across 28 products. A simple Spearman's rank correlation analysis revealed a weak correlation (= 0.2799) between CVSS scores and vulnerability bounties. Hypothesizing that the weak correlation may be due to differences in concerns, we qualitatively compared the severity assessment criteria in CVSS to those in bounty determination guidelines. We found that 22 of the 34 bounty determination criteria are not explicitly mentioned in the CVSS specification. The results indicate that CVSS metrics lack the specificity of bounty determination guidelines, leading to a tenuous relationship between CVSS score and bounty.

18. Semantic Similarity in Security Regulations

Sarah Elder, Hui Guo, Munindar Singh, Laurie Williams

Security requirements for a product are often influenced by the federal, state, and local laws, organizational policies, and other regulations the product must comply with. Frequently, a single product must comply with multiple different regulatory documents which may or may not contain related statements. The goal of this research is to facilitate analysis of security regulations by automatically identifying relations between security regulations using natural language processing and machine learning.

19. Toward Effective Adoption of Security Practices

Nirav Ajmeri, Shams Al-Amin, Emily Berglund, Jon Doyle, Munindar Singh

Security tools guide developers to identify potential vulnerabilities in their codes. However, use of security tools is not very common. Sanctions are a way to enforce adoption of security practices. We address the research question of which sanctioning mechanism promotes adoption of security practices, and propose a simulation framework to explore sanctioning mechanisms for greater adoption.

20. Toward Normative Threat Models to Prevent Misuse

Özgür Kafali, Munindar Singh, Laurie Williams

Creating a set of comprehensive security requirements is the first step for implementing the necessary protocols to prevent, detect, and respond to misuses. Current approaches to tackle this task via threat modeling techniques lack formal semantics, which prevents formal understanding and aggregation of threat models needed to help security practitioners make informed decisions. We propose threat models formalized via norms, where a norm captures who is accountable to whom and for what. We discuss the merits of a normative representation and list promising directions of research.

21. Towards Privacy-Preserving Mobile Apps: A Balancing Act

Wing Lam, Dengfeng Li, Zhengkai Wu, Xusheng Xiao, Tao Xie, Wei Yang

As mobile apps are increasingly becoming data-driven, these apps tend to collect much app usage data to carry out their promised utilities and enhance user experiences. Unfortunately, some highly sensitive information in the data provides little or no benefit towards delivering the apps' utilities. For instance, for an app whose purpose is to show video game trailers, it is unnecessary to request and send its users' phone number and contact list to a remote server. There is a strong need for a framework to help protect users' app usage data while retaining the app's utility efficacy (e.g., the number of enabled features). There are three main challenges in realizing such a framework: First, it is difficult to correctly identify security-sensitive information in the app usage data. For instance, user input text can contain sensitive information, and the framework needs to understand the semantic meaning of such text in order to know whether sensitive information is present or not. Second, because utilities of apps vary dramatically, there is a need for generically applicable program analysis to measure the impact of information anonymization on the level of utility efficacy. Third, balancing privacy preservation and utility efficacy requires fine-grained analysis on privacy specification (such as a privacy policy declared by the app's developers) and the app. To address these challenges, we propose a privacy framework that enables a mobile app's developers to determine what sensitive information can be anonymized while maintaining a desirable level of utility efficacy.

22. User Interactions and Permission Use on Android

Jeffrey Foster, Nikolaos Kofinas, Michelle Mazurek, Kristopher Micinski, Rock Stevens, Daniel Votipka

Android has a permission system that asks users for authorization before an app uses sensitive resources such as contacts or GPS location. A key challenge in such authorization systems is balancing user interruptions with making sensitive resource use transparent. We hypothesize that Android's existing authorization systems (install-time permission lists or run-time dialog boxes, depending on the version) could achieve a better balance by integrating with the app's User Interface (UI) because the UI deeply informs the user's mental model of the app's behavior, including security-relevant behavior. In particular, in our work we ask whether user interactions (button clicks, page changes, dialog boxes, etc.) can be taken as evidence of authorization to use certain sensitive resources. If so, this could reduce the need for separate authorization requests. Conversely, we ask whether sensitive resource use without an associated interaction suggests a need for additional authorization requests. Note that while our studies are heavily influenced by Android, we believe our results will generalize to related mobile OS and similar settings like web apps. To answer these questions, we conducted two related studies. First, we reviewed 150 popular Android apps to determine whether sensitive resource uses are related to user interactions in existing apps. If so, an authorization mechanism integrated with the UI could work well with existing app designs. To carry out this study, we developed AppTracer, a dynamic analysis tool that instruments Android apps to log UI actions and resource uses, and then visualizes the logs as graphs that show temporal ordering of logged events. We used AppTracer to determine whether each observed resource use in each app was interactive, meaning either it was immediately preceded by a related UI event (e.g., accessing contacts after clicking a button marked "contacts"), or it was the main focus of the current screen (e.g., using location on a map screen). We found that, across our subject apps, several resources (microphone, camera, external storage, and calendar) are used almost exclusively interactively; several others (including bluetooth and phone state) are used mostly non-interactively (which we call in the background even if the app itself is on screen); and several resources (most notably contacts and location) exhibit a mix of interactive and background uses. These results suggest interactive and background uses may call for different authorization mechanisms, and that these mechanisms cannot necessarily be divided strictly by resource. These results informed the design of our second study, a 961 participant online survey investigating participants' expectations about interactive and background permission uses. Each participant viewed a slideshow of two usage scenarios for a mock mobile app, where each scenario shows a short interaction (e.g., launching the app, clicking a button, etc.) and then asks if the participant expects microphone, location, and/or contacts to be used after the interaction. We chose these resources to reflect a range of interactivity as measured in our app study. We aimed to gain insight into how different factors affect user expectations, and therefore which authorization mechanisms might be appropriate for different usage patterns. As we anticipated, we found that users are much more likely to expect resources to be accessed after a related interaction than in the background. However, we also found that seeing one interactive use does not prime the user to expect a future background use, indicating a potential weakness in the Android M request-on-first-use authorization model. In contrast, our findings show that an authorization request at launch does increase expectations for both interactive and background accesses, perhaps because it better conveys the idea that the resource could be accessed at any time. Drawing on the results of our studies, we make three design recommendations: First, resource uses should be made after associated interactions as much as possible. Given the current makeup of apps, this should be achievable for many commonly used resources without extensive effort. Second, separate authorization dialogs might be unnecessary for resources that are accessed mostly interactively. Finally, authorization for background resource uses should be distinct from authorization for interactive uses, and these background authorizations may be most effective when the app is launched.

23. What makes Air Force cyber warfare training hard?

Amber Bloomfield, Susan Campbell, Lelyn Saner

As part of the process of constructing an aptitude test to determine who would be the best fit for cyber warfare operator training in the United States Air Force (USAF), we conducted qualitative interviews with instructors and analyzed existing quantitative course success data. The goal was to determine what makes cyber warfare operator training hard and to see if existing tests could predict who would succeed. The results provide parameters for the goals and current baselines of the United States Air Force cyber training efforts.

Tutorials

Two tutorials were presented by speakers from NIST and Case Western Reserve University.

1. The Bugs Framework (BF) "Hands-On"

Irena Bojanova and Paul E. Black, National Institute of Standards and Technology

Advancements of scientific foundation in cybersecurity rely on the availability of accurate, precise, and non-ambiguous definitions of software weaknesses (bugs) and descriptions of software vulnerabilities. The Bugs Framework (BF) organizes software weaknesses into distinct classes, such as Buffer Overflow (BOF), Injection (INJ), Faulty Operation (FOP), and Control of Interaction Frequency (CIF). Each BF class has an accurate and precise definition and comprises the following: attributes that identify the software fault; causes that bring about the fault; consequences the fault could lead to; and sites in code where the fault might occur. Through a "hands-on" approach, the attendees were able to analyze definitions and (static) attributes of bugs' classes, along with their related dynamic properties, such as proximate, secondary and tertiary causes, consequences, and sites. The focus was on three of the developed BF classes, as well as on examples of applying the BF taxonomy to describe vulnerabilities such as Heartbleed and Ghost. The audience was involved in describing particular software vulnerabilities and the benefits of BF. The organizers are the BF Principal Investigators at the National Institute of Standards Technology and are proposing this tutorial as a way to help researchers and practitioners more accurately and quickly diagnose, describe, and measure security vulnerabilities.

2. System Monitoring for Security

Xusheng Xiao, Case Western Reserve University

Intrusive multi-step attacks, such as Advanced Persistent Threat (APT) attacks, have caused significant financial losses for many well-protected businesses. These advanced attacks are sophisticated and stealthy, and can remain undetected for years as individual attack steps may not be suspicious enough for early detection. To counter them, a recent trend is to leverage ubiquitous system monitoring for collecting the attack provenance for a long period of time and perform attack investigation for identifying risky system behaviors. This tutorial presented an introduction to popular techniques and tools of collecting system monitoring data, such as auditd, Sysinternals/ETW, and sysdig. The tutorial also addressed how dependency analysis can be applied on the system monitoring data for attack investigation, and how to perform data reduction techniques to scale up the analysis to monitor more hosts was displayed. Another technique that mines the patterns of the low-level system behaviors collected by system monitoring (e.g., a process reads a file) and uses these patterns to identify high-level software behaviors that security analysts are interested in (e.g., file compressions and ssh login) came next. Finally, there was a discussion of the current challenges of analyzing system monitoring data for security problems and future research directions.

SOS
SCIENCE OF SECURITY AND PRIVACY

OLD BAY
SEASONING
30% Less Sodium
For Seafood, Poultry, Salads, and Meats

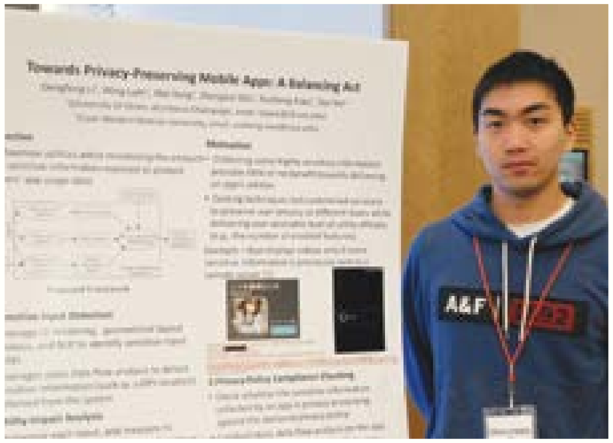
Towards Privacy-Preserving Mobile Apps: A Balancing Act
Motivation
A&F

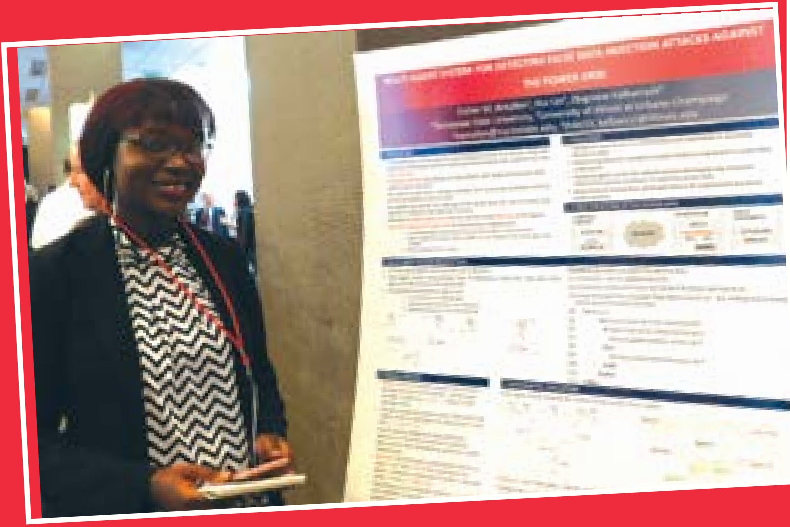
MULTI-AGENT SYSTEM FOR DETECTING FALSE DATA INJECTION ATTACKS AGAINST THE POWER GRID

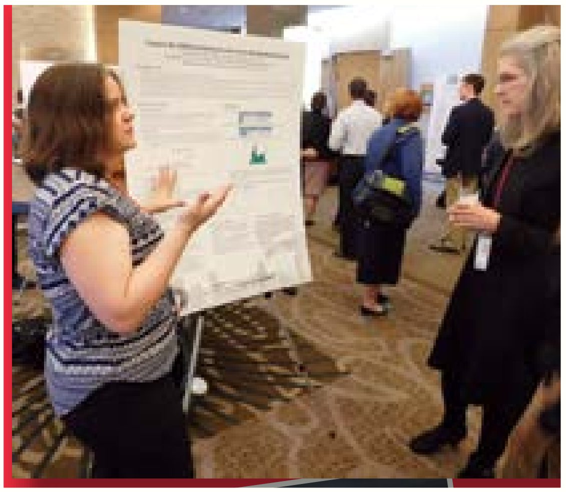

HOTSOS
HOT TOPICS in the SCIENCE OF SECURITY
SOS
SCIENCE OF SECURITY AND PRIVACY

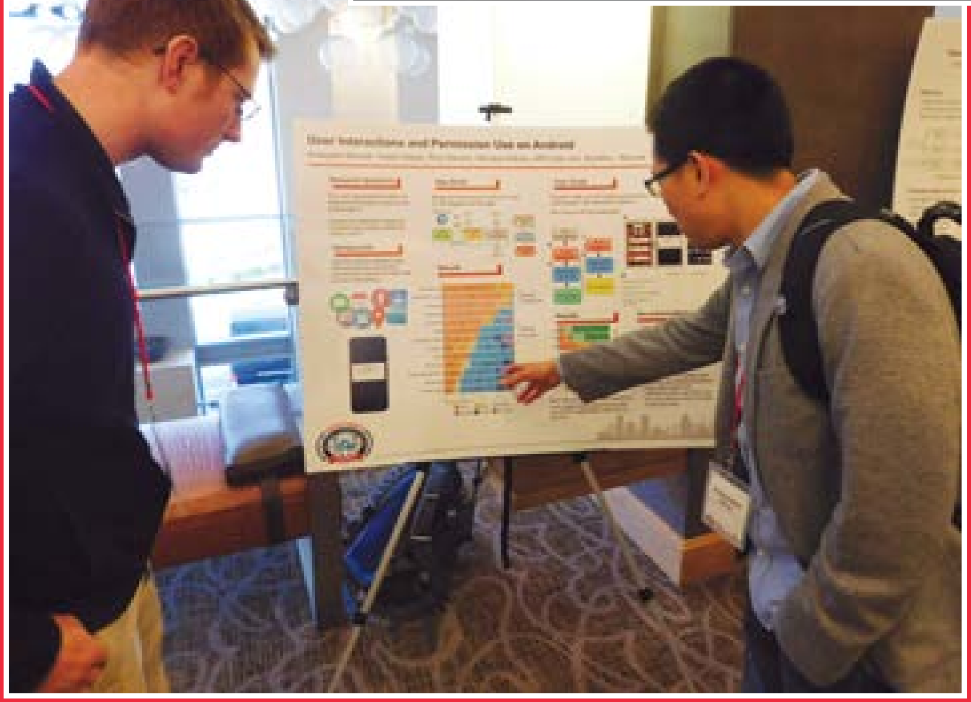
User Interactions and Permission Use on Android

The Science of Security Virtual Organization (SoS-VO)

The Science of Security Virtual Organization (SoS-VO) was established to provide a focal point for information about ongoing activities related to cybersecurity science and as a repository for significant research results. It emphasizes community development, information sharing, and interaction among researchers in the field. With over 1300 members and 6500 cybersecurity related items, the SoS-VO expands the SoS presence to universities, research centers, private companies, and government agencies worldwide. The SoS-VO provides a forum to discover resources, connect to others, and share and survey cybersecurity research. The goal of the SoS-VO is to help establish and support true collaboration in advancing cybersecurity science. In addition to hosting information on the SoS Lablets' and SURE research and coordinating Lablet research reporting to NSA, the SoS-VO enables its members to post research findings and publications done elsewhere. The site provides information on SoS activities, to include HotSoS, Lablet quarterly meetings, the annual Best Scientific Cybersecurity Paper Competition, and the International Science and Engineering Fair. The site also hosts chats, blogs, and forums, and provides information on upcoming events, position openings, calls for papers for conferences, and general cybersecurity news. New members are encouraged and can join by signing up via the SoS VO website at http://cps-vo.org/group/SoS

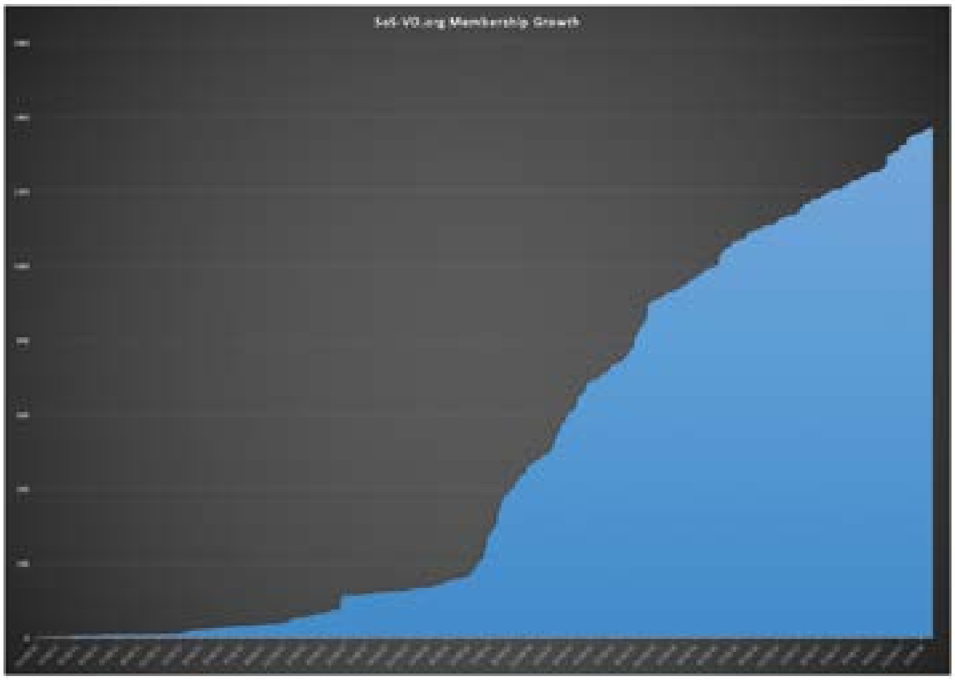

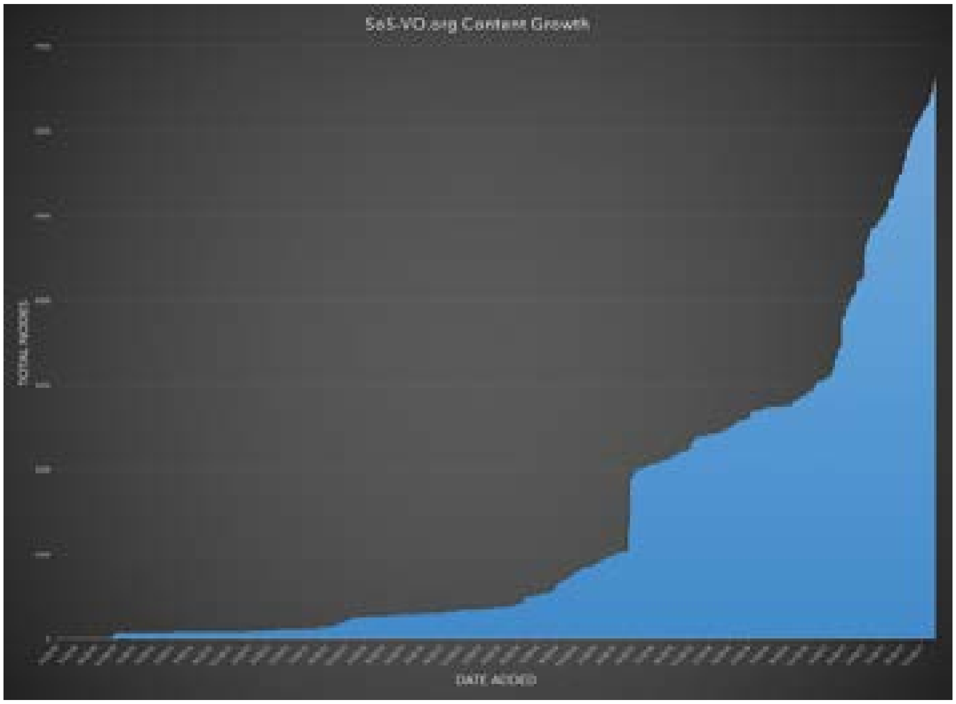

Outreach

In addition to the activities described elsewhere in this report, the Science of Security (SoS) initiative has employed multiple other means to expand awareness of Science of Security. Research Directorate (RD) SoS personnel give presentations on the Science of Security to representatives of other government agencies and at workshops and conferences. Researchers working at the Lablets and Sub-Lablets and collaborators from elsewhere in academia, government and industry serve as Science of Security ambassadors within their own organizations and at local and international symposia and conferences. One of the Science of Security Virtual Organization (SoS-VO) members, for example, was inspired to submit a Science of Security talk to RSA 2018 in response to the call for speakers.

Again in 2017, Science of Security personnel presented a talk at the Information Assurance Symposium which was held June 19-21 in Baltimore, MD. The 2017 Symposium was attended by representatives from across the U.S. Government, industry, and academia. Dr. Adam Tagert, the Science of Security Technical Director, gave a presentation entitled "Laying the Cyber Foundation: The Science of Security & Privacy". The briefing introduced SoS activities and engagement with academia, highlighted breakthrough research results, and described how to get involved.

Dr. Tagert also chaired a breakout session on Automating the Hard Parts of Security Measurement during the Maryland Cybersecurity Center (MC2) Workshop on Data-Driven Security held January 12-13 at the University of Maryland, College Park. The invitation-only workshop was aimed at researchers interested in studying security empirically, using data-driven techniques. The objective of the 2017 workshop was to identify the hard, open, problems in the field and to articulate tangible goals for advancing security through data-driven methods. Workshop participants, pictured above, were from universities, research institutes, government agencies, and think tanks from around the world.

The Science of Security Newsletter transitioned to SoS Reviews and Outreach (SoS R&O) in 2017. The SoS R&O is sent out monthly to nearly 1500 recipients in academia, government, and industry who were identified through their membership in the SoS Virtual Organization (SoS-VO), participation in an SoS-sponsored event, or a request to receive the mailing. 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. The sections of the R&O are 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.
- **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.
- **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 and addresses US and international policy issues, proposed regulations here and abroad, congressional inquiries and testimony, and indepth articles from non-technical publications.
- **Musings:** Brief articles on areas of concern or interest related to Science of Security

The initiatives noted above, combined with events such as HotSoS, the Annual Best Scientific Cybersecurity Paper Competition, SoS sponsorship of Intel ISEF awards, the SoS-VO, and Lablet activities, have expanded the SoS community to almost 200 institutions worldwide and an untold number of individual researchers. The 567 publications done by the Lablets since the inception of the Science of Security initiative have been presented and shared at research institutions, conferences, symposia, workshops, and government agencies around the world. Almost one fourth of the 59 cybersecurity analysis and design tools that have been developed by the Lablets have been released for use by other researchers. Lablet summer programs aimed at K-12 students and SoS engagement at the Intel ISEF provide opportunities to engage the next generation of researchers, while 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 disciplines. The SoS outreach efforts increase the likelihood that ad hoc and common practice approaches to security will be replaced by scientifically supported best practice methods established through rigorous research. By developing strategic rather than tactical methods of approaching cybersecurity, the practice of cybersecurity can be transformed to become efficient and proactive in both attack and defense.

Produced by Cyber Pack Ventures, Inc.