Annual Report

2015

VISION

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

Science of Security (SoS) Initiative

Annual Report

2015

Table of Contents

EXECUTIVE SUMMARY

The National Security Agency (NSA) Research Directorate sponsors the Science of Security (SoS) initiative for the promotion of a foundational cybersecurity science that is needed to mature the cybersecurity discipline and to underpin advances in cyber defense. Unlike the efforts undertaken by other parts of the government and industry, the SoS initiative focuses on foundational research that can fundamentally change the approach to cybersecurity, developing strategic rather than tactical approaches. The SoS initiative, established in 2012, has three goals: (1) engage the academic community for foundational research; (2) promote rigorous scientific principles; and (3) grow the SoS community. Over the past year, significant progress against each of the goals by expanding research based on the application of scientific principles and engaging academia, government, and industry in promoting the growth of Science of Security community.

When the Science of Security initiative was established, the three original SoS lablets, Carnegie Mellon Univeristy (CMU), North Carolina State University (NCSU), University of Illinois Urbana-Champaign (UIUC) worked with NSA to identify five "Hard Problems" that present significant technical challenges that will benefit from scientific research methods:

The Five Hard Problems are: (1) Scalability and Composability; (2) Policy-Governed Secure Collaboration; (3) Security Metrics Driven Evaluation, Design, Development, and Deployment; (4) Resilient Architectures; and (5) Understanding and Accounting for Human Behavior. In addition to 320 publications that were a direct result of hard problem research, there have been SoS hard problem research and written and released 310 research publications. These publications have tangible impacts on cybersecurity research and development including the following:

- Mathematical models have been developed to determine whether secure collaboration requirements, including priorities between them, are mutually consistent.
- Empirical studies have given evidence as to how the cyber threat landscape has changed following the introduction of various security technologies, whereas everything prior to this work was speculation.
- Before the lablet work, there were hundreds of disparate publications about intrusion-detection systems, each with varying methods and evaluation approaches. Our work has led to a taxonomy to compare those studies and to systematize that knowledge.
- Identified metrics that can predict vulnerabilities at the method level have been identified. Recent results from a large, open-source project show that the metrics increase just before a vulnerability is found, and decrease after a vulnerability is fixed, giving empirical evidence that new metrics are useful predictors of vulnerabilities.
- Lablet work has empirically demonstrated (using both the WINE dataset (Semantic Web Ontology) as well as network measurements on PKI (Public Key Infrastructure)revocation data following the Heartbleed incident) that software patches for known vulnerabilities are either not applied in a timely fashion, or are applied incorrectly. Prior to this work, it was unknown how quickly software patches were applied in the global environment,

or what techniques would be most beneficial for incentivizing faster patching.

- Before lablet work began, means to specify resiliency properties and requirements were not sufficiently precise or detailed to serve as a basis for rigorous systems engineering. We have developed a formal mathematical framework to enable more precise specification of the full range of properties of affordability, reliability, availability, safety, usability, scalability, evolvability, and resilience.

- Lablet work has discovered that a top-down strategy can be used to deploy policy enforcement across a network with greater efficiency and scalability than traditional, ingress-only deployments. This supports policy enforcement that can better absorb and adapt to adversarial traffic patterns.

- Lablet research has broken ground in developing practical mathematical frameworks that support reasoning about how robust a cyber-physical system might be to disruption. Previously, understanding was particularly lacking in how to approach reasoning about how an attack on the cyber component might be effected by manipulation of the physical component. Using this framework, we have developed algorithms that measure bounds on "how close" a physical disturbance can push a cyber-physical system near deleterious states.

After the addition of The Universityof Maryland, the four SoS Lablets and the 26 sub lablets have engaged more than 75 faculty and 50 graduate and post-doctoral students on more than 40 projects that have led to 166 publications. There have been over 320 relevant publications since the program was initiated.

In addition to the lablets, NSA funded a project on the System Science of SecUrity and REsilience for cyber-physical systems (SURE) to develop foundations and tools for designing, building, and assuring cyber-physical systems (CPS) that can maintain essential system properties in the presence of adversaries. Led by Vanderbilt University, SURE also involves researchers at Massachusetts Institute of Technology (MIT), University of California, Berkeley (UC Berkeley), and University of Hawaii (UH).

To promote rigorous research methods, the SoS initiative sponsored the third Competition for Best Scientific Cybersecurity Paper. This initiative recognized the best scientific cybersecurity paper published in 2014. The SoS initiative also awarded three research projects at the 2015 Intel International Science and Engineering Fair (ISEF), recognizing outstanding scientific accomplishments in cybersecurity at the high school level.

As well as lablet and SURE activities that have served to grow the community, the third annual Symposium and Bootcamp on the Science of Security (HotSoS) brought together researchers from numerous disciplines seeking a methodical, rigorous, scientific approach to identifying and removing cyber threats. The Science of Security (SoS) Newsletter, which provides cybersecurity news, upcoming events, and publications of interest, is read by Virtual Organization (VO) members and distributed to over 2,000 contributing authors and research faculty through direct mailings. Through all of the SoS activities, the SoS-VO has increased to over 900 members who are engaging in discussions, blogs and forums.

In FY16, SoS will continue lablet activities, SURE, SoS-VO, and HotSoS 2016, sponsor the 4th Annual Best Scientific Cybersecurity paper and ISEF, and continue to reach out to academia, government, and industry to grow the community and advance the Science of Security.

Details on progress against each of the hard problems can be found on page six Section I. Engage the Academic Community for Foundational Research.

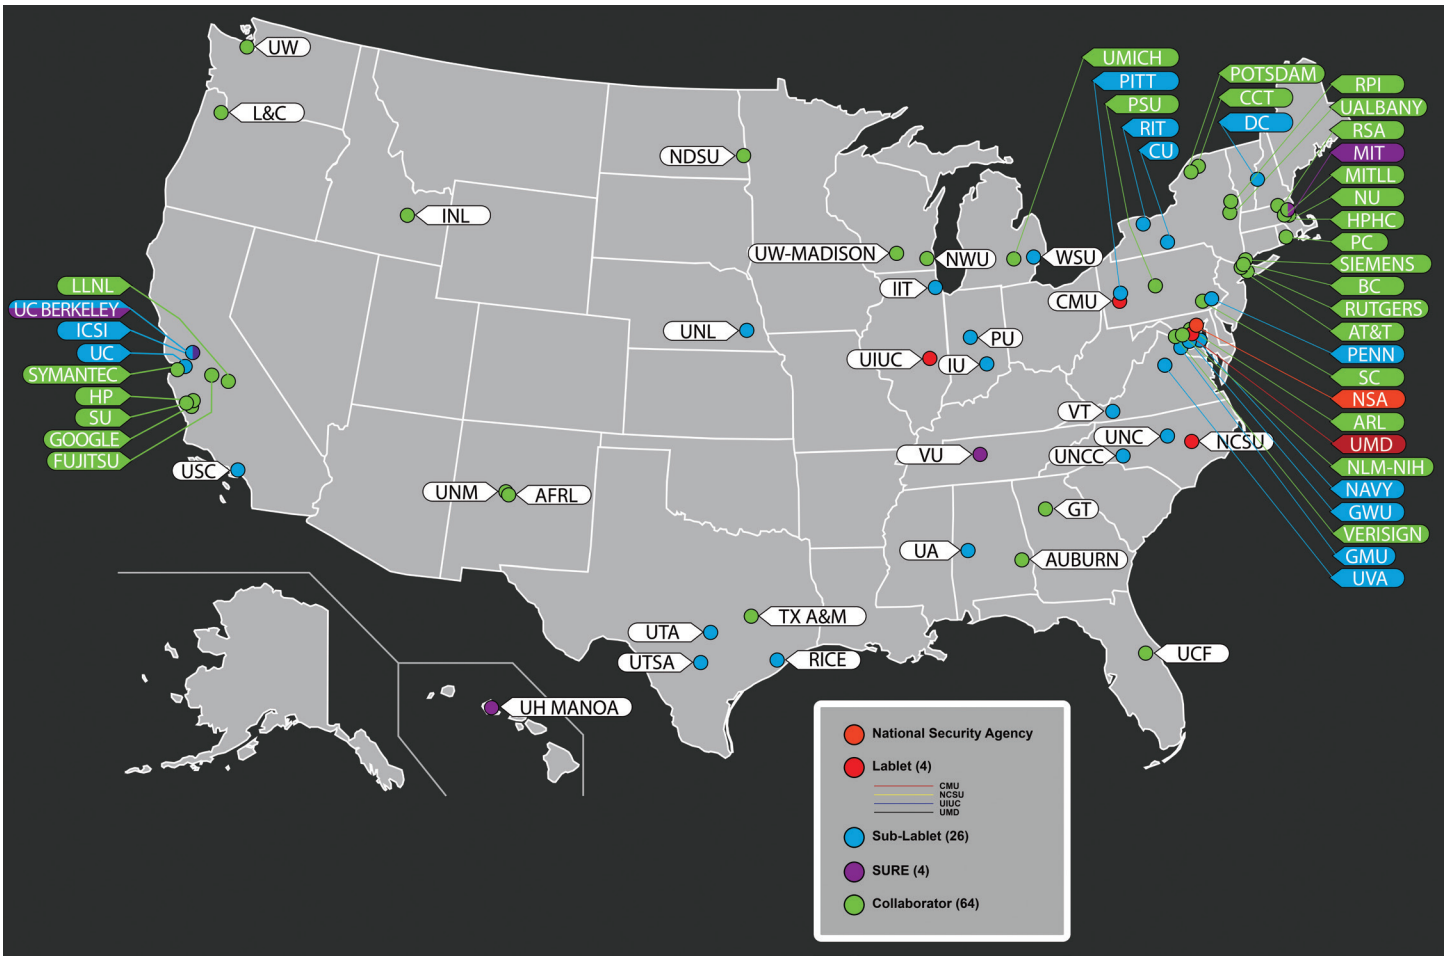

Impact of the Science of Security Program across academic, industry, and government organizations.

For more information about the Science of Security program, browse the SoS website at http://sos-vo.org

Section 1

Engage the Academic Community For Foundational Research

Science of Security

1110
0011
0001

LABLETS

Lablets, started in 2012, are dedicated to furthering Science of Security (SoS) goals for foundational research, enhancing the scientific rigor of cybersecurity, and growing the Science of Security community. The four lablets, Carnegie Mellon University (CMU), North Carolina State University (NCSU), University of Maryland (UMD), and University of Illinois at Urbana-Champaign (UIUC), work with a network of collaborating institutions, a total of 26 in all. The lablets meet together on a quarterly basis to present updates on current research against hard problems, exchange perspectives on progress in Science of Security, and strengthen the Science of Security community. Lablets also provide quarterly and annual reports on their foundational research in projects that help move security science forward, community outreach efforts to extend scientific rigor in the community and culture, educational activities to include changes to curriculum that indicate increased training or rigor in security research, and publications. Over the past 15 months (the lablets' 2014-15 contract year through June 2015), the lablets have published 120 papers associated with hard problems and the science of security. At the 2015 Summer Quarterly Lablet meeting, two panels were convened to address "What is Science of Security?" and "Is there a Science of Privacy?" both engaging in lively discussions. The Science of Security panelists noted that

Principal Investigators

PI's from the four lablets highlight their network of expanding relationships working on Science of Security activities.

it is misleading to compare this field to other fields because other fields are more restricted domains and more mature, while the SoS field is moving quickly. The expanding SoS community is making unique contributions with respect to empirical studies to close gaps between assumptions, models and how systems work against adversaries. Panelists noted that common lexicon and agreed models give us a chance to establish a basis for creating more fields, and describing a pathway to coherence by coalescing around hard problems so that we can come to conclusions about models. The Science of Privacy Panel noted that privacy has become more of a tangible issue and not only is more science needed to enable better decisions, but there needs to be a better understanding of human expectations and personal norms as well.

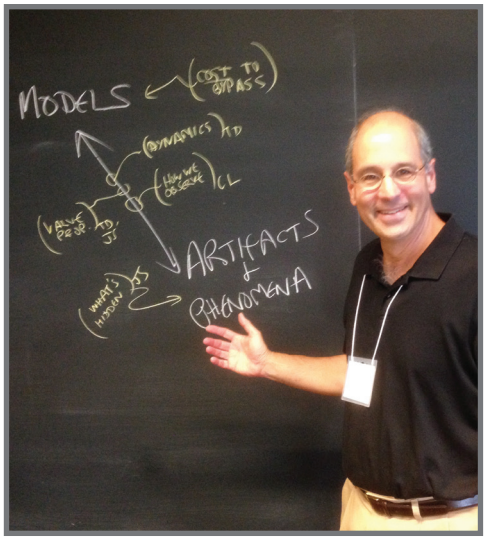

PI William Scherlis illustrates his thoughts on a model for Science of Security

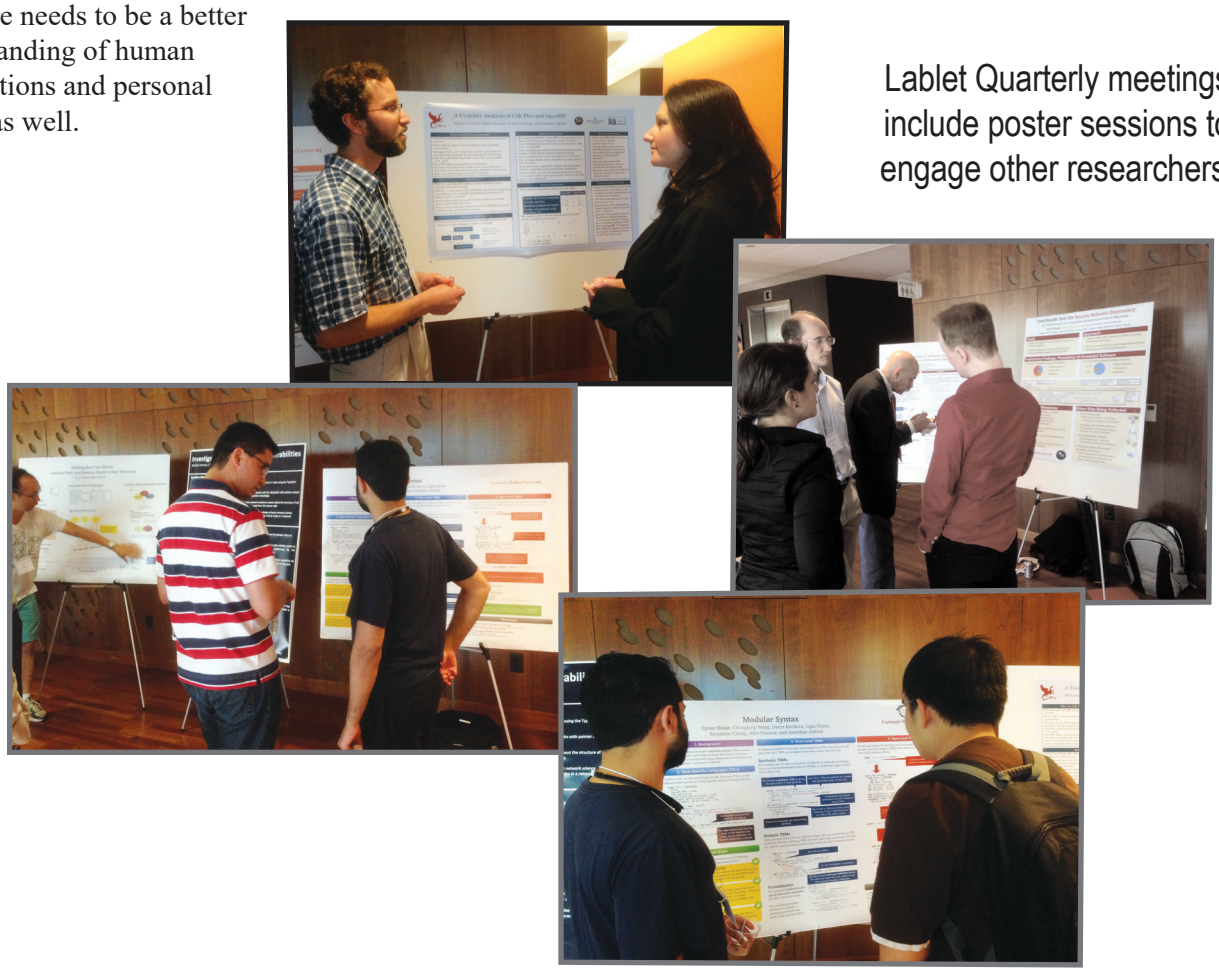

Lablet Quarterly meetings include poster sessions to engage other researchers.

Carnegie Mellon University

PI WILLIAM SCHERLIS

The Carnegie Mellon University (CMU) Science of Security (SoS) Lablet lead by Principal Investigator William Scherlis fosters scientific research in support of the cybersecurity mission of the National Security Agency. The CMU SoS Lablet has the mission to advance the state of cybersecurity research by focusing on the hardest technical problems, with emphasis on scale through *composability* of modeling and reasoning and on *human behavior* and *usability* for developers, evaluators, operators, and end users. The CMU Lablet also addresses SoS hard problems security metrics and policy-governed secure collaboration. Each of these five areas features a high level of technical challenge, a real opportunity for operational significance, and a significant likelihood of benefit from the synergetic Lablet approach.

The CMU SoS Lablet includes a set of interlinked technical projects that involve fifteen faculty from seven CMU academic departments, and seven other partner universities, including Cornell University, University of California at Irvine and Berkeley, University of Pittsburgh, Wayne State University, University of Nebraska, and University of Texas, San Antonio. Within CMU, seven departments and three colleges are involved. Eleven projects are underway and major papers have been published. Ten PhD students are supported, as well as four postdoctoral researchers, for the purpose of supporting the growth of researchers in SoS. The Lablet hosts the Security Behavior Observatory (SBO). The SBO is collecting behavioral data from users to identify missing knowledge about how people react or interact with security.

Carnegie Mellon University

FUNDAMENTAL RESEARCH

Race Vulnerability Study and Hybrid Race Detection

PIs –- Jonathan Aldrich (CMU) and Witawas Srisa-an (University of Nebraska, Lincoln)

Race conditions, since they are non-deterministic in nature, are notoriously hard to detect, yet they are a common cause of security vulnerabilities. In this project, we aim to improve the tradeoff between scalability and precision in race detectors. Existing race detectors suffer from either many false positives or unacceptably high overhead, which impedes their use in real world systems. Our hybrid race detection technique aims to be efficient and precise enough for practical large-scale applications.

In addition, we aim to better understand the human factors that cause race conditions. To do so, we are assembling a dataset of reproducible race vulnerabilities in real world programs and cataloging them to identify recurring problems. We are also running programming experiments to compare concurrency paradigms (e.g., the Cilk vs. OpenMP) to understand how these paradigms impact the security of the resulting code. If the study uncovers interesting relationships between races and security attacks, it can further contribute to security metrics research by providing another dimension of security assessment criteria. We also hope these studies will open up opportunities to mitigate race-related vulnerabilities.

A Language and Framework for Development of Secure Mobile Applications

PI –- Jonathan Aldrich

The goal of this project is to produce a language and framework that enables the construction of secure mobile applications with known security properties. The framework-based approach requires composition between the framework and its plugins, with modular type system and analysis techniques that avoid reanalyzing each individual component.

For example, providing programmers with mechanisms for constructing commands that are as convenient as strings while being as secure as prepared SQL statements is a promising mechanism for eliminating injection vulnerabilities. Mechanisms for embedding domain-specific languages (DSLs) for constructing commands within a programming language exist, but none have achieved widespread use, in part because prior techniques were unmodular, so that separately-defined embedded DSLs could not be used together. In this project we address these limitations using a novel mechanism called type-specific languages that supports modular DSL embeddings by associating a unique DSL with appropriate types.

Key issues include scalability (e.g., between client, server, and other distributed resources), tool usability, and security attributes such as confidentiality, integrity, and mitigation of common vulnerabilities seen in present technologies. Human behavior is also secondarily relevant, both because human limitations are the cause of many of the security vulnerabilities that we hope to eliminate, but also because an understanding of the way humans develop software is key to ensuring that our approach is usable and enhances (rather than detracts from) developer productivity.

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

PI –- Travis Breaux

When software developers reuse existing libraries, frameworks, platforms or services, they often cannot assess the security requirements satisfied by those third-party programs. We developed techniques for expressing data flow requirements and checking whether these flows preserve the use and collection limitation principles from international privacy standards. The approach is based on Description Logic and supports verification across requirements compositions, which occur when mobile apps are run on mobile devices or other platforms, or when software employs remote services for storage, authentication, etc. The results will appear in the proceedings of the 23rd IEEE International Requirements Engineering Conference, and the source code and a demonstration have been made available online.

Security is based on defense-in-depth, which requires the satisfaction of multiple security requirements to mitigate security threats. However, the prevailing mechanism for information assurance is based on checklists, which assume a single threat context (often the union of all possible threats and mitigations). We developed a technique to measure the impact of composing security requirements on perceived security risk in the presence of changing threats. The technique is based on multi-level modeling, a statistical technique that can measure the correlation of multiple factors on a dependent variable, such as risk perception. As we scale this approach to an increasing

number of threats, the results can be used to propose a minimal set of security requirements likely believed to mitigate the most relevant threats. The results will appear in the proceedings of the 23rd IEEE International Requirements Engineering Conference.

Codifying privacy regulations to be consistent with their high-level intuition can be difficult for security architects or application domain experts of information systems. We provide a formal framework to use UML sequence diagrams as a practical means to graphically express privacy regulations and policies to enable domain experts to verify and confirm the codification is valid. Once intuitively confirmed, our framework introduces an algorithmic approach to formalizing the semantics of sequence diagrams in terms of linear temporal logic (LTL) templates. In all the templates, different semantic aspects are expressed as separate, yet simple LTL formulas that can be composed to define the complex semantics of sequence diagrams. We leverage the analytical powers of automated decision procedures for LTL formulas to determine if a collection of sequence diagrams is consistent, and independent, and also to verify if a system design conforms to the privacy policies. This work will appear in IEEE Transactions on Dependable and Secure Computing in 2015.

USE: User Security Behavior

PI –- Lorrie Cranor

The Security Behavior Observatory addresses the hard problem of "Understanding and Accounting for Human Behavior" by collecting data directly from paid volunteers' home computers, thereby capturing people's computing behavior "in the wild." This data is the closest to the ground truth of the users' everyday security and privacy challenges that the research community has ever collected. We expect the insights discovered by analyzing this data will profoundly impact multiple research domains, including but not limited to, behavioral sciences, computer security and privacy, economics, and human-computer interaction. The Security Behavior Observatory will also address the "Predictive Security Metrics" hard problem in two ways: first, by evaluating established metrics for measuring their user-reported security behaviors by comparing them to users' actual computing behavior; and second, by proposing new predictive metrics on user behavior with respect to computer security and privacy founded on highly ecologically-valid data analyses.

Secure Composition of Systems and Policies

PI(s) –- Anupam Datta, Limin Jia

Interface-confinement is a common mechanism that secures untrusted code by executing it inside a sandbox. The sandbox limits (confines) the code's interaction with key system resources to a restricted set of interfaces. This practice is seen in web browsers, hypervisors, and other security-critical systems. Motivated by these systems, we develop a program logic, called System M, for modeling and proving safety properties of systems that execute adversary-supplied code via interface-confinement.

In addition to using computation types to specify effects of computations, System M includes a novel invariant type to specify the properties of interface-confined code. System M supports compositional proof–-security proofs of sequentially composed programs are built from proofs of their sub-programs. System M also admits concurrent composition–-properties proved of a program hold when that program executes concurrently with other, even adversarial, programs. System M can be used to model and verify protocols as well as system designs.

We demonstrate the reasoning principles of System M by verifying the state integrity property of the design of Memoir, a previously proposed trusted computing system.

Science of Secure Frameworks

PI(s) –- David Garlan, Bradley Schmerl

In this project, we are building a scientific basis for the security of framework-based applications. Software frameworks, such as Android, are used ubiquitously in modern applications because they offer a unique means for achieving composition and reuse at scale. Achieving security in framework-based applications can be challenging because of the close coupling between a framework and its plugins: Plugin developers must understand and obey the constraints in the framework's security model, which can be quite complex, in order to achieve security of the resulting application. We are investigating a combination of static and dynamic checking of framework security rules that can be used to provide secure frameworks that still maintain the flexibility to allow extensive plugins.

Our work in this project is focused on Android, along the following thrusts: (1) use of static analysis and model checking techniques to identify potential vulnerabilities that a set of apps installed on a device (we call this an app ecosystem) may exhibit; (2) architectural modeling and analysis of the app ecosystem to augment vulnerability checking; (3) use of architectural models at run time to guide run time analysis

and mitigation of these vulnerabilities. Furthermore, we are (4) applying sandboxing techniques at the framework level to provide better protection on mobile platforms.

This project is being done in conjunction with our subcontractors: Prof. Sam Malek at George Mason University and Prof. Marwan Abi Antoun at Wayne State University.

Epistemic Models for Security

PI –- Robert Harper

We are using a combination of methods to achieve modular analysis of the security properties of concurrent imperative programs. We use a type system based on the lax modality to express and enforce information flow constraints. We use a novel linear epistemic logic to analyze the execution traces of programs to derive which principals learn which secrets in a given run. We prove that for a well-typed program, a low-security principal can learn a high-security secret only if it is explicitly authorized to do so by an integrated authorization logic.

All of these methods, being grounded in logic and type theory, are inherently compositional in nature.

Multi-Model Run-Time Security Analysis

PI –- Juergen Pfeffer

Our research focuses on creating the scientific foundations to support model-based, run-time diagnosis and the repair of security attacks. Specifically, our research develops models that (a) scale gracefully with the size of system and have appropriate real-time characteristics for run-time use, and (b) support composition through multi-model analysis. In this research, we develop a rigorous basis for composing architectural models with organizational network models to provide much richer capabilities than is available from either in isolation. The following hard problems are addressed in this project. Composability through multiple semantic models (here, architectural, organizational, and behavioral), which provide separation of concerns, while supporting synergistic benefits through integrated analyses. Scalability to large complex distributed systems using architectural models. Resilient architectures through the use of adaptive models that can be used at run-time to predict, detect, and repair security attacks. Predictive security metrics by adapting social network-based metrics to the problem of architecture-level anomaly detection.

We have developed an approach to detect insider threats by clustering the paths on the architecture graph so generated. We represent the entire activity log over an underlying software system as a graph. For every user, a sequence of observed activities becomes a path on the architecture graph. We developed a clustering approach to cluster these paths. Anomalous paths are further investigated by incorporating organizational context (e.g., role of user). The clustering method is built on a generative model to generate sequences. A core challenge of investigating insider threats is the availability of datasets. In this project, a simulator for web systems has been developed for generating data that include malicious behavior. Furthermore, real-world datasets, including one from Los Alamos, one from CERT, and one from the “Vegas” lab have been evaluated for their suitability for insider attack research.

Highly Configurable Systems

PI –- Juergen Pfeffer

This project complements the Science of Security endeavor with a focus on the often overlooked problems of configuration options in systems. Whereas current approaches work on specific snapshots and require expensive recertification, our approaches extend underlying mathematical models (data-dependence graphs) with configuration knowledge and will thus scale analyses and reduce the need for repeating analyses. Furthermore, we explore whether configuration complexity and configuration-specific program-dependence is a suitable empirical predictor for the likelihood and severity of vulnerabilities in complex systems. Finally, technical and empirical results of our work will also bring new approaches to the field of social network analysis that can be very powerful and applicable for Science of Security far beyond the scope of the current Lablet. The SoS hard problems are addressed as follows. Scalability and Composability: Isolating configuration options or controlling their interactions will lead us toward composable analysis with regard to configuration options. Predictive security metrics: To what degree can configuration-related indicate implementations that are more prone to vulnerabilities or in which vulnerabilities have more severe consequences?

In this project we finished implementing and testing a tool to extract precise call graphs with function pointers for product lines/compile-time variability. This overcomes a key limitation of previous approaches, which are inaccurate due to their lack of pointer analysis and allow for more precise composability analysis. The tools will be available as part of the next version of TypeChef. Using this infrastructure we extracted a variety of network models based on files, functions, and features from

Linux and other systems (e.g., Busybox) and started with assessment of reasonability of selected network metrics for analyzing software systems. We also compiled lists of known Linux vulnerabilities (CVEs) and started to correlate them with results from graph analysis and other metrics. The goal is to identify whether certain characteristics, e.g., position in the call graph, increase the chance for a function to be a security risk.

Security Reasoning for Distributed Systems with Uncertainty

PI –- André Platzer

This project has made two contributions that will help reasoning about security goals in uncertain systems. First, we described a new problem class called #E-SAT (Zawadzki, Platzer, & Gordon, A Generalization of SAT and #SAT for Robust Policy Evaluation, 2013), which is a quantified generalization of SAT, one of the most important problems in computer science. Our extension includes both counting (#) and search (E) quantifiers. We demonstrated that a number of questions about the robustness and reliability of policies can be set up as #E-SAT instances. We also gave a rigorous theoretical characterization of the problem's worst-case complexity. Furthermore, we designed an algorithm that, despite the problem class's formidable worst-case complexity, performed extremely well empirically. The empirical success was due to exploiting a type of structure that seems common in practice.

Our second contribution is ongoing work on a set of approximate optimization techniques. 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 currently focused on applying these methods to anomaly detection problems, but hope to extend our work to policy synthesis questions shortly since the solver technique is general. 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). The LCP subsumes a broad class of optimization problems that includes quadratic programming. We have already theoretically investigated this method; implemented five different solvers based on these techniques, and are currently evaluating the method empirically. We are experimenting using anomaly detection LCP instances that represent one-class kernel support vector machines (SVMs). One-class SVMs are used in practice to perform anomaly detection, such as in phishing detection. Additionally, we have a new result that characterizes the class of functional approximation methods that are compatible with fast, iterative methods for the LCP. What makes the LCP approach particularly interesting is that it provides general approximation techniques

- Panelist: Jonathan Aldrich, "Language Composition," DSLDI 2015 Workshop in association with the ECOOP 2015 Conference.
- Invited Talk: Jonathan Aldrich, "Safely Composable Type-Specific Languages," and "Structuring Documentation to Support State Search: A Laboratory Experiment about Protocol Programming," IFIP Working Group on Language Design, April 2015.

Geo-Temporal Characterization of Security Threats

PI –- Kathleen Carley

This project aims to develop a global characterization of cyber-security threats. Using data from Symantec, and other indicators at the nation-state level, a threatened and threatening profile per country is produced. Questions addressed include, but are not limited to, which countries are most vulnerable to which types of threats? Are cyber threats following traditional lines of hostilities at the global level, or are new threats emerging? Social network and statistical techniques are used to assess the overall threat profile and theoretical results about error bounds that hold for general classes of LCPs (not just SVMs) and so can apply to a broad range of security problems stemming from policy optimization, classification, and anomaly detection.

COMMUNITY ENGAGEMENTS

- Invited Talk: Jonathan Aldrich, "Searching the State Space: A Qualitative Study of API Protocol Usability," IFIP Working Group on Language Design, June 2014.
- Workshop Organizer: Joshua Sunshine, Co-organized the 2014 Workshop on Evaluation and Usability of Programming Languages and Tools (PLATEAU) co-located with OOPSLA SPLASH. The workshop focused on applying the research techniques from two lablet-funded papers ("Structuring Documentation to Support State Search: A Laboratory Experiment about Protocol Programming," and "Searching the State Space: A Qualitative Study of API Protocol Usability") to further programming languages and language features. The workshop was attended by 34 researchers from around the world.

- Invited Talk: Travis D. Breaux, "Hermeneutics of Information Privacy," at the IFIP Working Group 2.9 on Requirements Engineering, Cozumel, Mexico, February 18, 2015.
- Keynote Talk: Travis D. Breaux, "Going Native - Relying on Pidgins and Creoles to Construct High Confidence Software," High Confidence Software and Systems Conference, Annapolis, MD, May 6, 2014.
- Keynote Talk: Travis D. Breaux, "Verifying Data Protection Rules in Complex Data Ecosystems," NSA/CSS Mission Compliance Conference, Baltimore, May 7, 2014.
- Panelist: Travis D. Breaux, "Overview of Privacy Engineering Approaches," NIST Privacy Engineering Workshop, Gaithersburg, MD, April 9, 2014.
- Invited Talk: Arbob Ahmad, 2015 SoS Lablet Summer Workshop, Carnegie Mellon University, Pittsburgh, PA, July 14-15, 2015.
- Distinguished Talk: Lorrie Cranor, "Security, Privacy, and Human Behavior," Women in Cybersecurity, Atlanta, GA, March 27, 2015.
- Keynote Talk: Alessandro Acquisti, Closing Workshop: "You Are Not Alone," SPION (Security and Privacy in Online Social Networks), Leuven, December 2014.
- Invited Talk: Sam Malek, "A Tool for Automated Detection of Inter-Application Security Vulnerabilities in Android," National Security Agency, College Park, Maryland, March 2015.
- International Workshop on Software Development Lifecycle for Mobile. Hong Kong, China, November 2014.
- Keynote Talk: Bradley Schmerl, "Challenges in Engineering Dependable Self-Adaptive System. International Workshop on Recent Advances in the Dependability Assessment of Complex systEms (RADIANCE)," Rio de Janeiro, Brazil, June 22, 2015.
- Invited Talk: Bradley Schmerl, "Reasoning about Human Involvement in Self-Adaptive Systems," University of Campinas, Campinas, Brazil, June 26, 2015.
- Invited Talk: David Garlan, "Identifying and resolving consistency issues between model representations," NASA Jet Propulsion Lab, July 2015.
- Invited Talk: David Garlan, "Self-Adaptive Systems," 2nd Latin-American School on Software Engineering (ELA-ES 2015), Porto Alegre, Brazil, June 2015.
- Keynote Talk: David Garlan, "Modeling Challenges for Cyber-Physical Systems," The International Workshop on Software Engineering for Smart Cyber-Physical Systems (SEsCPS), Florence, Italy, May 17, 2015.
- Invited Talk: David Garlan, "Software Architecture: A Travelogue," The Future of Software Engineering, Hyderabad, India, May 2014.
- Invited Talk: Limin Jia, "Design, Implementation, and Verification of XMHF," WiCys (Women in Cyber security), Nashville, TN, March 2014.
- Invited Talk: Limin Jia, "Proving Trace Properties of Programs that Execute Adversary-Supplied Code," INRIA Prosecco Seminar, INRIA, Paris, France, June 2014.
- Invited Talk: Limin Jia, "Proving Systems Secure Against Adversaries," Lehigh University, Lehigh, PA, Oct. 2014.
- Keynote Talk: Kathleen M. Carley, "Twitter, Terror and Terms: Network Analytics for Assessing Large Scale Media Data," PolNet, Portland, OR, June 2015.
- Keynote Talk: Kathleen M. Carley, "Social Media Analytics Using Dynamic Network Methodologies," DHS CCICADA, Pittsburgh, PA, March 2015.
- Invited Talk: Kathleen M. Carley, "Dynamic Network Analysis & Global Mapping," National Research Council and NGA, Washington DC, April 2015.
- Invited Talk: Kathleen M. Carley, "Cyber Security & Dynamic Network Methodologies," CERT, Pittsburgh, PA, April 2015.
- Invited Talk: Kathleen M. Carley, "Characterizing Insider Threats Using Network Analytics," Science of Security Lablet Meeting, Carnegie Mellon University, Pittsburgh, PA, July 14-15, 2015.
- Invited Talk: Kathleen M. Carley, "DNA, Networks and Simulation," CASOS Summer Institute, Carnegie Mellon University — training to approximately 25 attendees from around the world, July 5-12, 2015.
- Invited Talk: Kathleen M. Carley, "Dynamic-Network Analysis and *ORA," INSNA Sunbelt XXXV,

EDUCATIONAL

Brighton, UK — training to approximately 15 attendees from around the world, June 23-28, 2015.

- Invited Talk: Kathleen M. Carley, "Social Network Analysis for Science of Security," HotSOS, UIUC, Urbana-Champaigne, IL — dydactic seminar to approximately 35 people at HotSoS, April 21-22, 2015.
- Jonathan Aldrich. "Secure Coding," Course Number 14-735, Spring 2015.
- Juergen Pfeffer. "Introduction to Network Science," Course Number(s) 08-622/08-302, Fall 2014.
- Lorrie Cranor. "Usable Privacy and Security," Course Number(s): 05-436 / 05-836 / 08-534 / 08-734, Spring 2015.
- Co-taught by David Garlan and Bradley Schmerl. "Self-Adaptive Systems," Course Number 17-707, Spring 2015.
- Anupam Datta. "Secure Software Systems, Building Verifiable Systems," Course Number 18-732, Spring 2015
- Kathleen M. Carley. "New Module was Developed for Dynamic Network Analysis," Course Number(s) 08-801/08-640, Spring 2015
- Co-taught Kathleen M. Carley, Travis Breaux, and Lorrie Cranor. "Special Practicum Project based on insider espionage developed based on SoS research," Course Number 08-999, Fall 2014

Carnegie Mellon University

PUBLICATIONS

- J. Shahen, J. Niu, M. Tripunitara. "Mohawk+T: Efficient Analysis of Administrative Temporal Role-Based Access Control (ATRBAC) Policies," 20th ACM Symposium on Access Control Models and Technologies (SACMAT), pp. 15-26. Vienna, Austria, June 1 – 3, 2015. Presented by Jonathan Shahen.
- H. Hibshi, T. Breaux, M. Riaz, L. Williams. "Discovering Decision-Making Patterns for Security Novices and Experts." TR CMU-ISR-15-101, March 2015.
- Hui Shen, Ram Krishnan, Rocky Slavin, and Jianwei Niu. "Sequence Diagram Aided Privacy Policy Specification," IEEE Transactions on Dependable and Secure Computing, Issue 99, December 19, 2014.
- H. Hibshi, T. Breaux, M. Riaz, L. Williams. "A Framework to Measure Experts' Decision Making in Security Requirements Analysis," IEEE 1st International Workshop on Evolving Security and Privacy Requirements Engineering, pp. 13-18, Karlskrona, Sweden, August 25 - 29, 2014. Presented by Hanan Hibshi.
- R. Slavin, J.-M. Lehker, J. Niu, T. Breaux. "Managing Security Requirement Patterns Using Feature Diagram Hierarchies," IEEE 22nd International Requirements Engineering Conference, pp. 193-202, Karlskrona, Sweden, August 25 – 29, 2014. Presented by Rocky Slavin.
- Rao, H. Hibshi, T. Breaux, J-M. Lehker, J. Niu, "Less is More? Investigating the Role of Examples in Security Studies using Analogical Transfer," 2014 Symposium and Bootcamp on the Science of Security (HotSoS), Article 7. Raleigh, NC, April 8 -9, 2014. Presented by Hanan Hibshi.
- Hamid Bagheri, Eunsuk Kang, Sam Malek, and Daniel Jackson. "Detection of Design Flaws in the Android Permission Protocol through Bounded Verification," Proceedings of the 20th International Symposium on Formal Methods (FM 2015), pp. 73-89, Oslo, Norway, June 24 – 26, 2015. Presented by Eunsuk Kang.

- Javier Cámara, Gabriel A. Moreno and David Garlan. “Reasoning about Human Participation in Self-Adaptive Systems,” Proceedings of the 10th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS 2015), Florence, Italy, May 18 – 19, 2015. Presented by Javier Cámara.

- Alireza Sadeghi, Hamid Bagheri, and Sam Malek. “Analysis of Android Inter-App Security Vulnerabilities Using COVERT,” 37th International Conference on Software Engineering (ICSE), Tool Demo Track, Florence, Italy, May 16 – 24, 2015. Presented by Alireza Sadeghi.

- Bradley Schmerl, Jeff Gennari, David Garlan. “An Architecture Style for Android Security Analysis,” (Poster), Symposium and Bootcamp on the Science of Security (HotSoS), Urbana IL, April 21-22, 2015. No presenter.

- Khalaj, E., Wang, Y., Giang, A., Abi-Antoun, M., and Rajlich, V. “Impact Analysis Based on a Global Hierarchical Object Graph,” 22nd IEEE International Conference on Software Analysis, Evolution, and Reengineering (SANER), pp. 221-23. Montreal, Canada, March 2 – 6, 2015. Presented by Marwan Abi-Antoun.

- Hamid Bagheri, Alireza Sadeghi, Joshua Garcia, and Sam Malek. “COVERT: Compositional Analysis of Android Inter-App Security Vulnerabilities,” Technical Report GMU-CS-TR-2015-1. Accepted to appear in the IEEE Transactions on Software Engineering.

- Vanciu, R., Khalaj, E. and Abi-Antoun, M. “Comparative Evaluation of Architectural and Code-Level Approaches for Finding Security Vulnerabilities,” Workshop on Security Information Workers, co-located with ACM Conference on Computer and Communications Security (CCS), pp. 27-34, Scottsdale, AZ, November 3 – 7, 2014. Presented by E. Khalaj.

- Riyadh Mahmood, Nariman Mirzaei, and Sam Malek. “EvoDroid: Segmented Evolutionary Testing of Android Apps,” Proceedings of the 22th ACM SIGSOFT International Symposium on the Foundations of Software Engineering (FSE 2014), pp. 599-609. Hong Kong, November 16 – 22, 2014. Presented by Riyadh Mahmood.

- Marwan Abi-Antoun, Sumukhi Chandrashekar, Radu Vanciu, and Andrew Giang. “Are Object Graphs Extracted Using Abstract Interpretation Significantly Different from the Code?,” IEEE International Working Conference on Source Code Analysis and Manipulation, pp. 245-54. Victoria, British Columbia, Canada, September, 28 – 29, 2014. Presented by Marwan Abi-Antoun.

- Marwan Abi-Antoun, Sumukhi Chandrashekar, Radu Vanciu, and Andrew Giang. “Are Object Graphs Extracted Using Abstract Interpretation Significantly Different from the Code?,” (Extended Version). Technical report, Wayne State University, September 2014.

- Javier Cámara, Antonia Lopes, David Garlan and Bradley Schmerl. “Impact Models for Architecture-Based Self-Adaptive Systems,” Proceedings of the 11th International Symposium on Formal Aspects of Component Software (FACS2014), pp. 89-107. Bertinoro, Italy, September 10 – 12, 2014. Presented by Javier Cámara.

- Jonathan Aldrich, Cyrus Omar, Alex Potanin, and Du Li. “Language-Based Architectural Control,” 6th International Workshop on Aliasing, Capabilities, and Ownership (IWACO), Uppsala, Sweden, July 28 – August 1, 2014. Presented by Jonathan Aldrich.

- Khalaj, E., Vanciu, R., and Abi-Antoun, M. “Is There Value in Reasoning about Security at the Architectural Level: a Comparative Evaluation?,” Poster at Symposium and Bootcamp on the Science of Security (HotSoS), Article 30, Raleigh, NC, April 8 – 9, 2014. Presented by E. Khalaj.

- Khalil Ghorbal, Jean-Baptiste Jeannin, Erik W. Zawadzki, André Platzer, Geoffrey J. Gordon, and Peter Capell. “Hybrid Theorem Proving of Aerospace Systems: Applications and Challenges,” Journal of Aerospace Information Systems. Vol 11:10, pp. 702 – 13, October 2014.

- Forget, Alain, Komanduri, Saranga, Acquisti, Alessandro, Christin, Nicolas, Cranor, Lorrie, Telang, Rahul. 2014. “Building the Security Behavior Observatory: An Infrastructure for Long-term Monitoring of Client Machines,” IEEE Symposium

and Bootcamp on the Science of Security (HotSoS) 2014. Raleigh, NC, August 8 – 9, 2014. Presented by Alain Forget.

- Forget, S. Komanduri, A. Acquisti, N. Christin, L.F. Cranor, R. Telang. "Security Behavior Observatory: Infrastructure for Long-term Monitoring of Client Machines," Carnegie Mellon University CyLab Technical Report CMU-CyLab-14-009. July 14, 2014. Presented by Alain Forget.

- S. Zhou, J. Al-Kofahi, T. Nguyen, C. Kästner, and S. Nadi. "Extracting Configuration Knowledge from Build Files with Symbolic Analysis," Proceedings of the 3rd International Workshop on Release Engineering (Releng), New York, NY: ACM Press. Florence, Italy, May 19, 2015. Presented by Shurui Zhou.

- Ferreira , Gabriel & Kästner, Christian & Pfeffer, Jürgen & Apel, Sven. "Characterizing Configuration Complexity in Highly-Configurable Systems with Variational Call Graphs," HotSoS 2015 Symposium and Bootcamp on the Science of Security, Article 17, Urbana-Champaign, IL, April 21-22, 2015. Presented by Gabriel Ferreira.

- C. Hunsen, J. Siegmund, O. Lessenich, S. Apel, B. Zhang, C. Kästner, and M. Becker. "Preprocessor-Based Variability in Open-Source and Industrial Software Systems: An Empirical Study," Empirical Software Engineering (ESE), Special Issue on Empirical Evidence on Software Product Line Engineering, Springer Science+Business Media, New York, April 14, 2015. Presented by Claus Hunsen.

- S. Nadi, T. Berger, C. Kästner, and K. Czarnecki. "Where do Configuration Constraints Stem From? An Extraction Approach and an Empirical Study," IEEE Transactions on Software Engineering (TSE), Issue 99. March 23, 2015. Presented by Sarah Nadi.

- Max Lillack, Christian Kästner, Eric Bodden. "Tracking Load-time Configuration Options," In Proceedings of the 29th IEEE/ACM International Conference on Automated Software Engineering (ASE), Vasteras, Sweden, September 15 – 19, 2014. Presented by Max Lillack.

- Kästner, Christian & Pfeffer, Jürgen. "Limiting Recertification in Highly Configurable Systems. Analyzing Interactions and Isolation among Configuration Options," HotSoS 2014: 2014 Symposium and Bootcamp on the Science of Security, Article 23, Raleigh, NC, April 8-9, 2014. Presented by Juergen Pfeffer.

- Hemank Lamba, Thomas J. Glazier, Bradley Schmerl, Jürgen Pfeffer, David Garlan. "Detecting Insider Threats in Software Systems using Graph Models of Behavioral Paths (short paper) ," HotSoS 2015 Symposium and Bootcamp on the Science of Security, Article 20, Urbana-Champaign, IL, April 21-22, 2015. Presented by Hemank Lamba.

- Limin Jia, Shayak Sen, Deepak Garg, and Anupam Datta. "System M: A Program Logic for Code Sandboxing and Identification," Carnegie Mellon University, Technical Report CMU-CyLab-13-001, 2013 (updated July 2014).

- Ghita Mezzour. "Assessing the Global Cyber and Biological Threat," Ph.D. Thesis, Carnegie Institute of Technology, Electrical and Computer Engineering & School of Computer Science, Institute for Software Research, Carnegie Mellon University, Pittsburgh, PA, April 2015. Presented by Ghita Mezzour.

- Ghita Mezzour, Kathleen M. Carley, and L. Richard Carley. "An Empirical Study of Global Malware Encounters," Proceedings of ACM Symposium and Bootcamp on the Science of Security (HotSoS), Article 8, Urbana, IL, April 21 – 22, 2015. Presented by Ghita Mezzour.

- Ghita Mezzour, Kathleen Carley. "Spam Diffusion in a Social Network Initiated by Hacked Email Accounts," International Journal of Security and Networks, 9(3), pp. 144-53. November 2014.

- Mezzour, Ghita & Carley, Richard L. & Carley, Kathleen M. "Global Mapping of Cyber Attacks," Carnegie Mellon University, School of Computer Science, Institute for Software Research, Technical Report CMU-ISR-14-111. October 2014.

- Mezzour, Ghita & Carley, Richard L & Carley, Kathleen M. "Longitudinal Analysis of a Large Corpus of Cyber Threat Descriptions," Journal of Computer Virology and Hacking Techniques, Published online by Springer, pp. 1–12. June 9, 2014, Paris.

- Tingting Yu. "SimExplorer: A Testing Framework to Detect Elusive Software Faults," University of Nebraska at Lincoln Doctoral Dissertation, August 2014. Presented by Tingting Yu.

- Tingting Yu, Witawas Srisa-an, and Gregg Rothermel. "SimRT: An Automated Framework to Support Regression Testing for Data Races," In Proceedings of the International Conference on Software Engineering (ICSE), pp. 48-59, Hyderabad, India, May 31 – June 7, 2014. Presented by Tingting Yu.

- Joshua Sunshine, James D. Herbsleb, and Jonathan Aldrich. "Searching the State Space: A Qualitative Study of API Protocol Usability," International Conference on Program Comprehension (ICPC), co-located with ICSE, Florence, Italy, May 18 – 19, 2015. Presented by Joshua Sunshine.

- Cyrus Omar, Chenglong Wang, and Jonathan Aldrich. "Composable and Hygienic Typed Syntax Macros," 30th Symposium on Applied Computing (SAC), Salamanca, Spain, April 13 – 7, 2015. Presented by Cyrus Omar.

- Nathan Fulton, Cyrus Omar, and Jonathan Aldrich. "Statically Typed String Sanitation Inside a Python," Workshop on Privacy and Security in Programming (PSP), pp. 3-10, Portland, OR, Oct 20 – 24, 2014. Presented by Nathan Fulton.

- Darya Kurilova, Alex Potanin, and Jonathan Aldrich. "Wyvern: Impacting Software Security via Programming Language Design," Workshop on Evaluation and Usability of Programming Languages and Tools (PLATEAU), Portland, OR, October 20 – 24, 2014. Presented by Darya Kurilova.

- Michael Coblenz, Jonathan Aldrich, Brad Myers, and Joshua Sunshine. "Considering Productivity Effects of Explicit Type Declarations," Workshop on Evaluation and Usability of Programming Languages and Tools (PLATEAU), pp. 59-61, Portland, OR, October 20 – 24, 2014. Presented by Michael Coblenz.

- Cyrus Omar, Darya Kurilova, Ligia Nistor, Benjamin Chung, Alex Potanin, and Jonathan Aldrich. "Safely Composable Type-Specific Languages," European Conference on Object-Oriented Programming (ECOOP), pp. 105-30, Uppsala, Sweden, July 28 – August 1, 2014. Presented by Cyrus Omar.

- Michael Maass, Bill Scherlis, and Jonathan Aldrich. "In-Nimbo Sandboxing," Symposium and Bootcamp on the Science of Security (HotSOS), 2014. Raleigh, NC, August 8 – 9, 2014. Presented by Michael Maass.

- Joshua Sunshine, James D. Herbsleb, and Jonathan Aldrich. "Structuring Documentation to Support State Search: A Laboratory Experiment about Protocol Programming," Proceedings of the European Conference on Object-Oriented Programming (ECOOP), pp. 157-81, Uppsala, Sweden, July 28 – August 1, 2014. Presented by Joshua Sunshine.

- F. Medeiros, C. Kästner, M. Ribeiro, S. Nadi, and R. Gheyi. "The Love/Hate Relationship with The C Preprocessor: An Interview Study," In Proceedings of the 29th European Conference on Object-Oriented Programming (ECOOP), Berlin/Heidelberg: Springer-Verlag, 2015.

- C. Hunsen, J. Siegmund, O. Lessenich, S. Apel, B. Zhang, C. Kästner, and M. Becker. "Preprocessor-Based Variability in Open-Source and Industrial Software Systems: An Empirical Study," Empirical Software Engineering (ESE), Special Issue on Empirical Evidence on Software Product Line Engineering, 2015.

- H. Hibshi, T. D. Breaux, S. B. Broomell, "Assessment of Risk Perception in Security Requirements Composition." To Appear: IEEE 23rd International Requirements Engineering Conference (RE'15), 2015.

- T. D. Breaux, D. Smullen, H. Hibshi. "Detecting Repurposing and Over-collection in Multi-Party Privacy Requirements Specifications," To Appear: IEEE 23rd International Requirements Engineering Conference (RE'15), Ottawa, Canada, Sep. 2015.

- Du Li, Alex Potanin, and Jonathan Aldrich. "Delegation vs Inheritance for Typestate Analysis," In the 17th Workshop on Formal Techniques for Java-like Programs (FTfJP), 2015.

- Joseph Lee, Jonathan Aldrich, Troy Shaw, and Alex Potanin. "A Theory of Tagged Objects," In European Conference on Object-Oriented Programming (ECOOP), 2015.

- Michael Coblenz, Robert Seacord, Brad Myers, Joshua Sunshine and Jonathan Aldrich. "A Course-Based Usability Analysis of Cilk Plus and OpenMP," To appear in IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC).

NC STATE UNIVERSITY

NORTH CAROLINA STATE UNIVERSITY

PI Laurie Williams

North Carolina State University's (NCSU) Science of Security Lablet (SoSL), led by Principal Investigator Laurie Williams has embraced and helped build a foundation for NSA's vision of the Science of Security (SoS) and of a SoS community. NCSU lablet has emphasized data-driven discovery and analytics to formulate, validate, evolve, and solidify the theory and practice of security. Efforts in the current lablet have yielded significant findings, providing a deeper understanding of users' susceptibility to deception, developers' adoption of security tools, how trust between people relates to their commitments. Motivated by NSA's overarching vision for SoS and building on our experience and accomplishments, (1) NCSU is developing a science-based foundation for the five hard problems that we previously helped formulate; and (2) fostering a SoS community with high standards for reproducible research. NCSU's approach involves a comprehensive, rigorous perspective on SoS, including an integrated treatment of technical artifacts, humans (both stakeholders and adversaries) along with relationships and processes relevant to the hard problems. Continual evaluation of our research and community development efforts is key to our approach. NCSU's collaboration has engaged 55 authors and 13 institutions. The Lablet also hosted an NSA strategy meeting on the Science of Privacy. This workshop shared hard problem strategy and research methods on this emerging topic.

We have formed teams to conduct scientific research and evaluate progress on hard problems: Security Metrics and Models; Humans; Policy; and Resilient Architectures. The Scalability and Composability hard problem has no explicit team since we address it as a secondary hard problem in several of our projects. Each Hard Problem team is composed of three or four projects researching complimentary aspects of the Hard Problem. We also have additional teams for Research Methods, Community Development and Support, and for Evaluation.

Security Metrics and Models

Attack Surface and Defense-in-Depth Metrics: Rochester Institute of Technology: Andy Meneely, NC State University: Laurie Williams

Systematization of Knowledge from Intrusion Detection Models: NC State University: Huaiyu Dai, Rochester Institute of Technology: Andy Meneely

Vulnerability and Resilience Prediction Models: NC State University: Mladen Vouk, Laurie Williams

Human Behavior

Warning of Phishing Attacks: Supporting Human Information Processing, Identifying Phishing Deception Indicators, and Reducing Vulnerability: NC State University: Christopher B. Mayhorn, Emerson Murphy-Hill

A Human Information-Processing Analysis of Online Deception Detection: Purdue University: Robert W. Proctor, Ninghu Li

Leveraging the Effects of Cognitive Function on Input Device Analytics to Improve Security: NC State University: David L. Roberts, Robert St. Amant

Policy-Governed Secure Collaboration

Understanding Effects of Norms and Policies on the Robustness, Liveness, and Resilience of Systems: NC State University: Emily Berglund, Jon Doyle, Munindar Singh

Formal Specification and Analysis of Security - Critical Norms and Policies: NC State University: Jon Doyle, Munindar Singh, Rada Chirkova

Scientific Understanding of Policy Complexity: Purdue University: Ninghui Li, Robert Proctor, NC State University: Emerson Murphy-Hill

Resilient Architectures

Resilience Requirements, Design, and Testing: University of Virginia: Kevin Sullivan, NC State University: Mladen Vouk, University of North Carolina at Charlotte: Ehab Al-Shaer

Redundancy for Network Intrusion Prevention Systems (NIPS): University of North Carolina: Mike Reiter

Smart Isolation in Large-Scale Production Computing Infrastructures: NC State University: Xiaohui (Helen) Gu, William Enck

Automated Synthesis of Resilient Architectures: University of North Carolina at Charlotte: Ehab Al-Shaer

Additional Teams

Research Methods, Community Development and Support: University of Alabama: Jeff Carver, NC State University: Lindsey McGowen, Jon Stallings, Laurie Williams, David Wright

Evaluation: NC State University: Lindsey McGowen, Jon Stallings, David Wright, University of Alabama: Jeff Carver

Our collective efforts have advanced the Science of Security by bringing forth best practices and research methodologies to the hard problems in security.

- We are developing an agent-based simulation methodology for secure collaboration. Agent-based simulation is an established methodology for understanding complex systems formed of the interactions of autonomous parties and where analytical solutions are not expected to be found. It has been used in areas such as epidemiology and, given the complexity of systems in connection with cybersecurity, should be a component of security research.
- We are advancing the nature and rigor of empirical studies of end users via lab experiments and surveys.
- We are advancing the empirical study of system administration via evaluations of artifacts such as software and policies through tools as well as through experimental studies in which participants apply competing approaches to create specifications (thereby helping us address insidious security errors through errors in modeling or configuration).
- Through example, we are advancing the systematic development of survey papers in the science of security.

FUNDAMENTAL RESEARCH

Security Metrics and Models

We have results related to attack surface metrics and vulnerabilities. Many organizations prioritize security efforts around the general idea of attack surface (entry and exit points of a software program), considering areas of the code not reachable by an attacker to be a lower priority. However, the process for practically identifying what part of the code is on the attack surface and specific attack surface metrics have not been validated. We have results (weakly) associating our attack surface metrics with vulnerabilities. Additionally, our analysis of crash dumps at Microsoft indicates that the code identified in crash dumps accounts for most (94.6%) of the vulnerabilities, indicating that crash dumps may be used to indicate whether a piece of code is on the attack surface or not.

Projects:

- Attack Surface and Defense-in-Depth Metrics
- Systematization of Knowledge from Intrusion Detection Models
- Vulnerability and Resilience Prediction Models

Human Behavior

We have developed a cognitive model of users based on the well-known ACT-R framework. We have developed an understanding of observable human behaviors that indicates the level of thought a user puts into an action (as a measure of the naturalness of the action). These contributions provide some of the bases of the science underlying the humans hard problem by leading us to an understanding of (1) how users process information and make security-relevant mistakes and the bases on which we may identify such mistakes; (2) how to distinguish potential deceptive user behaviors through largely unobtrusive observations of users; and (3) how to generate cognitively and ecologically relevant warnings to users to assist them in their security-relevant decision making.

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

- A Human Information-Processing Analysis of Online Deception Detection
- Leveraging the Effects of Cognitive Function on Input Device Analytics to Improve Security

Policy

We have developed metrics of policy complexity that capture how difficult a set of policies is for people to comprehend, which are indicative of configuration errors that lead to vulnerabilities. We have identified such errors in practical enterprise policies. We have developed and evaluated an argumentation-based approach for capturing firewall requirements that reduces errors and improves comprehensibility over traditional methods. We are studying policy errors in software and how to ameliorate such errors based on principles pertaining to software analytics. We have developed a normative formulation of accountability that captures its essential features separately from traceability; we have additionally developed a (partial) approach that relates normative relationships to data representations as a basis for logging and analytics. We have developed a simulation framework in which to study the robustness and resilience of norms that modulate the security-relevant behaviors and interactions of users, as a basis for understanding and exploring potential norms.

Projects:

- Understanding Effects of Norms and Policies on the Robustness, Liveness, and Resilience of Systems
- Formal Specification and Analysis of Security-Critical Norms and Policies
- Scientific Understanding of Policy Complexity

Resilient Architectures

We have created a classification scheme of existing isolation techniques. The purpose of the scheme is to enable the identification of underlying design principles the tradeoffs between them. Discovering these principles will aid in the design of the next generation of smart isolation techniques to support resilient architectures. Similarly, we have created a taxonomy of resiliency metrics. The purpose of the metrics is to estimate the resiliency level that a system exhibits given a specific attack model or scenario, and the implied recommended actions to improve resiliency.

Projects:

- Resilience Requirements, Design, and Testing
- Redundancy for Network Intrusion Prevention Systems (NIPS)
- Smart Isolation in Large-Scale Production Computing Infrastructures
- Automated Synthesis of Resilient Architectures

COMMUNITY ENGAGEMENTS

- In April 2014, we organized the first HotSoS in Raleigh. Over 130 leaders from government, industry, and the academic community met to discuss new and ongoing programs in security science. The presentations emphasized a broad range of topics including computing architectures, networks, software engineering practices, models of human interaction and behavior, organizational models, and evaluation methodologies.
- In Summer 2014, we conducted a two-day research workshop for lablet participants on best practices for the science of security. This workshop included sessions conducted by a statistician on experimental design; by a "science of science" methodologist on the ways in which we can collectively advance the science of security; by a computer scientist on conducting empirical software engineering research. A similar workshop was held May 2015.
- In October 2014, we conducted an Industry Day workshop whose first part involved Pecha Kucha presentations by all lablet projects; presentations by invited industry speakers; and a poster session by students. We used this workshop as a way to engage more closely with industry colleagues, both in advancing cybersecurity and in promoting the science of security.
- We are developing a rubric for reviewing research publications in a way that seeks to bring out and evaluate their core scientific claims and findings. We offered a workshop at the January 2015 quarterly meeting based on this rubric. The idea is that this rubric

would sensitize researchers to the scientific aspects of security research and thereby lead to papers and peer reviews that are more clearly scientific and thus lead to improved scientific research overall.

- We have taken numerous opportunities to give keynote and other invited lectures on the Science of Security to broader computer science communities as a way to bring them into the fold.
- We have identified a seed list of publication venues where Science of Security research appears. We are in the midst of engaging the community (at other lablets) on refining and ranking a list of venues.
- On June 23–24, we hosted an invitation-only planning workshop for an upcoming NSA workshop on Science of Privacy. This gave us an opportunity to discuss Science of Security with visitors and to present posters on Lablet research.
- Held our Community Day on October 29, 2015, where we presented our research and held discussions with local industry and government colleagues.

Invited NSU Talks at the SoS Quarterly Meetings

Science of Security Quarterly Lablet PI Meeting, July 2014:

Lindsey McGowen, "Evaluating the Development of a Science of Security: A Plan for Measuring and Demonstrating Lablet Contributions."

Science of Security Quarterly Lablet PI Meeting, October 2014:

Andrew Meneely, "Developing Security Metrics."

Munindar Singh, "Survey on Policy-Governed Secure Collaboration."

Science of Security Quarterly Lablet PI Meeting, January 2015:

William Enck, "Systematizing Isolation Techniques."

Lindsey McGowen, "Customized Bibliometrics for Evaluating Computer Science Research."

Ehab Al-Shaer, "On Objective Resiliency Analysis of Smart Grid Energy Management Systems."

EDUCATIONAL

Hold (approximately) bi-weekly research seminars in the Fall 2014 and Spring 2015 academic semesters. These seminars consist of two main types of discussions and are attended by NCSU as well as remote participants at our collaborating institutions.

- In research design seminars, students present their designs for their proposed study, including not only the motivation and existing theoretical frameworks, but also details of the theoretical or empirical investigations they plan to carry out. Our motivation for discussing research designs is, first, to reflect on the nature of an investigation before launching into the effort and, second, to vet the proposed design in consultation with peers in the lablet. The intended benefit is in strengthening the scientific basis of the research by improving clarity of the hypotheses and metrics underlying the research as well as ensuring the design would help evaluate those hypotheses.
- In manuscript review seminars, students make a presentation about a manuscript they are preparing for submission for peer review. The intended benefit is in strengthening the positioning of the research with respect to the literature and in discussing the robustness of the claimed evaluation of the hypotheses.
- We continued to collect and organize feedback to student presenters during our regular seminars. We further refined feedback instruments with a view to guiding presenters and the audience toward best practices in the science of security.
- On May 27–28, 2015, we conducted a two-day summer workshop for the purpose of increasing our collective knowledge and experience with scientific research in security.
    - Our research methods team had previously published a template of components of a well-structured scientific research paper. Participants provided feedback on the template, customized the template for each hard problem, and evaluated previously published security research for the components in the template.
    - Two tutorials were given on statistical methods and bibliometrics.
    - In addition to lablet faculty and students, the workshop was attended by an industry researcher and some NSA personnel (primarily from the NCSU Laboratory for Analytic Sciences).

PUBLICATIONS

- Amant, R. S. & Goodwin, P. R. & Domínguez, I. & Roberts, D. L. (2015) "Toward Expert Typing in ACT-R." *Proceedings of the 2015 International Conference on Cognitive Modeling (ICCM 15)*.
- Chopra, A. K. & Singh, M. P. (2015) "Cupid: Commitments in Relational Algebra." *Proceedings of the 23rd Conference on Artificial Intelligence (AAAI)*.
- Chopra, A. K. & Singh, M. P. (2014) "The Thing Itself Speaks: Accountability as a Foundation for Requirements in Sociotechnical Systems." *Requirements Engineering and Law (RELAW), 2014 IEEE 7th International Workshop on.*
- Domínguez, I. X. & Goel, A. & Roberts, D. L. & Amant, R. S. (2015) "Detecting Abnormal User Behavior Through Pattern-mining Input Device Analytics." *Proceedings of the 2015 Symposium and Bootcamp on the Science of Security (HotSoS 2015)*.
- Dou, K. & Wang, X. & Tang, C. & Ross, A. & Sullivan, K. (2015) "An Evolutionary Theory-Systems Approach to a Science of the Ilities." *Proceedings of the Conference on Systems Engineering Research (CSER 2015)*.
- Du, H. & Narron, B. Y. & Ajmeri, N. & Berglund, E. & Doyle, J. & Singh, M. P. (2015) "ENGMAS textendash Understanding Sanction under Variable Observability in a Secure Environment." *Proceedings of the 2nd International Workshop on Agents and CyberSecurity (ACySE)*.
- Du, H. & Narron, B. Y. & Ajmeri, N. & Berglund, E. & Doyle, J. & Singh, M. P. (2015) "Understanding Sanction under Variable Observability in a Secure, Collaborative Environment." *Proceedings of the International Symposium and Bootcamp on the Science of Security (HotSoS)*.
- Heorhiadi, V. & Fayaz, S. & Reiter, M. K. & Sekar, V. (2014) "SNIPS: A Software-Defined Approach for Scaling Intrusion Prevention Systems via Offloading." *10th International Conference on Information Systems Security, ICISS 2014*.
- Huang, Y. & He, X. & Dai, H. (2015) "Systematization of Metrics in Intrusion Detection Systems." (Poster) *ACM Proc. Of the Symposium and Bootcamp on the Science of Security (HotSoS), University of Illinois at Urbana-Champaign, IL*
- Kim, D. & Vouk, M. (2014) "A survey of common security vulnerabilities and corresponding countermeasures for SaaS." *Second IEEE International workshop on Cloud Computing Systems, Networks, and Applications (CCSNA-2014)*.
- Mayhorn, C. B. & Welk, A. K. & Zielinska, O. A. & Murphy-Hill, E. (2015) "Assessing individual differences in a phishing detection task." *Proceedings of the Annual International Ergonomics Association Conference*.
- Rahman, M. A. & Al-Shaer, E. & Bobba, R. B. (2014) "Moving Target Defense for Hardening the Security of the Power System State Estimation." *First ACM Workshop on Moving Target Defense*.
- Rahman, M. A. & Al-Shaer, E. & Kavasseri, R. G. (2014) "Impact Analysis of Topology Poisoning Attacks on Economic Operation of the Smart Power Grid." *Proceedings of the 34th International Conference on Distributed Computing Systems (ICDCS)*.
- Rivers, A. T. & Vouk, M. A. & Williams, L. A. (2014) "On Coverage-Based Attack Profiles." *Software Security and Reliability-Companion (SERE-C), 2014 IEEE Eighth International Conference on*.
- Singh, M. P. (2015) "Norms as a Basis for Governing Sociotechnical Systems: Extended Abstract." *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*.
- Singh, M. P. (2015) "Cybersecurity as an Application Domain for Multiagent Systems." *Proceedings of the 14th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*.
- Subramani, S. (2014) "A Study of Fedora Security Profile."
- Venkatakrishnan, R. (2014) "Redundancy-Based Detection of Security Anomalies in Web-Server Environments."

- Welk, A. K. & Mayhorn, C. B. (2015) "All Signals Go: Investigating How Individual Differences Affect Performance on a Medical Diagnosis Task Designed to Parallel a Signal Intelligence Analyst Task." *Symposium and Bootcamp on the Science of Security (HotSoS)*.
- Zielinska, O. A. & Tembe, R. & Hong, K. W. & Ge, X. & Murphy-Hill, E. & Mayhorn, C. B. (2014) "One Phish, Two Phish, How to Avoid the Internet Phish: Analysis of Training Strategies to Detect Phishing Emails." *Human Factors and Ergonomics Society Annual Meeting*.
- Zielinska, O. & Welk, A. & Mayhorn, C. B. & Murphy-Hill, E. (2015) "Exploring expert and novice mental models of phishing." *Proceedings of the 2nd HotSoS: Symposium and Bootcamp on the Science of Security*.
- Knight, J., Xiang, J., and Sullivan, K., "Real-World Types and Their Applications." to appear, *Proceedings of The International Conference on Computer Safety, Reliability, and Security (SAFECOMP 2015)*. Delft, The Netherlands. Sept. 23–25, 2015.
- Donghoon Kim, Henry E. Schaffer, and Mladen A. Vouk, "About PaaS Security." *Proceedings of the 3rd International IBM Cloud Academy Conference (ICACON 2015)*, Budapest, Hungary. May 21–23, 2015.
- Anoosha Vangaveeti. "An Assessment of Security Problems in Open Source Software," (M.S. Thesis, NC State University, 2015).
- Chen, J., Yang, W., Xiong, A., Li, N., & Proctor, R. W. (August 5, 2015). "Warning users of phishing attacks with a Google Chrome extension." Talk presented at *Human-Computer Interaction International 2015*, Los Angeles, CA.
- Proctor, R. W., & Chen, J. (2015). "The role of human factors/ergonomics in the science of security: Decision making and action selection in cyberspace." *Human Factors, 57*, 721–727.
- Welk, A., Zielinska, O., Tembe, R., Xe, G., Hong, K. W., Murphy-Hill, E., & Mayhorn, C. B. (in press). "Will the 'Phisher-men' Reel you in? Assessing Individual Differences in a Phishing Detection Task." *International Journal of Cyber Behavior, Psychology, and Learning.*
- Roopak Venkatakrishnan, Mladen Vouk, "Using Redundancy to Detect Security Anomalies – Towards IoT Security Attack Detectors," *ACM Ubiquity*, to appear, 2015.
- Donghoon Kim and Mladen A. Vouk, "Securing Scientific Workflows." *Proceedings of the 2015 IEEE International Conference on Software Quality, Reliability and Security (QRS)*, Vancouver, Canada, August 2–5, 2015 as part of the QRS workshop on Trustworthy Computing, 2015.
- Yu Xianqing, Peng Ning, Vouk, M.A., "Enhancing security of Hadoop in a public cloud," in the *Proceedings of the 6th International Conference Information and Communication Systems (ICICS)*, 7–9 April 2015, pp. 38–43.

NC STATE UNIVERSITY

ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

PI DAVID NICOL

The University of Illinois at Urbana-Champaign (UIUC) Lablet, led by Principal Investigator (PI) David Nicol, is contributing broadly to the development of security science while leveraging Illinois expertise in resiliency, which in this context means a system's demonstrable ability to maintain security properties even during ongoing cyber attacks. The Lablet's work draws on several fundamental areas of computing research. Some ideas from fault-tolerant computing can be adapted to the context of security. Strategies from control theory are being extended to account for the high variation and uncertainty that may be present in systems when they are under attack. Game theory and decision theory principles are being used to explore the interplay between attack and defense. Formal methods are being applied to develop formal notions of resiliency. End-to-end system analysis is being employed to investigate resiliency of large systems against cyber attack. The Lablet's work also draws upon ideas from other areas of mathematics and engineering as well. UIUC is actively engaged in five projects addressing metrics (primary), human behavior, policy, and resiliency. These projects are looking at data-driven models of attacker behavior, human circumvention of security, and data-driven model-based decision-making. Twenty students are supported.

The team is comprised of mostly faculty and researchers from the University of Illinois at Urbana-Champaign. Project by project details of the personnel are listed below.

A Hypothesis of Testing and Framework for Network Security. Illinois: Brighten Godfrey, Matt Caesar, David Nicol, Bill Sanders; Illinois Institute of Technology: Dong (Kevin) Jin

Data-Driven Model-Based Decision-Making. Illinois: Bill Sanders, Masooda Bashir, David Nicol; Newcastle University, UK: Aad Van Moorsel

Data Driven Security Models and Analysis. Illinois: Ravi Iyer, Zbigniew Kalbarczyk, and Adam Slagell

Science of Human Circumvention of Security. Illinois: Tao Xie; University of Southern California: Jim Blythe; University of Pennsylvania: Ross Koppel; Dartmouth College: Sean Smith

Static-Dynamic Analysis of Security Metrics for Cyber-Physical Systems. Illinois: Sayan Mitra and Geir Dullerud; Rice University: Swarat Chaudhuri

The Science of Security (SoS) has many attributes that range from the 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 the 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.

Illinois hosted the *2015 Symposium and Bootcamp on the Science of Security (HotSoS)* in April, reached out to new students through a graduate Science of Security seminar, and provided summer internships to graduate students. The four summer interns were a diverse group from two other universities. Their work was presented in the poster session at HotSoS 2015.

FUNDAMENTAL RESEARCH

Project: A Hypothesis of Testing and Framework for Network Security. Illinois: Brighten Godfrey, Matt Caesar, David Nicol, Bill Sanders; Illinois Institute of Technology: Dong (Kevin) Jin

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 problem of predictive security metrics. To realize NetHTM, we need the ability to model and predict behavior of networked systems. We made advances in modeling and enforcing correct behavior in dynamic networks. This required a model of network behavior under timing uncertainty; that is, in a dynamic network, we will have only imperfect information about the exact time network events take place, which makes reasoning about properties difficult. We used our model and verification algorithms on top of it to develop network control algorithms which preserve specified properties across time. A paper on this project was submitted in 2014 and accepted to one of the two top venues in computer networking, *USENIX NSDI '15, the 12th USENIX Symposium on Networked Systems Design and Implementation.*

In addition, we made progress on modeling virtualized networks, with an emphasis on determining when virtual and physical networks may differ, and resolving these inconsistencies. A paper on this work appeared at the ACM Workshop on Hot Topics in Software Defined Networks (HotSDN) in August 2014, where it received the Best Paper Award.

Project: Data-Driven Model-Based Decision-Making. Illinois: Bill Sanders, Masooda Bashir, David Nicol; Newcastle University, UK: Aad Van Moorsel

System security analysis requires a holistic approach that considers the behavior of non-human subsystems, bad actors or adversaries, and expected human participants such as users and system administrators. Modeling and evaluating human behavior is challenging, but it is an imperative component in security analysis. We have developed and implemented a modeling formalism to formally describe the behavior of human participants and how their decisions affect overall system performance and security. With the HITOP modeling formalism and its implementation in the Mobius modeling framework, we are able to produce quantitative security metrics for cyber-human systems. HITOP evaluates a human's opportunity, willingness, and capability to perform individual tasks in their daily behavior. Partnered with an effective data collection strategy to validate model parameters, we have made good progress toward a sound model of human behavior. Our next steps include development of a case study to validate our approach, as well as further refinement of the HITOP methodology based on the experience we gain from the study.

Project: Data Driven Security Models and Analysis. Illinois: Ravi Iyer, Zbigniew Kalbarczyk, and Adam Slagell

We developed and evaluated AttackTagger, a Factor Graph based framework for accurate and preemptive detection of attacks, i.e., before the system misuse. A Factor Graph is a type of probabilistic graphical model that can describe complex dependencies among random variables using an undirected graph representation, specifically a bipartite graph. The bipartite graph representation consists of variable nodes representing random variables, factor nodes representing local functions (or factor functions), and edges connecting the two types of nodes. In our model, random variables correspond to observable events (e.g., alerts generated by an intrusion detection system) and unknown user states (benign, suspicious, or malicious). The factor functions represent dependencies between events and user states. By evaluating the constructed graph, we can determine the user state at each stage of an attack. We used security logs on real-incidents that occurred over a six-year period at the National Center for Supercomputing Applications (NCSA) to evaluate AttackTagger. Our data consist of security incidents that led to the target system being compromised, i.e., the attacks were detected after the fact. AttackTagger detected 74% of attacks; a vast majority of them were detected before the system misuse (minutes or hours). Importantly, AttackTagger uncovered hidden malicious users that were missed by the intrusion detection systems during the incident and by security analysts in post-incident forensic analysis.

Project: Science of Human Circumvention of Security. Illinois: Tao Xie; University of Southern California: Jim Blythe; University of Pennsylvania: Ross Koppel; Dartmouth College: Sean Smith

We continue to study people's trust in cyber security, websites, their organization's databases, and use of the Internet. We focus especially on passwords as a prime 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 designing a mechanical Turk experiment/simulation to further examine users' use of passwords, workarounds, cyber trust, and strategies. Results include duplication in our simulation of a version of "uncanny descent" in which making constraints on passwords more complex can decrease overall security. To study people's trust in cyber security, especially mobile app security, we focus on exposing contextual information to enable mobile app users to make informed decisions on mobile app security.

Project: Static-Dynamic Analysis of Security Metrics for Cyber-Physical Systems. Illinois: Sayan Mitra and Geir Dullerud; Rice University: Swarat Chaudhuri

Addressing the hard problem of developing predictive security metrics, in this collaborative project, we have formulated the general problem of controller synthesis in the presence of resource constrained adversaries; namely, given an adversary of a certain class, parameterized 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 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. Going forward, we will expand this research to address significantly more complex systems involving more general nonlinear dynamics, and apply them to controller synthesis for, and security evaluation of, autonomous and semi-autonomous unmanned vehicles.

COMMUNITY ENGAGEMENTS

The UIUC Lablet team has put forth outreach efforts throughout the Science of Security community. The *2015 Symposium and Bootcamp on the Science of Security (HotSoS)* was held at the University of Illinois at Urbana-Champaign on April 21-22. The symposium brought together researchers from numerous disciplines seeking a comprehensive and methodical approach to identifying and removing threats.

UIUC SoS Lablet Bi-weekly Research Seminars September 2014 – April 2015

- Yu Wang, "Entropy-minimizing Mechanism for Differential Privacy of Discrete-time Linear Feedback Systems"
- Zhenqi Huang, "Verification from Simulations and Modular Annotations"
- David Nicol, "Science of Security Hard Problems: A Lablet Perspective"
- Soudeh Ghorbani, "Towards Correct Network Virtualization"
- Ravi Iyer, "Resiliency Survey: Challenges Going Forward"
- Ken Keefe, "Making Sound Security Decisions Using Quantitative Security Metrics"
- Tao Xie, "AppContext: Differentiating Malicious and Benign Mobile App Behavior Under Contexts"
- Mohammad Noureddine, "Human Aware Science of Security"
- Tao Xie, "Science of Human Circumvention of Security"
- Brighten Godfrey, "Hypothesis Testing for Network Security"
- Phuong Cao, "Preemptive Intrusion Detection: Theoretical Framework and Real-world Measurements"
- Geir Dullerud, "Static Dynamic Analysis of Security Metrics for Cyber Physical Systems"
- Wenxuan Zhou, "Enforcing Customizable Consistency Properties in Software-Defined Networks"
- Mohammad Noureddine, "A Taxonomy of Human Behavior in Cybersecurity"

SoS Quarterly Meetings

- July 2014, NSA SoS Quarterly Meeting, Bill Sanders, "Making Sound Design Decisions Using Quantitative Security Metrics"
- October 2014, NSA SoS Quarterly Lablet Meeting, Ravi Iyer, "Survey on Resilience"
- October 2014, NSA SoS Quarterly Lablet Meeting, Sayan Mitra, "Static-Dynamic Analysis of Security Metrics for Cyber-Physical Systems"
- January 2015, NSA SoS Quarterly Lablet Meeting, Matt Caesar, "Hypothesis Testing for Network Security"
- January 2015, NSA SoS Quarterly Lablet Meeting, Ravi Iyer, "Preemptive Intrusion Detection: Theoretical Framework and Real-world Measurements"

SoS Speaker Series

- Somesh Jha, "Thoughts on Retrofitting Legacy Code for Security," University of Wisconsin, April 2015

Other Presentations

- Keynote Talk: Ross Koppel, "Software Loved by its Vendors and Disliked by 70% of its Users: Two Trillion Dollars of Healthcare Information Technology's Promises and Disappointments," *2014 USENIX Summit on Health Information Technologies,* August 2014.
- Presentation: Ross Koppel, "Ethnography of Computer Security Evasions in Healthcare Organizations: Circumvention and Cyber Controls," *European Sociological Association Midterm Conference*, August 2014.
- Tutorial: Tao Xie. "Text Analytics for Security," *21st ACM Conference on Computer and Communications Security (CCS)*, November 2014.
- Invited Talk: Sean Smith, "Circumvention: Why Do Good People Do Bad Things, and What Can We Do About It?," Rutgers University, Department of Electrical and Computer Engineering Colloquium, December 2014.
- Keynote Talk: Ross Koppel, "Healthcare Software Usability and the Influence on Compliance with Cybersecurity Rules," Royal College of Physicians (Edinburgh), February 2015.
- Invited Seminar: Ross Koppel, "Healthcare Software Usability and the Influence on Compliance with Cybersecurity Rules, Wales Health Trust at Prince of Wales Hospital, February 2015.
- Invited Tutorial: Jim Blythe and Sean Smith, "Understanding and Accounting for Human Behavior," *Symposium and Bootcamp on the Science of Security (HotSoS)*, April 2015.
- Presentation: Ross Koppel, Sean W. Smith, and Harold Thimbleby, "What You See Is What You See: Misinforming Displays in Electronic Health Care Records and Medical Devises," *2015 International Symposium on Human Factors and Ergonomics in Health Care: Improving Outcomes (HFES)*, April 2015.
- Invited Talk: Sean W. Smith, "Trust Challenges in Massive Multi-organization Distributed Systems," Dagstuhl Seminar: Assuring Resilience, Security and Privacy for Flexible Networked Systems and Organizations, April 2015.
- Presentation: Wenxuan Zhou, "Enforcing Customizable Consistency Properties in Software-Defined Networks," *USENIX Symposium on Network Systems Design and Implementation (NSDI)*, April 2015.
- Invited Talk: Kevin Jin, "Enforcing Customizable Consistency Properties in Software-Defined Networks," 4th Greater Chicago Area Systems Research Workshop (GCASR), April 2015.
- Presentation: Wei Yang (advised by Tao Xie), "AppContext: Differentiating Malicious and Benign Mobile App Behavior Under Contexts," *37th International Conference on Software Engineering (ICSE)*, Florence, Italy, May 2015.
- Presentation: Kevin Jin, "VT-Mininet: Virtual-time-enabled Mininet for Scalable and Accurate Software-Define Network Emulation," *ACM SIGCOMM Symposium on SDN Research 2015 (SOSR)*, Santa Clara, CA, June 2015.
- Presentation: Jiaqi Yan, "A Virtual Time System for Linux-container-based Emulation of Software-defined Networks," *ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, London, UK, June 2015

NSA SoS HotSoS Presentations, April 2015

- John C. Mace, Charles Morisset, and Aad van Moorsel, "Modelling User Availability in Workflow Resiliency

Analysis," *Symposium and Bootcamp on the Science of Security (HotSoS)*, April 2015.

- Mohammad Noureddine, Ken Keefe, William H. Sanders and Masooda Bashir, "Quantitative Security Metrics with Human in the Loop," Poster presented at *Symposium and Bootcamp on the Science of Security (HotSoS)*, April 2015.

- Robert Cain and Aad van Moorsel, "Optimisation of Data Collection Strategies for Model-Based Evaluation and Decisions-Making," Poster presented at *Symposium and Bootcamp on the Science of Security (HotSoS),* April 2015.

- Tao Xie, Judith Bishop, Nikolai Tillmann, and Jonathan de Halleux, "Gamifying Software Security Education and Training via Secure Coding Duels in Code Hunt," Poster presented at *Symposium and Bootcamp on the Science of Security (HotSoS)*, April 2015.

- Tutorial: Zbigniew Kalbarczyk, "Resilience of Cyber-Physical Systems and Technologies," *Symposium and Bootcamp on the Science of Security, (HotSoS)*, April, 2015

EDUCATIONAL

[Godfrey, Caesar, Nicol, Sanders, Jin] David Nicol developed and taught a graduate course in the Science of Security for spring 2015. The seminar, **ECE 598**, examined a number of security papers from the literature and for each discussed the questions: what attributes of this paper either study security properties themselves as first class objects, or use scientific methodologies to identify and assess security properties; and in what ways is this paper lacking in scientific foundations for the work it presents. Discussions were lively, and exhibited a variance in students' understandings and expectations of "Science of Security." The security of computers, communications, and data is of great concern to our society. Decades of research have produced solutions to a variety of isolated problems, some of which have been produced using techniques that are recognizable as "scientific," others of which appear to be ad-hoc. There is a growing sentiment in the community that research in security should be conducted when possible on a scientific or engineering basis. This course examined the questions of what might constitute a science of security, framing the questions around five "hard areas" proposed by the NSA: Composition, Policy, Metrics, Resiliency, and Human Factors. The students read and presented papers from the literature that exemplified a scientific approach to security, and wrote essays on the questions raised by the course. The course was intended for graduate students interested in trustworthy systems research.

[Xie, Blythe, Koppel, Smith] PI Tao Xie is designing teaching materials on **Code Hunt** (https://www.codehunt.com/) released by Microsoft Research for teaching and training students on software security. The teaching materials incorporate educational gamification to teach students on improving their software security skills. Some initial designs are described in the HotSoS 2015 poster paper.

[Godfrey, Caesar, Nicol, Sanders, Jin] Kevin Jin has developed and taught a new graduate-level course, **CS558 Advanced Computer Security**, at the Illinois Institute of Technology. A key topic in this course is network security, which covers some of the research results of the project. Kevin Jin received the CS Teacher of the Year Award in May 2015, mainly because of his contribution to the cyber security curriculum at the Information Trust Institute (IIT)**.**

[Xie, Blythe, Koppel, Smith] Ross Koppel is developing a course on the ethnography of organizational workflow and cyber workarounds. This course will involve approximately 20 students interviewing workers about password circumvention and ways of accessing information that is not part of official policy. These findings will help to continue our work of discovering ways well-indented workers create vulnerabilities in cyber security.

[UIUC SoS Lablet] Three 2015 undergraduate and one graduate intern have been selected for the UIUC SoS Lablet Summer Internship Program. The interns have been working on their own research projects guided by SoS faculty. The internship program began on June 1 and concluded on July 24 with a poster sessiions.

PUBLICATIONS

- Dong Jin and Yi Ning. "Securing Industrial Control Systems with a Simulation-based Verification System," 2014 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation, May 2014.

- Vijay Kothari, Jim Blythe, Sean W. Smith, Ross Koppel. "Agent-Based Modeling of User Circumvention of Security," 1st International Workshop on Agents and CyberSecurity (ACySE '14), Article 5, 4 pages, May 2014.

- Cuong Pham, Zachary Estrada, Zbigniew Klabarczyk, and Ravishankar Iyer. "Reliability and Security Monitoring of Virtual Machines using Hardware Architectural Invariants," 44th International Conference on Dependable Systems and Networks, June 2014. William C. Carter Award for Best Paper based on PhD work and Best Paper Award voted by conference participants.

- G. Wang, Zachary Estrada, Cuong Pham, Zbigniew Klabarczyk, and Ravishankar Iyer. "Hypervisor Introspection: Exploiting Timing Side-Channels against VM Monitoring," 44th International Conference on Dependable Systems and Networks, June 2014.

- Soudeh Ghorbani and Brighten Godfrey, "Towards Correct Network Virtualization," ACM Workshop on Hot Topics in Software Defined Networks (HotSDN), August 2014. Best Paper Award.

- Jim Blythe, Ross Koppel, Vijay Kothari and Sean Smith. "Ethnography of Computer Security Evasions in Healthcare Settings: Circumvention as the Norm," 2014 USENIX Summit on Health Information Technologies, August 2014.

- Cuong Pham, Zachary Estrada, Phuong Cao, Zbigniew Kalbarczyk, and Ravishankar Iyer. "Building Reliable and Secure Virtual Machines using Architectural Invariants," IEEE Security and Privacy Magazine, vol. 12, no. 5, pp. 82-85, September – October 2014.

- Ross Koppel, Sean Smith, Jim Blythe and Vijay Kothari. "Workarounds to Computer Access in Healthcare Organizations: You Want My Password or a Dead Patient?" Driving Quality in Informatics: Fulfilling the Promise, Series on Technology and Informatics, vol. 208, Feb.-Mar, 2015.

- Zhengi Huang, Yu Wang, Sayan Mitra and Geir Dullerud. "Controller Synthesis for Linear Time-varying Systems with Adversaries," January 2015. http://arxiv.org/abs/1501.04925

- Ross Koppel, Sean Smith, James Blythe, and Vijay Kothari. "Workarounds to Computer Access in Healthcare Organizations: You Want My Password or a Dead Patient?," Information Technology and Communications in Health (ITCH 2015), February – March 2015.

- Sean Smith, Ross Koppel, Jim Blythe and Vijay Kothari. "Mismorphism: A Semiotic Model of Computer Security Circumvention," Technical Report TR2015-768, Dartmouth College, March 2015.

- Vijay Kothari, Jim Blythe, Sean Smith and Ross Koppel. "Measuring the Security Impacts of Password Policies Using Cognitive Behavioral Agent Based Modeling," Symposium and Bootcamp on the Science of Security (HotSoS), April 2015.

- John C. Mace, Charles Morisset, and Aad van Moorsel. "Modelling User Availability in Workflow Resiliency Analysis," Symposium and Bootcamp on the Science of Security (HotSoS), April 2015.

- Phuong Cao, Eric Badger, Zbigniew Kalbarczyk, Ravishankar Iyer, Alexander Withers and Adam Slagell. "Towards a Unified Security Testbed and Security Analytics Framework," Symposium and Bootcamp for the Science of Security (HotSoS), April 2015.

- Sean Smith, Ross Koppel, Jim Blythe and Vijay Kothari. "Mismorphism: A Semiotic Model of Computer Security Circumvention," Symposium and Bootcamp on the Science of Security (HotSoS), April 2015.

- T. Xie, J. Bishop, N. Tillmann and J. de Halleux. "Gamifying Software Security Education and Training via Secure Coding Duels in Code Hunt," Symposium and Bootcamp on the Science of Security (HotSoS), April 2015.

- Phuong Cao, Eric Badger, Zbigniew Kalbarczyk, Ravishankar Iyer and Adam Slagell. "Preemptive Intrusion Detection: Theoretical Framework and Real-World Measurements," Symposium and Bootcamp for the Science of Security (HotSoS), April 2015.

- Wei Yang, Xusheng Xiao, Benjamin Andow, Sihan Li, Tao Xie, and William Enck. "AppContext: Differentiating Malicious and Benign Mobile App Behavior Under Context," 37th International Conference on Software Engineering (ICSE 2015), Florence, Italy, May 2015.

- Wenxuan Zhou, Matthew Caesar, Brighten Godfrey, and Dong Jin. "Enforcing Generalized Consistency Properties in Software-Defined Networks," 12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2015), May 2015.

- Zhenqi Huang, Chuchu Fan, Alexandru Mereacre, Sayan Mitra, and Marta Kwiatkowska. "Simulation-based Verification of Cardiac Pacemakers with Guaranteed Coverage," Appeared in a Special Issue of the IEEE Design and Test, June 2015. DOI 10.1109/MDAT.2015.2448543

continued on page 43

UNIVERSITY OF MARYLAND

PI JONATHAN KATZ

The newest lablet added in 2014, the University of Maryland (UMD) Lablet, led by (PI) Jonathan Katz currently has 20 faculty, including 15 at UMD from the departments of Computer Science, Electrical and Computer Engineering, Information Science, Criminology, and Reliability Engineering. In addition, there are five collaborators at other universities.

The ten projects currently underway support more than 15 PhD students and have generated more than a dozen publications. These efforts include workshops on data-driven approaches to security and privacy. Their strengths are human behavior and policy-governed collaboration and security metrics.

FUNDAMENTAL RESEARCH

The UMD Lablet involves several projects looking at different aspects of the five hard problems.

Verification of Hyperproperties

Michael Clarkson, Cornell University, and Michael Hicks, UMD, are attacking the problem of compositional security by trying to develop a verification methodology based on hyperproperties, a generalization of the classical notion of properties. This methodology, if successful, would enable verification that software systems satisfy security policies, thus providing predictable security and increasing the trustworthiness of software. It would, in particular, allow program components to be analyzed independently and then securely composed to build larger systems in a scalable fashion. They have recently been working on an automated verification methodology for security. In this methodology, security policies are expressed as logical formulas, and a model checker verifies those formulas. The formulas are expressed in a new logic named HyperLTL, which generalizes linear-time temporal logic (LTL). A paper on this work was published at the 3rd Conference on Principles of Security and Trust (POST 2014).

Trustworthy and Composable Software Systems with Contracts

David Van Horn et al. are investigating compositional-verification techniques using language-based mechanisms for specifying and enforcing program properties called contracts. Initial results confirm that behavioral properties of programs can be verified using this approach and they are now trying to scale the approach to cover multi-language programs and security properties. This team recently made a theoretical breakthrough by showing how to efficiently generate counterexamples witnessing contract violations. This is important for testing and debugging software that uses contracts. They have been able to prove that their method is both sound and relatively complete. A paper describing these results has been accepted to the 36th annual ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2015); and prior work, published at the 14th ACM SIGPLAN International Conference on Functional Programming (ICFP'14), was invited to be included in a special issue of the Journal of Functional Programming. The PIs are in the process of launching an interactive web service for experimenting with their system at http://scv.umiacs.umd.edu, and plan to advertise the service widely to the Programming Languages Community in the coming months.

Empirical Models for Vulnerabilities and Attacks

Tudor Dumitras et al. are working to design more-informative metrics to quantify security of deployed systems. This work addresses the hard problem of developing quantifiable metrics for assessing the security of systems, and understanding how those metrics evolve in the real world. The research team has 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 evaluated these metrics on nearly 300 million reports of intrusion-protection telemetry, collected on more than six million hosts. They have found several interesting results so far, including: (1) The exploitation ratio and the exercised attack surface tend to decrease with newer product releases. (2) Hosts that quickly upgrade to newer product versions tend to have a reduction in exercised attack surfaces. (3) Quantitative improvements with respect to the metrics they have studied are often associated with the introduction of new security technologies, thus demonstrating the effectiveness of those technologies in the field. A paper on this work was published at The 17th International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2014).

Human Behavior and Cyber Vulnerabilities

V. S. Subrahmanian et al. are using an empirical approach to study factors that affect the rate at which security patches are deployed. A longer-term goal is to correlate this data with sociological studies of network administrators, to determine why patches are not deployed more quickly. This work addresses, in part, the hard problem of developing quantifiable metrics for assessing the security of systems, and understanding how those metrics evolve in the real world. As an example, the PIs carried out a large-scale measurement of security-patch deployment. They collected a corpus of daily patch-deployment measurements for 1,593 vulnerabilities from 10 popular client applications that are difficult to monitor through network scanning, and are often targeted in spear-phishing attacks. The team found that for 77 percent of the vulnerabilities analyzed, patching started within 7 days of the disclosure date. But they also observed that none of the applications considered, except for the Chrome browser (which employs automated updates), were able to reach 90 percent of the vulnerable host population for more than 90 percent of the patches released during the 5-year observation period. A paper describing this work was presented at the IEEE Symposium on Security and Privacy 2015.

Does the Presence of Honest Users Affect Intruder Behavior?

Michel Cukier and David Maimon are applying a criminological viewpoint to develop a better understanding of attackers' behavior. Using honeypots deployed at the University of Maryland, they are studying how different system-level aspects affect intruders' behavior. The main accomplishment of this quarter is a submission to DSN reporting on their study using data collected over 31 months from two target computer configurations: one presenting attackers with a warning banner, and one that does not. They observed significant ($p<0.05$) differences based on country-of-origin of the attacker and also found that the use of a banner altered behavior originating in China.

User-Centered Design for Security

Adam Aviv, United States Naval Academy, and Jennifer Golbeck, UMD, are focusing on using empirical studies (surveys) to understand users' perceptions of security and usability. The overarching goal is to apply what they learn to predict user perceptions, and to use those predictions to design better policies, better user interfaces, and more-secure systems generally. This would enable the design of systems in which users' perceptions of security match some known metric of security, thus inducing security by design. In one recent work, they have studied perceptions of security and usability for Android's graphical password mechanism. They found that users' perceptions of security are unaffected by spatial shifting, but greatly affected by "complexity." Most surprisingly, they were able to predict perceptions and found that none of the tested features alone impacted perceptions, but rather the total length of the password was the most predictive of security perceptions. A paper on this topic was accepted to Annual Computer Security Applications Conference (ACSAC) 2014.

Understanding Developers' Reasoning about Privacy and Security

Katie Shilton et al. have begun undertaking qualitative studies of users and developers in an effort to discover factors that encourage or discourage privacy and security by design. This work is directed at the broader goal of understanding human behavior and its impact on security. They have continued interviews with mobile application developers focused on cultural and workplace dynamics, and these are expected to progress over the course of the coming year. They are also working on contextual privacy software whose goal is to make information sharing more transparent and user-friendly.

Trust, Recommendation Systems, and Collaboration

John S. Baras and Jennifer Golbeck are studying the fundamental notion of trust, and seeking to develop appropriate models that can be applied to study the dynamics of small groups of parties exploring mechanisms for collaboration based on their local policies. They have used game theory to characterize the costs and benefits of collaboration as a function of the level of trust, and have proved formally the conjecture that "trust is a lubricant for cooperation." This work directly addresses the hard problem of policy-governed secure collaboration, among others.

Reasoning about Protocols with Human Participants

Jonathan Katz and Poorvi Vora, George Washington University, have adapted a protocol for remote electronic voting based on physical objects like scratch-off cards. What is particularly novel here is that the human voter is explicitly modeled as a participant in the protocol, taking into account limitations on the kinds of computations humans can be expected to perform. In this sense, this work related to the general problem of modeling human behavior and appropriately taking human behavior into account when designing security protocols.

COMMUNITY ENGAGEMENTS

Tudor Dumitras and Jonathan Katz organized a two-day workshop on "data-driven security" to which all lablet members were invited. Researchers from outside the lablet also attended. See http://www.umiacs.umd.edu/~tdumitra/data-driven for further information.

Jennifer Golbeck and co-authors submitted a survey paper to IEEE Security & Privacy, covering the "human behavior" hard problem.

David Van Horn's results on contract verification were presented at the National Institute of Informatics special meeting on "Software Contracts for Communication, Monitoring, and Security," as well as a Dagstuhl seminar on analysis and verification. The results were also presented at ACM SIGPLAN International Conference on Functional Programming (ICFP). The presentation of the ICFP 2014 paper was recorded and made available on YouTube; it has been circulated among the programming languages community.

Elaine Shi, Cornell University, co-organized The National Science Foundation Secure and Trustworthy Cyberspace Principal Investigators (NSF SATC PI) meeting that took place in January 2015. She is also serving as program co-chair for The International Conference on Information, Communications and Signal Processing (ICICS), and the ACM Symposium on InformAtion, Computer and Communications Security The Fourth International Workshop on Security in Cloud Computing (AsiaCCS-SCC).

Adam Aviv is serving as program co-chair for The Workshop on Cyber Security Evaluation and Testing (CSET).

Jonathan Katz served as program chair for The IACR International Conference on Practice and Theory of Public-Key Cryptography (PKC), and was also on the program committee for the "2015 IEEE Security & Privacy Conference."

A paper by Elaine Shi and Michael Hicks was selected as the winner of the NSA's second annual "Best Scientific Cybersecurity Paper Competition."

EDUCATIONAL

Michael Hicks ran a "Build It, Break It, Fix It" contest that drew the interest of students from 61 universities across the US, including 39 students from the University of Maryland. In contrast to security competitions, which encourage teams only to attack systems and find vulnerabilities, the goal of this contest is to encourage writing of secure programs in the first place.

Michael Hicks is serving as program chair for the 2015 Computer Security Foundations Workshop, and is also on the program committee for the "2015 IEEE Symposium for Security & Privacy." He also serves on the Institute for Defense Analyses/ Center for Computing Sciences (IDA/CCS) program review committee. He has been blogging about programming-language security at http://www.pl-enthusiast.net/.

Michael Hicks, Jennifer Golbeck, and Jonathan Katz are offering massive open online courses (MOOCs) in computer-security on Coursera. These courses will cover programming-language security, cryptography, and usable security.

Adam Aviv sponsored a female high-school student for an

8-week internship. He is also developing a senior-level elective on cybersecurity, as well as one focusing on usable security.

David Van Horn has incorporated his work into his graduate class on "Program Analysis and Understanding." He will also work to incorporate this into the pedagogically-oriented programming environment accompanying his textbook, How to Design Programs. He was invited to lecture about his work at a graduate summer school at the University of Utah in July.

Katie Shilton has developed a module focusing on the human side of implementing security that will be incorporated into two courses: a Masters-level course on "Policy Issues on Digital Curation," and a post-graduate course on the same topic. She will be incorporating her findings from her current research into this module. A project based on the "Bubbles" work is being incorporated into a mobile-development undergraduate course offered this semester.

Tudor Dumitras is offering a course on distributed-systems security that incorporates a discussion of security metrics and empirical studies of security properties.

Michel Cukier leads the Advanced Cybersecurity Experience for Students (ACES) undergraduate honors program in cybersecurity, which incorporates a holistic approach to cybersecurity covering technical, policy, and behavioral aspects of the problem.

PUBLICATIONS

- Yabing Liu, Will Tome, Liang Zhang, David Choffnes, Dave Levin, Bruce Maggs, Alan Mislove, Aaron Schulman, Christo Wilson. 2015. "On SSL Certificate Revocation: The Race to the Bottom in Securing the Web's PKI." Submitted to ACM Internet Measurement Conference (IMC).
- Phuc C. Nguyen and David Van Horn. 2015. "Relatively Complete Counterexamples for Higher-Order Programs." PLDI 2015, the 36th Annual ACM SIGPLAN International Conference on Programming Language Design and Implementation. http://dx.doi.org/10.1145/2737924.2737971
- Adam J. Aviv and Dane Fichter. 2014. "Understanding Visual Perceptions of Usability and Security of Android's Graphical Password Pattern." In Proceedings of the 30th Annual Computer Security Applications Conference (ACSAC '14).
- Andrew G. West and Adam J. Aviv. "Measuring Privacy Disclosures in URL Query Strings." In Internet Computing, IEEE, 18(6): 52-59, 2014.
- Andrew G. West and Adam J. Aviv. "On the Privacy Concerns of URL Query Strings." In W2SP '14: Proceedings of the Workshop on Web 2.0 Security and Privacy, May 2014.
- Jeanne Luning-Prak and Adam J. Aviv. "A Self-Reporting Survey of Android's Unlock Password." Poster presented at the 30th Annual Computer Security Applications Conference (ACSAC '14), 2014.
- Krontiris, I., Langheinrich, M. & Shilton, K. 2014. "Trust and Privacy in Mobile Experience Sharing – Future Challenges and Avenues for Research." In IEEE Communications, August 2014. http://cps-vo.org/node/17109
- Martin, K. and Shilton, K. "Why Experience Matters to Privacy: How Context-Based Experience Moderates Consumer Privacy Expectations for Mobile Applications." Journal of the Association for Information Science & Technology. Forthcoming 2015.
- Antonio Nappa, Richard Johnson, Leyla Bilge, Juan Caballero, and Tudor Dumitras. "The Attack of the Clones: A Study of the Impact of Shared Code on Vulnerability Patching." In IEEE Symposium on Security and Privacy, San Jose, CA, May 2015.
- Elissa M. Redmiles, Amelia Malone , and Michelle L. Mazurek (UMD). "How I Learned to Be Secure: Advice Sources and Personality Factors in Cybersecurity." Poster presented at the Symposium on Usable Privacy and Security (SOUPS), July 2015.

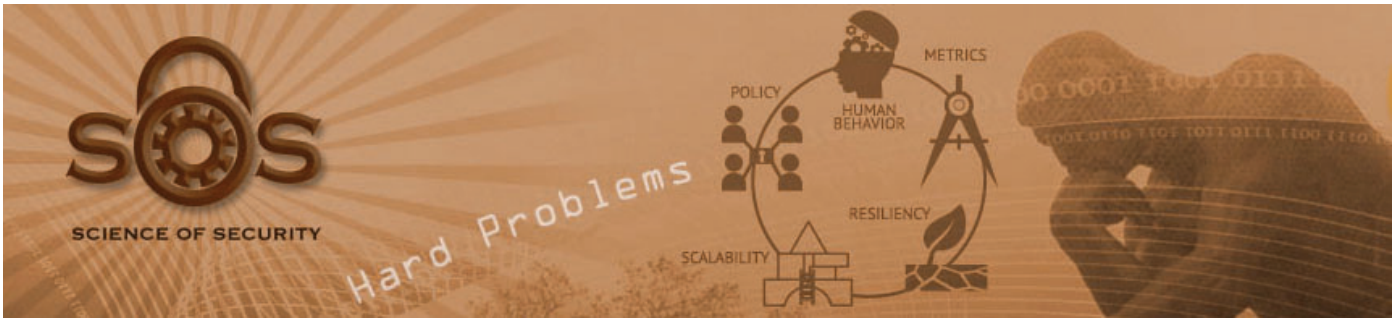

Status of Hard Problems

The National Security Agency (NSA) sponsors the Science of Security (SoS) Initiative to promote the underpinning of cybersecurity with a discipline built on scientific foundations. A component of this initiative funds small, multi-disciplinary labs at four universities, Carnegie Mellon, North Carolina State, Illinois-Urbana Champaign, and Maryland-College Park, to discover foundational research, promote rigorous scientific principles, and grow a SoS community. In 2012 as a measure to establish the beginnings of a common language and gauge progress, the lablet Principal Investigators (PIs) developed, in collaboration with NSA, five Hard Problems in security needing science to advance. A paper defining the hard problems was publicly released in November 2012.[1]

These five are not all encompassing of cybersecurity, but are specific challenges for cybersecurity that need research for advancement, and the lablet research is focused in these areas. The five Hard Problems are Scalability and Composability; Policy-Governed Secure Collaboration; Security-Metrics-Driven Evaluation, Design, Development, and Deployment; Resilient Architectures; and Understanding and Accounting for Human Behavior. The following details how SoS Lablet (SoSL) research had improved the challenges of the hard problems.

[1] The five Hard Problems paper can be found on the Science of Security Virtual Organization Webpage: http://www.sos-vo.org.

The hard problem of Scalability and Composability involves techniques for constructing and analyzing software that scale to large systems, and furthermore, that allow analysis of the entire system to proceed by analyzing its parts independently. The two properties are closely related, because in many cases compositionality is the key to achieving scalability in practice. Lablet research has produced fundamental advances in our ability to perform compositional security-related analyses on both software and data.

- Before our work, there was no way to compositionally guarantee safety when executing unknown code provided by an adversary. We developed a theory of compositional security called "adversary-aware assume guarantee," developed an accompanying program logic that allows compositional proofs of safety for programs that execute adversary-supplied code, and applied the approach to hypervisors and trusted computing systems. Relevant publications:
    - o L. Jia, S. Sen, D. Garg, A. Datta, "A Logic of Programs with Interface-confined Code," in Proceedings of 28th IEEE Compute Security Foundations Symposium, July 2015
    - o A. Vasudevan, S. Chaki, L. Jia, J. McCune, J. Newsome, A. Datta, "Design, Implementation, and Verification of an eXtensible and Modular Hypervisor Framework," in Proceedings of 34th IEEE Symposium on Security and Privacy, May 2013.
    - o A. Datta, J. Franklin, D. Garg, L. Jia, D. Kaynar, "On Adversary Models and Compositional Security," IEEE Security and Privacy 9 (3): 26-32 (2011) (Special Issue on the Science of Security).
- Before our research, integrating syntax from two languages within the same program was difficult to do compositionally; it was only possible by following arcane and/or verbose syntactic conventions. As a result, developers in today's programming systems tend to use strings to write domain-specific syntax (e.g., for SQL queries) within a host language, leading to command injection attacks. We developed a way to compose domain-specific syntaxes within a single program, allowing developers to write programs in more natural ways while at the same time mitigating command injection attacks. Relevant publications:
    - o C. Omar, D. Kurilova, L. Nistor, B. Chung, A. Potanin, J. Aldrich, "Safely Composable Type-Specific Languages," in Proceedings of European Conference on Object-Oriented Programming (ECOOP), 2014.
    - o C. Omar, C. Wang, J. Aldrich, "Composable and Hygienic Typed Syntax Macros," in Proceedings of Symposium on Applied Computing (SAC), 2015.
- Before our work, approaches to information flow security did not have a semantic basis for authorizing exceptions to a non-interference property. We use a novel linear epistemic logic to analyze the execution traces of programs and enforce information flow constraints using a type system based on the lax modality. This approach is inherently compositional, and the connection to logic allows us to generate proofs in the authorization policy that provides a justification for violating a too rigid non-interference property. Relevant publication:
    - o A. Ahmad, R. Harper, "An Epistemic Formulation of Information Flow Security," Draft paper under review, 2015.
- Before our research, the best-known ways to do certain attack surface analysis had combinatorial complexity. Our work has found a way to do this analysis with only linear complexity, and we can characterize the scale of the information that is given up in this tradeoff. This work is under preparation for publication in early winter 2016.
- Before our work in the domain of operational coverage-based attack surfaces, it was not clear how to translate static and dynamic analysis of resource-constrained attacks into efficient attack detectors and filters.

- A. Rivers, M. Vouk, L. Williams, "On Coverage-Based Attack Profiles," Fast Abstract in the Software Security and Reliability-Companion (SERE-C), 2014 IEEE Eighth International Conference on, 2014.

- Our work has also improved the scalability of security-based planning and learning algorithms, and applied graph clustering algorithms to detect insider attacks in more scalable ways. Relevant publication:
  - H. Lamba, T. Glazier, B, Schmerl, J. Pfeffer, D. Garlan. 2015. "Detecting Insider Threats in Software Systems using Graph Models of Behavioral Paths," (short paper), 2015 Symposium and Bootcamp on the Science of Security (HotSoS '15).

The hard problem of Policy-Governed Secure Collaboration seeks to develop the science underlying methods for expressing and enforcing normative requirements and policies for handling data with differing usage needs and among users in different authority domains. Over the course of the SoSL efforts, we have deepened the scientific foundations that help us address key limitations of the state of the art.

- Prior to the SoSL efforts, although policy approaches existed that handled authentication and authorization of users for performing data operations based on attribute or role-based credentials, they did not adequately and explicitly characterize the correctness requirements for secure collaboration and their impact on security. We have advanced our understanding of how to express and validate normative requirements that are left implicit in previous research. Specifically, we have developed elements of a formal language that helps capture various subtleties of secure collaboration requirements, including priorities between them. We are developing mathematical models for determining whether those requirements are mutually consistent. We are developing methods to determine whether participants are interacting in a way that complies with the stated requirements or deviates from the requirements only when necessary to satisfy higher priority requirements. Relevant publications:
  - A. Chopra, M. Singh, "Cupid: Commitments in Relational Algebra," in Proceedings of 23rd Conference on Artificial Intelligence (AAAI), 2015.
  - M. Baldoni, C. Baraglio, A. Chopra, M. Singh, "Composing and Verifying Commitment-Based Multiagent Protocols," in Proceedings of 24th International Joint Conference on Artificial Intelligence (IJCAI), 2015.

- Previously, there was inadequate scientific basis for judging if security policies were comprehensible. We have developed a model of policy complexity (as perceived by humans), yielding a modular form for firewall policies that we have empirically shown helps improve user understanding.

- Previously, there was inadequate study of how security policies are developed and whether they capture stakeholder requirements. We have empirically evaluated a method for capturing design rationales for security policies and shown empirically (for firewall policies) that it helps enhance comprehension, thereby improving maintenance and reducing errors. Relevant publications:
  - O. Kafali, N. Ajmeri, M. Singh, "NoRest: Norm Revision for Sociotechnical Systems—a Formal Approach for Secure, Privacy-enhanced Collaboration," unpublished manuscript, 2015.
  - N. Ajmeri, C. Hang, S. Parsons, M. Singh, "Creating and Maintaining Firewall Policies: An Approach based on Argumentation and its Empirical Evaluation," unpublished manuscript, 2015.

- Previous research did not study social architectures requisite for the adoption and enforcement of normative requirements and the creation of flexible trust relationships. We have begun to create models of social architectures in qualitative terms as a basis for further empirical validation with users. Relevant publication:
  - H. Du, B. Narron, N. Ajmeri, E. Berglund,

J. Doyle, M. Singh. 2015. "Understanding Sanction under Variable Observability in a Secure, Collaborative Environment," in Proceedings of 2015 Symposium and Bootcamp on the Science of Security (HotSoS '15).

- Previously, there was inadequate understanding of how security requirements may interact to impact security, and risk scoring methods did not operationalize defense-in-depth, a general theory of security. We have developed a method for computing the extent to which two or more security requirements interact to impact security based on actual analyst threat perceptions; this method operationalizes defense-in-depth.

- Historically, security analysts assess system security using control checklists. These lists, such as the "NIST Special Publication 800-53," describe controls for access control, encryption, auditing, etc., in a largely independent manner. However, defense-in-depth suggests that security controls interact to improve security. Based on our recent work, we discovered that security requirements "compose" together to yield increases and decreases in overall security, and that some requirements provide no additional benefit to security based on assumptions about threats and use cases. Relevant publication:

    - H. Hibshi, T. Breaux, S. Broomell, "Assessment of Risk Perception in Security Requirements Composition," IEEE 23rd International Requirements Engineering Conference (RE'15), 2015.

The hard problem of Security Metrics involves techniques for effectively measuring and quantifying the extent to which a given system satisfies a particular set of security properties. Lablet work in this area is making progress on several aspects of the problem: for example, performing statistical analyses of real-world datasets to understand and quantify factors leading to vulnerabilities and exploits; developing metrics (whether software-, network-, or system-based) to predict vulnerabilities and/or measure effectiveness of countermeasures; and measuring humans' perceptions of various security measures. Highlights of this work include the following:

- Before our work, one could only speculate about the real-world effect of new security technologies that have been introduced. Empirical studies using the WINE dataset have been used to give evidence as to how the cyber threat landscape has changed following the introduction of various security technologies. This work found, for example, evidence that exploits decreased following the introduction of sandboxing techniques in Windows Internet Explorer 7 (IE7) and Adobe Reader 10. On the other hand, the overall field discovery rate of security problems for open source operating systems and applications, such as the Fedora platform set, does not seem to have changed over a long time (20 or so releases). It is low (orders of magnitude below non-security problems), but persistent and almost constant. Furthermore, it would appear that programmers are making the same mistakes (e.g., those from the top 25 most frequent and dangerous security problems) over and over again, i.e., overall software development processes do not appear to be improving. Relevant publications:

    - K. Nayak, D. Marino, P. Efstathopoulos, T. Dumitras, "Some Vulnerabilities are Different than Others: Studying Vulnerabilities and Attack Surfaces in the Wild," International Symposium on Research in Attacks, Intrusions and Defenses (RAID), Gothenburg, Sweden, 2014.

    - S. Subramani, M. Vouk, L. Williams, "Non-Operational Testing of Software for Security Issues," in Software Reliability Engineering Workshops (ISSREW), 2013 IEEE International Symposium on, 2013.

    - D.Y. Lee, M. Vouk, L. Williams, "Using Software Reliability Models for Security Assessment—Verification of Assumptions," in Software Reliability Engineering Workshops (ISSREW), 2013 IEEE International Symposium on, 2013.

    - S. Subramani, M. Vouk, L. Williams. 2014. "An Analysis of Fedora Security Profile," in Proceedings of the 2014 Symposium and Bootcamp on the Science of Security (HotSoS '14).

- S. Subramani, "Security Profile of Fedora," MS Thesis, 2014.

- A. Vangaveeti, "An Assessment of Security Problems in Open Source Software," MS Thesis, 2015.

- Before our work, there were hundreds of disparate publications about intrusion-detection systems, each with varying methods and evaluation approaches. Our work has led to a taxonomy to compare those studies and to systematize that knowledge. We also examined the security of different cloud-layers and developed a taxonomy of vulnerabilities and related countermeasures, and applied that information in the context of Hadoop running in a cloud environment. Relevant publications:

    - Y. Huang, X. He, H. Dai. 2015. "Poster: Systematization of Metrics in Intrusion Detection Systems," in Proceedings of 2015 Symposium and Bootcamp on the Science of Security (HotSoS '15).

    - X. He, H. Dai, P. Ning, R. Dutta, "Dynamic IDS Configuration in the Presence of Intruder Type Uncertainty," IEEE Global Conference on Communications (GLOBECOM), 2015.

    - D. Kim, M. Vouk, "A Survey of Common Security Vulnerabilities and Corresponding Countermeasures for SaaS," GLOBECOM workshop on Cloud Computing Systems, Networks and Applications (CCSNA), 2014.

    - D. Kim, H. Schaffer, M. Vouk, "About PaaS Security," in Proceedings of 3rd International IBM Cloud Academy Conference (ICACON), 2015.

    - X. Yu, P. Ning, M. Vouk. 2014. "Securing Hadoop in Cloud," 2014 Symposium and Bootcamp on the Science of Security (HotSoS '14).

    - Y. Xianqing, P. Ning, M. Vouk, "Enhancing Security of Hadoop in a Public Cloud," in Proceedings of 6th International Conference Information and Communication Systems (ICICS), 2015.

- Before starting this project, only a few approaches existed for analyzing the attack surface of software systems. Those approaches were more at the system level and did not provide actionable feedback to software engineers as they develop code. Today, we have discovered that there are metrics that can predict vulnerabilities at the method level. By simulating the sequence of actions that an attacker might take when exploiting a vulnerability, we are able to estimate the areas of a system where vulnerabilities are likely to be found. Recent results from a large, open-source project show that our metrics increase just before a vulnerability is found and decrease after a vulnerability is fixed, giving empirical evidence that our new metrics are useful predictors of vulnerabilities. Relevant publication:

    - C. Theisen, K. Herzig, P. Morrison, B. Murphy, L. Williams, "Approximating Attack Surfaces with Stack Traces," International Conference on Software Engineering (ICSE), 2015.

- Before starting the project, the statistical distribution of cyber-attack events was not fully understood. Our work shows that many attacks have a hypergeometric character driven by resource and schedule constraints' issues, and that this information can be used to build efficient high-probability attack detection sensors.

    - A. Rivers, M. Vouk, L. Williams, "On Coverage-Based Attack Profiles," Fast Abstract in the Software Security and Reliability-Companion (SERE-C), 2014 IEEE Eighth International Conference on, 2014.

- Before our work, models to predict the presence of vulnerabilities and the resilience of systems were not accurate enough to make them actionable by practitioners. We have made progress in developing new metrics that can be used to more accurately evaluate the probability that a given host is vulnerable and, if so, whether it might be exploited. We have also confirmed that appropriately tuned "classical" reliability models can be used to assess "security reliability" of open source software, including prediction of the rate and numbers of field discoverable security problems in the follow-up releases of the software. Relevant publication:

    - K. Nayak, D. Marino, P. Efstathopoulos, T. Dumitras, "Some Vulnerabilities are Different than Others: Studying Vulnerabilities and Attack Surfaces in the Wild," International Symposium on Research in Attacks, Intrusions and Defenses (RAID), Gothenburg, Sweden, 2014.

- Before our work, it was unknown how quickly software patches were applied in the real world, or what techniques would be most beneficial for incentivizing faster patching. Lablet work has empirically demonstrated (using both the WINE dataset as well as network measurements on PKI revocation data following the Heartbleed incident) that software patches for known vulnerabilities are either not applied in a timely fashion, or are applied incorrectly. Beyond characterizing the rate of software patching, the work also aims to determine factors that influence this rate. Relevant publication:

    - Nappa, R. Johnson, L. Bilge, J. Caballero, T. Dumitras, "The Attack of the Clones: A Study of the Impact of Shared Code on Vulnerability Patching," in Security and Privacy (SP), 2015 IEEE Symposium on , vol., no., pp. 692-708, 17-21 May 2015.

- Other lablet work is investigating techniques for representing real-world security incidents, and designing sound methods for preemptive detection of such events. This work looks at relations between available evidence (e.g., a sequence of observable events) and hidden system states to determine the most probable sequence of state transitions consistent with the evidence. This, in turn, is used to automatically assess, based on the evidence, whether a system or a user account has been compromised. The resulting tool was validated with real-world incidents collected over a six-year period. Relevant publications:

    - P. Cao, E. Badger, Z. Kalbarczyk, R. Iyer, A. Slagell. 2015. "Preemptive Intrusion Detection: Theoretical Framework and Real-World Measurements," 2015 Symposium and Bootcamp on the Science of Security (HotSoS '15).

- Finally, lablet researchers are looking at quantifying users' perceptions of security, with current research focused on the graphical password system used by Android smartphones. By understanding what makes users perceive something as secure, researchers hope to design systems whose actual security aligns with those perceptions, thus influencing better choices on the part of users. Relevant publication:

    - A. Aviv, D. Fichter. 2014. "Understanding Visual Perceptions of Usability and Security of Android's Graphical Password Pattern," In Proceedings of the 30th Annual Computer Security Applications Conference (ACSAC '14).

The hard problem of Resilient Architectures has several attributes, with different emphases depending on the context and community. One attribute of resilience captures the notion of robustness. That is, the ability of the system to statically withstand attack, e.g., through diversity in implementation. Another attribute captures the notion that a system can continue to deliver essential services (albeit potentially at a diminished level) in the midst of an attack. Yet another attribute stresses how quickly a system can be restored to full functionality following an attack.

- Before our work began, means to specify resiliency properties and requirements were not sufficiently precise or detailed to serve as a basis for rigorous systems engineering. We have developed a formal mathematical framework to enable more precise specification of the full range of properties of affordability, reliability, availability, safety, usability, scalability, evolvability, and resilience.

- Previously, network policy enforcement was primarily deployed statically at network ingress points. We have now found that a top-down strategy can be used to deploy policy enforcement across a network with greater efficiency and scalability than traditional, ingress-only deployments. This supports policy enforcement that can better absorb and adapt to adversarial traffic patterns. Relevant publications:

    - V. Heorhiadi, S. Fayaz, M. Reiter, V. Sekar, "SNIPS: A Software-Defined Approach for Scaling Intrusion Prevention Systems via Offloading," in Information Systems Security, 10th International Conference on, ICISS 2014.

    - V. Heorhiadi, M. Reiter, V. Sekar, "Accelerating the Development of Software-Defined Network Optimization Applications using SOL," CoRR: ACM Computing Repository, abstract, 2015.

- Previously, the properties and tradeoffs among different security isolation techniques had not been made explicit. Based upon a literature survey, we have found that the security isolation problem comprises a large design space consisting of many design dimensions. Existing approaches fall short in terms of adaptability and measurability. Relevant publication:
    - R. Shu, P. Wang, S. Gorski, B. Andow, A. Nadkarni, L. Deshotels, J. Gionta, W. Enck, X. Gu, "A Systematic Study of Security Isolation," in submission to ACM Computing Surveys (CSUR).
- Previously, there existed no way to reason about how robust a cyber-physical system might be to disruption (robustness being one attribute of resilience). Understanding was particularly lacking in how to approach reasoning about how an attack on the cyber component might be effected by manipulation of the physical component. Lablet research has broken ground in developing practical mathematical frameworks that support this reasoning, within the context of a hybrid system model that conjoins the continuous control description that interacts with the physical realm, and the cyber component by which the control is implemented. Using this framework we have developed algorithms that measure bounds on "how close" a physical disturbance can push a cyber-physical system near deleterious states. Relevant publications:
    - Z. Huang, Y. Wang, S. Mitra, G. Dullerud, "Controller Synthesis for Linear Time-varying Systems with Adversaries," 2015. URL: http://arXiv:1501.04925v1 [cs.SY]
    - Z. Huang, C. Fan, A. Mereacre, S. Mitra, M. Kwiatkowska, "Simulation-based Verification of Cardiac Pacemakers with Guaranteed Coverage," Special Issue of IEEE Design and Test, 2015.
- Previously, analytic approaches to resilience emphasized the robustness attribute of resiliency. We have extended the foundations of analyzing resilient architectures by providing a reasoning approach to self-protection that uses stochastic multiplayer games to reason about human involvement, self-protection latency, and uncertainty, accommodating a much more dynamic viewpoint of resiliency.
- Previously, it was not clear how some of the existing high-assurance resilience methods, such as run-time fault-tolerance through back-to-back assessment, could be used to detect and counter security vulnerabilities and attacks. We show, using some of the recent vulnerabilities, breaches and attacks (such as Heartbleed) that back-to-back testing and run-time comparison-based analysis can discover a very large fraction of such issues including zero-day attacks. This opens the door to development of off-the-shelf security sensors for clouds, IoT, and stand-alone services. Relevant publications:
    - R. Venkatakrishnan, M. Vouk. 2014. "Diversity-based Detection of Security Anomalies," In Proceedings of the 2014 Symposium and Bootcamp on the Science of Security (HotSoS '14).
    - R. Venkatakrishnan, "Redundancy-Based Detection of Security Anomalies in Web-Server Environments," MS Thesis, 2014.
    - R. Venkatakrishnan, M. Vouk, "Using Redundancy to Detect Security Anomalies—Towards IoT Security Attack Detectors," ACM Ubiquity, to appear, 2015.

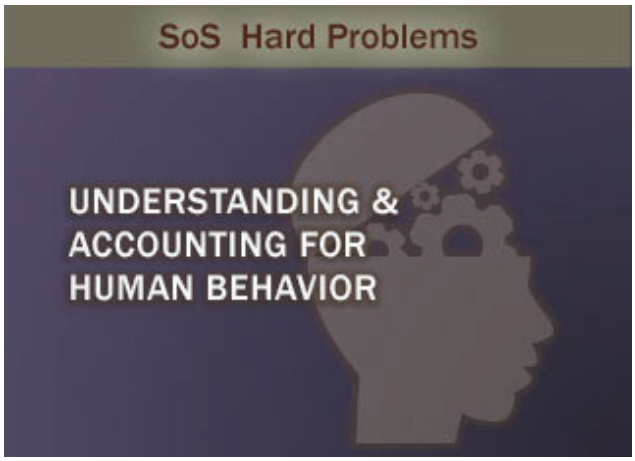

The hard problem of Human Behavior aids in the handling of the unpredictability of human actors in the security of systems. Computers do what we tell them to do, but humans do what they want to do, adding unpredictability and complexity to the design and implementation of computer systems. A variety of research projects were dedicated to developing models and insights of human behaviors that enable the design, modeling, and analysis of systems with specified security properties.

- The Security Behavior Observatory is giving us unique insights into the human behaviors that lead to security vulnerabilities on personal computers. For example, we are able to trace the steps that our participants have taken that have led to malware infections and to participants uninstalling anti-virus software.

- Before we began our work, security analysts nominally scored security requirements (e.g., control lists) to determine their independent impact on security. Today, we have a new method for computing the extent to which two or more security requirements interact to impact security based on actual analyst threat perceptions. Our method implements defense-in-depth, a general theory of security that had been spoken to, but not operationalized by, traditional risk scoring methods.

- Before our project began, there was ample speculation but little empirical evidence about how or why type-like specifications affect the productivity and accuracy of human programmers. We studied this question in the setting of application programming interfaces (APIs) that define protocols of interaction that API clients must follow—a problem with significant security and reliability implications in practice. Using laboratory studies and experiments, we found that leveraging type-like protocol specifications documentation can increase programmer productivity and reduce programmer errors, both by a factor of 2 or more, when developers are using particularly challenging APIs.

- Before our work began, researchers knew that phishing scams were problematic but they did not understand why people fell prey to them. We have now found that social factors such as trust and cognitive factors such as attention and impulsiveness influence the likelihood of falling prey to social engineering when phishing emails are received. We have worked to classify the message content of hundreds of archived phishing emails to determine how classical persuasion research predicts the likelihood of data loss. Our goal is to understand how social engineering occurs and to develop techniques to combat these tactics being used by cyber-criminals. Relevant publications:

    - C. Mayhorn, E. Murphy-Hill, O. Zielinska, A. Welk, “The Social Engineering Beyond Phishing,” The Next Wave, 2015.

    - A. Welk, O. Zielinska, R. Tembe, G. Xe, K. Hong, E. Murphy-Hill, C. Mayhorn, “Will the Phisher-men” Reel You in? Assessing Individual Differences in a Phishing Detection Task,” International Journal of Cyber Behavior, Psychology and Learning, 2015 (in press).

- Before our project began, biometrics research had found ways of processing low-level interaction with a computer—the mouse movements and the keys typed on a keyboard—to authenticate individual users, potentially as a substitute for typing passwords. Our work took on a related question: could such interaction patterns be used to distinguish ordinary users from malicious users, behaving deceptively? The algorithms we have developed can detect patterns that co-occur with different usage motivations, in the laboratory settings we have evaluated. These algorithmic models, which are informed by work in cognitive psychology, have accuracy above 80% to 90% in our experimental user studies. Relevant publications:

    - T. Barik, A. Chakraborty, B. Harrison, C. Roberts, R. St. Amant. 2013. “Modeling the Concentration Game with ACT-R,” The 12th International Conference on Cognitive Modeling (ICCM), 2013.

- R. St. Amant, P. Goodwin, I. Dominguez, D. Roberts. 2015. “Toward Expert Typing in ACT-R,” The 13th International Conference on Cognitive Modeling (ICCM), 2015.
- A. Chakraborty, B. Harrison, P. Yang, D. Roberts, R. St. Amant. 2014. “Exploring Key-Level Analysis for Computational Modeling of Typing Behavior,” 2014 Symposium and Bootcamp on the Science of Security (HotSoS ’14).
- I. Dominguez, A. Goel, D. Roberts, R. St. Amant. 2015. “Detecting Abnormal User Behavior through Pattern-mining Input Device Analysis,” 2015 Symposium and Bootcamp on the Science of Security (HotSoS ’15).

PUBLICATIONS

continued from page 29

- Z. J. Estrada, C. Pham, F. Deng, Z. Kalbarczyk, R. K. Iyer, L. Yan. “Dynamic VM Dependability Monitoring Using Hypervisor Probes,” to appear at 11the European Dependable Computing Conference- Dependability in Practice, EDCC 2015, Paris, France, Sept. 7-11, 2015.
- Jiaqi Yan and Dong Jin. “VT-Mininet: Virtual-time-enabled Mininet for Scalable and Accurate Software-Define Network Emulation,” ACM SIGCOMM Symposium on SDN Research 2015 (SOSR15), Santa Clara, CA, June 2015.
- Jiaqi Yan and Dong Jin, “A Virtual Time System for Linux-container-based Emulation of Software-defined Networks.” ACM SIGSIM Conference on Principles of Advanced Discrete Simulation, London, UK, June 2015 (Finalist for the Best Paper Award)
- Ning Liu, Adnan Haider, Xian-He Sun and Dong Jin. “FatTreeSim: Modeling a Large-scale Fat-Tree Network for HPC Systems and Data Centers Using Parallel and Discrete Event Simulation.” ACM SIGSIM Conference on Principles of Advanced Discrete Simulation, London, UK, June 2015. Best Paper Award
- Ning Liu, Xian-He Sun and Dong Jin. “On Massively Parallel Simulation of Large-Scale Fat-Tree Networks for HPC Systems and Data Centers,” Poster presented at ACM SIGSIM Conference on Principles of Advanced Discrete Simulation, London, UK, June 2015. Best Poster Award.
- Anduo Wang, Fan Yang, Mangesh Bendre, Brighten Godfrey, and Matthew Caesar. “Ravel: Orchestrating Software-Defined Networks,” software demo at ACM SIGCOMM Symposium on SDN Research (SOSR), June 2015.

Science of SecUrity and REsilience for Cyber-Physical Systems (SURE)

The project on the System Science of SecUrity and REsilience for cyber-physical systems (SURE) is developing foundations and tools for designing, building, and assuring cyber-physical systems (CPS) that can maintain essential system properties in the presence of adversaries. The technology base of SURE will provide CPS designers and operators with models, methods, and tools that can be integrated with an end-to-end model-based design flow and tool chain. SURE is an NSA-funded project aimed at improving scientific understanding of resiliency, described as having the attributes of functional correctness by design, robustness to reliability failures or faults, and survivability against security failures and attacks. Water distribution and traffic control architectures were offered as examples of the types of cyber physical systems to be examined.

On October 27, 2014 researchers from four universities—Vanderbilt, Hawai'I, California-Berkeley, and MIT—met to kick off the SURE project. According to Xenofon Koutsoukos, Professor of Electrical Engineering and Computer Science in the Institute for Software Integrated Systems (ISIS) at Vanderbilt University, the Principle Investigator (PI) for SURE, "The project aims to equip CPS designers and operators with theory-based comprehensive tools that improve resilience against faults and intrusions, and also enable designers to make security decisions and allocate resources in a decentralized manner." In addition to Professor Koutsoukos as PI, the SURE research team includes Saurabh Amin (MIT), Anthony Joseph (UC Berkeley), Gabor Karsai (Vanderbilt), Dusko Pavlovic (U. of Hawai'I), Larry Rohrbough (UC Berkeley), S. Shankar Sastry (UC Berkeley), Janos Sztipanovits (Vanderbilt), Claire Tomlin (Vanderbilt), Peter Volgyesi (Vanderbilt) Yevgeniy Vorobeychik (Vanderbilt), and Katie Dey (Vanderbilt) - Outreach.

The research problems and questions to be addressed include:

- Risk Analysis and Incentive Design
    - How can the collection of agents in CPS deal with strategic adversaries?
    - How can strategic agents contribute to CPS efficiency and safety, while protecting their conflicting individual objectives?
- Resilient Monitoring and Control
    - What are the control architectures that can improve resilience against intrusions and faults?
    - What types of dynamics can provide inherent robustness against impacts of faults and cyber-attacks?
    - What are the physics-based invariants that can be used as "ground truth" in intrusion detection?
- Decentralized Security
    - How can we design systems that are resilient even when there is significant decentralization of resources and decisions?
- Formal Reasoning about Security in CPS
    - How do we formally and practically reason about secure computation and communication?
- Integrative Research and Evaluation
    - How do we integrate and evaluate cyber & physical platforms and resilient monitoring & control architectures?
    - How do we interface and support human decision makers?

The research challenges facing the team include such problems as spatio-temporal dynamics, multiple strategic interactions with network interdependencies, inherent uncertainties in both public & private systems, and tightly coupled control and economic incentives.

On March 17 and 18, 2015 the SURE team met with members of NSA's R2 Directorate to review their first six months of work. Initially looking at water distribution and surface traffic control architectures, air traffic control and satellite systems are added examples of the types of cyber physical systems examined. Xenofon Koutsoukos, PI for SURE, indicated the use of these additional cyber physical systems is to demonstrate how the SURE methodologies can apply to multiple systems. Main research thrusts include hierarchical coordination and control, science of decentralized security, reliable and practical reasoning about secure computation and communication, evaluation and experimentation, and education and outreach. The centerpiece is their testbed for evaluation of CPS security and resilience.

The development of the Resilient Cyber Physical Systems (RCPS) Testbed supports evaluation and experimentation across the complete SURE research portfolio. This platform is being used to capture the physical, computational and communication infrastructure; describe the deployment, configuration of security measures and algorithms; and provides entry points for injecting various attack or failure events. "Red Team" vs "Blue Team" simulation scenarios are being developed. After the active design phase—when both teams are working in parallel and in isolation—the simulation is executed with no external user interaction, potentially several times. The winner is decided based on scoring weights and rules that are captured by the infrastructure model.

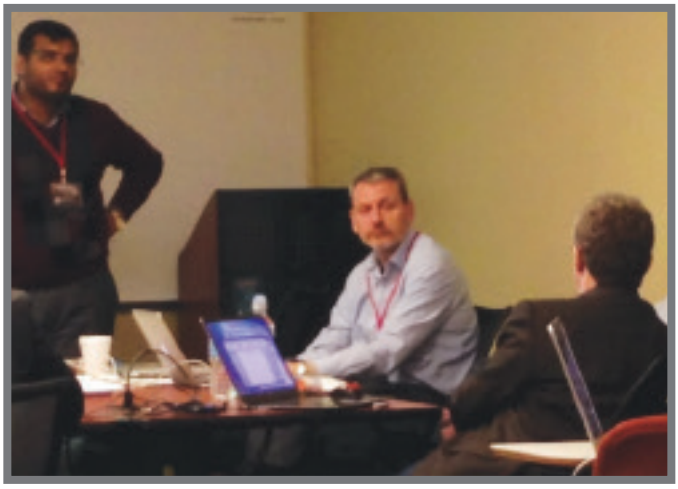

Waseem Abbas and lead PI Xenophon Koutsoukos, Vanderbilt, listen to comments about resilient sensor designs from David Corman, National Science Foundation.

In addition to the testbed, ten research projects on resiliency were presented. These presentations covered both behavioral and technical subjects including adversarial risk, active learning for malware detection, privacy modeling, actor networks, flow networks, control systems, software and software architecture, and information flow policies. The CPS-VO web site, its scope and format was also briefed.

Project Overview, Xenofon Koutsoukos (Vanderbilt University)
URL: http://cps-vo.org/node/18484
Project Thrustsare Hierarchical Coordination and Control; Risk analysis and incentive design that aim at developing regulations and strategies at the management level; Resilient monitoring and control of the networked control system infrastructure; Science of decentralized security which aims 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 which aims to contribute a formal framework for reasoning about security in CPS; Evaluation and experimentation using modeling and simulation integration of cyber and physical platforms that directly interface with human decision makers; and Education and outreach.

Evaluation Testbed, Peter Volgyesi and Himanshu Neema (Vanderbilt University) URL: http://cps-vo.org/node/18483
The objectives of the RCPS Testbed are to develop and maintain well-defined domains, language, rules, tools, and metrics; integrate existing robust domain tools and technologies, simulators, analysis tools, middleware; maintain model libraries and repositories; Red Team vs Blue Team scenarios and challenges; simulate real adversary behavior; integration technology; meta-programmable tools; strong versioning; web-based interfaces; and cloud-based, scalable services. Current research being conducted on the RCSP Testbed includes complex attack strategies; attack description language; using and orchestrating existing atomic action; adversarial risk analysis; repeated, automated simulation runs; probabilistic interdependency graphs; optimization; resilient monitoring; and control and science of decentralized security.

Demo: Resilient and Secure Component-Based Software for CPS Architectures, William Emfinger and Pranav Kumar (Vanderbilt University)
URL: http://cps-vo.org/node/18517
The RCPS Testbed consists of embedded system hardware with hosts running actual code; a physical system simulator; code running on the hosts that communicates with the physics simulator to get current sensor state and to control the actuators; a smart network switch that allows emulation of network resources

to accurately emulate the system's network; integrated analysis and measurement tools and modeling tools; code generators; and deployment/monitoring utilities. The demonstration showed many of these features using a simulated GPS satellite constellation.

Science of Adversarial Risk in CPS, Yevgeniy Vorobeychik (Vanderbilt University) URL: http://cps-vo.org/node/18479
CPS security relies on many individual decision makers making good choices. Risk stems from choices, which are optimal for individuals, but not for the system as a whole, but in most real CPS security, the system involves multiple defenders, with each defender "charged" with security for a subset of assets. When security decisions are decentralized and decision makers have different interests, system-level security can be sub-optimal. Next steps will be to use simulation as a "multi-defender" platform to form a bridge into the evaluation testbed and to develop automated methods for CPS model-based risk analysis in GME using the attack description language.

Incentive Mechanisms for CPS Security, Saurabh Amin (MIT)
URL: http://cps-vo.org/node/18489
Incentive mechanisms are needed to encourage the building of secure systems. For certain regulatory regimes, electricity distributors make sub-optimal investment in monitoring; user steals less when fines are higher or detection probability is higher. Distributor invests more in monitoring when costs of monitoring lower or user stealing is higher. Due to information deficiencies, R and S are interdependent. Equilibrium depends on relative frequencies of failures and reliability failure distribution. Defenders should co-design defenses against faults and attacks. Contributions of the work are a network game with interdependent reliability and security, full characterization of equilibria, and a polynomial-time algorithm for enumerating all equilibria. Future work will be to study defender interactions with multiple strategic attackers, game parameters not known to all players, link capacities, and edge reinforcement.

Vanderbilt graduate students Pranav Srinivas Kumar (L) and William Emfinger demonstrated the Resilient Cyber Physical Systems testbed.

Putting Humans in the Loop: Active Learning at Scale for Malware Detection, Anthony Joseph (UC Berkeley)
URL: http://cps-vo.org/node/18489
This study looks at use of Machine Learning to separate positive (malicious) from negative (benign) instances. Security Analytics: Using Robust ML for adversary resistant security, metrics and analytics; Pattern mining and prediction, at scale, on big data, with adversaries; Detecting and classifying malicious actions within Cyber-Physical Systems, malware, spam. Situational Awareness: Helping the humans-in-the-loop; Real-time, Machine Learning-based analytics for human domain experts; Interaction with multiple thrusts; Hierarchical Coordination and Control via a ML pipeline addressing CPS security needs for Resilient Monitoring and Control and Evaluation and experimentation using humans and real-world data (malware).

Modeling Privacy in Human CPS, Roy Dong (UC Berkeley)
URL: http://cps-vo.org/node/18486
From an engineering perspective, there are two dominant paradigms: control over information and secrecy. The author proposes privacy contracts since privacy is a good: higher privacy settings could cost more. There is asymmetric information in this problem, and adverse selection becomes an issue.

Secure Computation in Actor Networks, Dusko Pavlovic (U of Hawai'I)
URL: http://cps-vo.org/node/18512
Security is both a suitable subject for science and the process of security is also similar to the process of science, since both science and security depend on the methods of inductive inference. A scientific theory can never be definitely proved, but can only be disproved by new evidence, and improved into a better theory. Because of the same dependency, every security claim and method has a lifetime, and always eventually needs to be improved.

Resilient Sensor Network Design for Flow Networks, Waseem Abbas (Vanderbilt University)
URL: http://cps-vo.org/node/18515
Leakages and faults in flow networks cause commercial and physical losses. Using water supply information, they systematically examine early detection and localization mechanisms of reported and unreported breaks in an efficient way. Resilience issues include uncertainty in system response to burst pipes, inherent model uncertainty, transient system analysis, additional uncertainty in infrastructure topology and characteristics, underground infrastructure that is not visible and hard to access, and the spatial distribution of the networks and complex looped topology due to constant expansion and rehabilitation. This approach considers pipe burst events as opposed to previously majority of work considering water

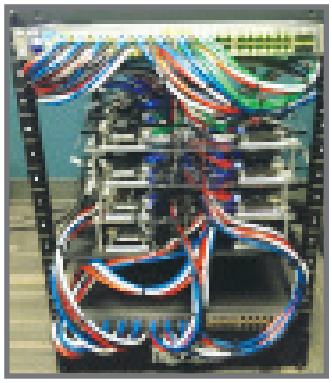

The Resilient Cyber Physical System Testbed hardware component

quality. There is very limited work on localization as compared to detection, and issue for resiliency.

Attack-Resilient Observation Selection, Aron Laszka (Vanderbilt University)
URL: http://cps-vo.org/node/18478
To dynamically control any system, accurate information about its evolving state systems to be monitored can extend over a vast area resulting in many possible points of observation. Focused on traffic patterns, this study posits that the resilience of monitoring to denial-of-service type attacks can be achieved by placing sensors in a resilient way. Resilient sensor placement is formulated as a constrained optimization problem based on a formal prediction model that is applicable to multiple domains. Previous work focused on observation selection while current work is addressing resilient observation selection. Future work will address unit costs of uncertainty for both the "no-attack" case and the "attacked" and selections minimizing the sum cost of both uncertainties.

Using Machine Learning to Improve the Resilience of Control, Claire Tomlin (UC Berkeley)
URL: http://cps-vo.org/node/18482
Using data from air traffic control, the authors use machine learning as a tool to visualize a model of resiliency. They conclude research in the security of control systems has assumed a fixed control algorithm, and considered attack of the sensors, algorithm, machine learning adapts the control based on data collected. In theory, the learning could be used to detect anomalies and intrusions. However, if an attacker knew the learning algorithm, it would be easier to spoof the system without detection

Resilient and Secure Component-Based Software for CPS Architectures, Gabor Karsai (Vanderbilt University)
URL: http://cps-vo.org/node/18511
The "CPS Cloud" is used as an open sensing/Computing/ Actuation Platform where various customer applications can run side-by-side. The physical world can be simulated in real-time with the desired degree of fidelity, including faults, the network can be emulated in real-time with a desired degree of fidelity, including cyber effects, and embedded computing platforms are very affordable. Some examples of potential CPS Cloud subjects include fractionated satellite—observation platforms, coordinated swarm of UAVs executing a mission, fleet of UUVs collecting data while in motion, and monitoring and control nodes on the Smart Grid. Challenges in building this CPS Cloud include networked, distributed control systems, fault-and security resilience, and applications with different trust and security levels that must run side-by-side.

System-Level Co-design for CPS Security, Janos Sztipanovits (Vanderbilt University) URL: http://cps-vo.org/node/18525
The traditional system-level synthesis problem for the "cyber" side of CPS is to dDerive specification for the behavior of the system components that will be implemented using networked computing, derive a functional model for the information architecture and componentize the system, select computing/ networking platform, derive deployment model assigning components of the information architecture to processing and communication platforms, generate code for software components, and perform timing analysis in order to make security part of system-level co-design processes. Mitigation of security vulnerabilities cost performance, timing, and functionality. Integration into design processes will reduce performance degradation.

Science of Security Virtual Organization, Katie Dey (Vanderbilt University)
URL: http://cps-vo.org/node/18528
The Cyber Physical Systems Virtual Organization is a tool to develop community, collaborate, and support technology transfer and translational research. The CPS-VO web page is the focal point for information sharing and community outreach and development. Nodes provide information about SURE activities, meetings, and research as well as general announcements about upcoming events, funding opportunities, discussion forums and chat, and a newsletter containing current research bibliographies about topics of interest to the Science of Security community.

Section 2
Promote Rigorous Scientific Principles

Annual Best Cybersecurity Paper Competition

The Best Scientific Cybersecurity Paper Competition is sponsored yearly by NSA's Research Directorate and reflects the Agency's desire to increase scientific rigor in the cybersecurity field. This competition was established to recognize current research that exemplifies the development of scientific rigor in cybersecurity research. Science of Security (SoS) is a broad enterprise, involving both theoretical and empirical work across a diverse set of topics. While there can only be one best paper, no single paper can span the full breadth of SoS topics. Nevertheless, work in all facets of security science is both needed and encouraged.

The third NSA Competition for Best Scientific Cybersecurity Paper recognized the best scientific cybersecurity paper published in 2014. Papers were nominated between December 1, 2014 through March 31, 2015 and 50 nominations were received. Three papers were selected for recognition, a winning paper and two papers for an honorable mention.

The 3rd annual competition winner, "Additive and Multiplicative Notions of Leakage and Their Capacities," is a research paper presented at the 2014 IEEE Computer Security Foundations Symposium written by Prof. Mario S. Alvim, Dr. Kostas Chatzikokolakis, Prof. Annabelle McIver, Prof. Carroll Morgan, Dr. Catuscia Palamidessi and Prof. Geoffrey Smith. This international team's research focused on information flows and theory, and proposed leakage measures to set bounds on the amount of information a vulnerability can divulge. They mathematically proved their measures were robust as they were resistant to limited knowledge of operation conditions and of attackers cost benefit calculations. In doing so they advanced information flow theory and Shannon capacity and g-leakage. This paper was selected for the award as the research shows strong scientific work and provides needed foundations for information flow and cyberscurity. Their theories can be applied to a wide range of applications such as helping to evaluate vulnerabilities for gauging the safeness of an application or to prioritize vulnerability remediation.

The first paper receiving an honorable mention, "Increasing Security Sensitivity with Social Proof: A Large-Scale Experimental Confirmation," was written by Sauvik Das, Dr. Adam D.I. Kramer, Prof. Laura Dabbish and Prof. Jason I. Hong. Their paper was presented at the 2014 ACM Computer and Communication Security Conference. They examined ways to motivate individuals to adopt security features by showing information about their friends' use of the security features. Particularly notable was the scale of this study, 50,000 people were studied, which is on a much larger scale than traditional human behavior studies. The work also showed scientific merit, analysis, and the paper clearly documents the study, results, and motivation of both the study and statistical approaches employed.

The second paper receiving an honorable mention, "Quantitative Evaluation of Dynamic Platform Techniques as a Defensive Mechanism," was written by Dr. Hamed Okhravi, Dr. James Riordan, and Dr. Kevin Cater and presented at the 17th International Symposium on Research in Attacks, Intrusions and Defenses. Their research studied the effectiveness of dynamic platforms where programs and computers are often changed as a way to prevent intrusions. They built an experimental testbed for evaluation and also simulated the dynamics. The two approaches led to similar results. The paper was selected as it provided scientific analysis of the dynamic platform approach to quantifiably measure increased resistance to compromise. The approach utilized in the paper is able to be applied as a way to evaluate effectiveness of dynamic platforms, which will help decision making and design choices.

NSA Competition Leads

- Dr. Deborah Frincke - Director of Research, NSA
- Dr. Adam Tagert - Science of Security, NSA Trusted Systems Research Group

Distinguished Expert Reviewers

- Dr. Whitfield Diffie - Cybersecurity Advisor
- Dr. Daniel Earl Geer, Jr, Sc.D. - Chief Information Security Officer at In-Q-Tel
- Dr. John D. McLean - Superintendent of the Naval Research Laboratory's Information Technology Division (ITD)
- Professor M. Angela Sasse - Professor of Human-Centered Technology and Head of Information Security Research in the Department of Computer Science at University College London (UCL), UK
- Professor Fred B. Schneider - Samuel B. Eckert Professor of Computer Science at Cornell University
- Mr. Phil Venables - Chief Information Risk Officer at Goldman Sachs
- Professor David A. Wagner - Professor in the Computer Science Division at the University of California, Berkeley
- Jeannette Wing - Vice President, head of Microsoft Research International

INTEL International Science and Engineering Fair (ISEF)

The 2015 National Security Research Directorate Science of Security Initiative Science Fair Award was presented to four high school students for their work on three research projects at the Intel International Science and Engineering Fair (ISEF) on May 14, 2015. The NSA Research Award at ISEF recognizes and encourages outstanding scientific accomplishments in cybersecurity. Criteria for selection include impact and generalization of results, novel aspects of project, quality of science communication, and a project reflective of scientific principles.

Steven Katz of NSA congratulates Elliot Gorokhovsky, the ISEF First Place Winner.

First Place ($3,000)**: A Novel Algorithm for #SAT Using Inclusion-Exclusion Principle and Memorization** by Elliot Gorokhovsky, 16, from Fairview High School in Boulder, Colorado.

Project Description: #SAT is a generalization of an important computer science problem of counting how many conditions satisfy a set of Boolean functions. One use of #SAT solvers is to verify the security properties of a cryptographic algorithm. Elliot developed a novel approach for the counting models of the conjunctive normal form (CNF) formulas that uses memorization (the principle of caching the results of function calls so that they may be reused) to exploit the order of the clauses in this formula and the structure they create. This makes possible the optimization of #SAT by focusing on clustering that maximizes the potential for memorization in the algorithm. His algorithm presents a new lens through which to view #SAT as an independent algorithm and as a complement to the depth-first search (DPLL) algorithm.

Honorable Mention ($1,000): **Capacity Limits of Working Memory: The Impact of Multitasking on Cognitive Control in Digital Natives and Digital Immigrants** by Sarayu Caulfield, 17 and Alexandra Ulmer, 18, from Oregon Episcopal School in Portland, Oregon

Project Description: This is a study of people's ability to process information, cognitive control, in respect to age and multitasking. Sarayu and Alexandra performed a human behavior study where they tested people on how well they could task switch and task filter controlling for multitasking activities, such as checking for an email. They found that high media multitaskers (most adolescents) have a lower ability to filter irrelevant information (task filter) and are less able to limit representation of irrelevant tasks (task switch). It suggests that adults more often could focus on a task while adolescents focused on distracting stimuli.

Honorable Mention ($1,000): **Development of an Authorship Identification Algorithm for Twitter Using Stylometric Techniques** by Cherry Ying Zou, 16, from Poolesville High School in Poolesville, Maryland

Project Description: This project worked on improving algorithms to detect authorship of Twitter posts. Traditional authorship algorithms require longer lengths of text to get accuracy but Twitter tweets are 140 characters or less. Cherry used different bigram (two letter combinations) as tokens in a modified Naïve Bayes classifier. She tested her approach using celebrity tweets as her dataset. Cherry's approach increased accuracy by around 25%. This approach can be applied to anonymous cyber crimes on Twitter.

Section 3

GROW THE SCIENCE OF SECURITY COMMUNITY

HotSoS Overview

HotSoS is the premier Science of Security community event. Co-sponsored by the Association for Computer Machinery (ACM), HotSoS aims to address the fundamental problems of security in a principled manner. The annual event brings together the diverse research community—academia, government, and industry. Participants strive for a comprehensive and methodical approach to identifying and removing threats to security.

The 2015 Symposium and Bootcamp on the Science of Security (HotSoS) was held April 21-22 at the University of Illinois at Urbana-Champaign National Center for Supercomputing Applications. This third annual conference brought together researchers from numerous disciplines seeking a methodical, rigorous scientific approach to identifying and removing cyber threats. Part of the Science of Security effort, the HotSoS goal is to understand how computing systems are designed, built, used, and maintained with an understanding of their security issues and challenges. It seeks not only to put scientific rigor into research, but also to identify the scientific value and underpinnings of cybersecurity.

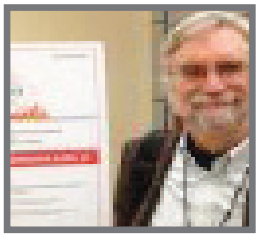

David Nicol, Director of the Illinois Trust Institute and co-PI for the Illinois Science of Security Lablet, was conference chair. Introducing the event, he called for participants to interact and share ideas, thoughts, and questions about the nature of security and the nascent science that is emerging. Kathy Bogner, Intelligence Community Coordinator for Cybersecurity Research, represented the NSA sponsor and welcomed the group, noting the government's long-term interest and commitment to their work. She challenged them to continue to address cybersecurity using strong scientific principles and methods and to share the fruits of that work. She cited the numbers of universities and individual collaborators engaged in Science of Security research as an indication of activity and growth in the field.

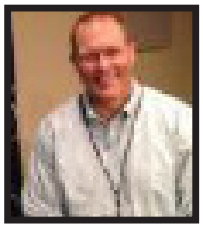

Michael Reiter, Lawrence M. Slifkin Distinguished Professor of Computer Science, University of North Carolina, delivered the keynote **"Is it Science or Engineering? A Sampling of Recent Research."** He said interest in a "Science of Security" is confusing to many researchers, in part due to a lack of clarity about what this "science" should be like and how it should differ from principled engineering. To help clarify the distinction, he described recent research projects about large-scale measurement, attack development, human-centric design, network defense, and provable cryptography to assess which ones, if any, constitute "science." A lively debate ensued.

Jonathan Spring, Researcher and Analyst for the CERT Division, Software Engineering Institute, Carnegie Mellon University, spoke on **"Avoiding Pseudoscience in the Science of Security."** In his view, we seek the philosophical underpinnings to science of security in an effort to avoid pseudoscience. We look at the philosophy of science to describe how "observation and reasoning from results" differs between computing and other sciences due to the engineered elements under study. He demonstrated the challenges in avoiding pseudoscience and some solutions with a case study of malware analysis.

Patrick McDaniel, Professor of Computer Science and Co-Director of the Systems and Internet Infrastructure Security Laboratory, Penn State University, addressed **"The Importance of Measurement and Decision Making to a Science of Security."** A "science" is based on a reasoned modification to a system or environment in response to a functional, performance, or security need. His talk highlighted

activities surrounding the Cyber-Security Collaborative Research Alliance, five universities working in collaboration with the Army Research Lab. Tutorials and a workshop were conducted with concurrent paper presentations.

Five tutorials covered social network analysis; human behavior; policy-governed secure collaboration, security-metrics-driven evaluation, design, development and deployment; and resilient architectures. The workshop focused on analyzing papers from the security literature to determine how completely authors describe their research methods.

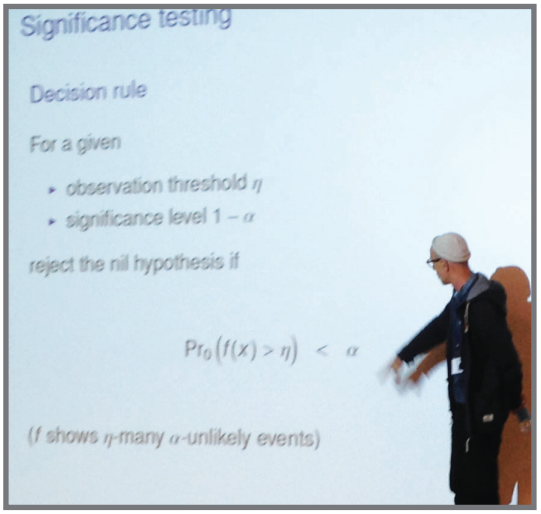

Dusko Pavlovic, a Professor in the Information and Computer Sciences Department, University of Hawai'i. presents his paper, **"Towards a Science of Trust."** The presentation was both animated and stimulating. He explored the idea that security is not just a suitable subject for science, but that the process of security is also similar to the process of science. His take away from the session: "Use methods of Science in the practice of Security."

Thirteen researchers from the United Kingdom and the United States presented individual papers on studies about signal intelligence analyst tasks, detecting abnormal user behavior, tracing cyber-attack analysis processes, vulnerability prediction models, preemptive intrusion detection, enabling forensics, global malware encounters, workflow resiliency, sanctions, password policies, resource-bounded systems integrity assurance, active cyber defense and science of trust.

Allaire Welk, NC State Universtiy addresses methods of learning for signals intelligence analysts.

The 2013 Best Scientific Cybersecurity Paper was an invited paper. Chang Liu, University of Maryland presented **"Memory Trace: Oblivious Program Execution for Cloud Computing."** Chang Liu is a PhD student in the Department of Computer Science at the University of Maryland. His research interest lies at the intersection of security and programming languages. Currently, he is visiting the University of California, Berkeley and working in the Computer Science Department.

Ignacio X. Dominquez, NC State University, listens to a question about his work on input device analytics.

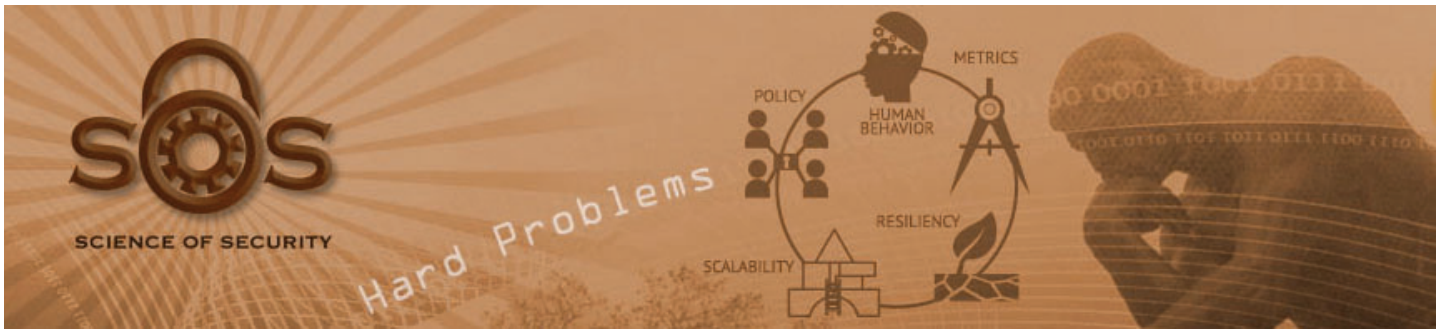

RESEARCH PAPER SESSIONS

Papers presented during the research paper sessions covered a range of scientific issues related to the 5 Hard Problems of cybersecurity: Scalability and Composability; Policy-Governed Secure Collaboration; Security-Metrics-Driven Evaluation, Design, Development, and Deployment; Resilient Architectures; and Understanding and Accounting for Human Behavior. The individual presentations are described below. They will be published in an upcoming ACM conference publication.

Integrity Assurance in Resource-Bounded Systems through Stochastic Message Authentication

Aron Laszka, Yevgeniy Vorobeychik, and Xenofon Koutsoukos

Assuring communication integrity is a central problem in security. The presenters propose a formal game-theoretic framework for optimal stochastic message authentication, providing provable integrity guarantees for resource-bounded systems based on an existing MAC scheme. They use this framework to investigate attacker deterrence, optimal design of stochastic message authentication schemes, and provide experimental results on the computational performance of their framework in practice.

Active Cyber Defense Dynamics Exhibiting Rich Phenomena

Ren Zheng, Wenlian Lu, and Shouhuai Xu

The authors explore the rich phenomena that can be exhibited when the defender employs active defense to combat cyber attacks. This study shows that active cyber defense dynamics (or more generally, cybersecurity dynamics) can exhibit bifurcation and chaos phenomena that have implications for cyber security measurement and prediction. First, that it is infeasible (or even impossible) to accurately measure and predict cyber security under certain circumstances, and second, that the defender must manipulate the dynamics to avoid unmanageable situations in real-life defense operations.

Towards a Science of Trust

Dusko Pavlovic

This paper explores the idea that security is not just a suitable subject for science, but that the process of security is also similar to the process of science. This similarity arises from the fact that both science and security depend on the methods of inductive inference. Because of this dependency, a scientific theory can never be definitely proved, but can only be disproved by new evidence and improved into a better theory. Because of the same dependency, every security claim and method has a lifetime, and always eventually needs to be improved.

Challenges with Applying Vulnerability Prediction Models

Patrick Morrison, Kim Herzig, Brendan Murphy, and Laurie Williams

The authors address vulnerability prediction models (VPM) as a basis for software engineers to prioritize precious verification resources to search for vulnerabilities. The goal of this research is to measure whether vulnerability prediction models built using standard recommendations perform well enough to provide actionable results for engineering resource allocation. They define "actionable" in terms of the inspection effort required to evaluate model results. They conclude VPMs must be refined to achieve actionable performance, possibly through security-specific metrics.

Preemptive Intrusion Detection: Theoretical Framework and Real-World Measurements

Phuong Cao, Eric Badger, Zbigniew Kalbarczyk, Ravishankar Iyer, and Adam Slagell

This paper presents a framework for highly accurate and preemptive detection of attacks, i.e., before system misuse. Using security logs on real incidents that occurred over a six-year period at the National Center for Supercomputing Applications (NCSA),

the authors evaluated their framework. The data consisted of security incidents that were only identified after the fact by security analysts. The framework detected 74 percent of attacks, and the majority of them were detected before the system misuse. In addition, six hidden attacks were uncovered that were not detected by intrusion detection systems during the incidents or by security analysts in post-incident forensic analyses.

Enabling Forensics by Proposing Heuristics to Identify Mandatory Log Events

Jason King, Rahul Pandita, and Laurie Williams

Software engineers often implement logging mechanisms to debug software and diagnose faults. These logging mechanisms need to capture detailed traces of user activity to enable forensics and hold users accountable. Techniques for identifying what events to log are often subjective and produce inconsistent results. This study helps software engineers strengthen forensic-ability and user accountability by systematically identifying mandatory log events through processing of unconstrained natural language software artifacts, and then, by proposing empirically-derived heuristics to help determine whether an event must be logged.

Modelling User Availability in Workflow Resiliency Analysis

John C. Mace, Charles Morisset, and Aad van Moorsel

Workflows capture complex operational processes and include security constraints that limit which users can perform which tasks. An improper security policy may prevent certain tasks being assigned and may force a policy violation. Tools are required that allow automatic evaluation of workflow resiliency. Modelling user availability can be done in multiple ways for the same workflow. Finding the correct choice of model is a complex concern with a major impact on the calculated resiliency. The authors describe a number of user availability models and their encoding in the model checker PRISM, used to evaluate resiliency. They also show how model choice can affect resiliency computation in terms of its value, memory, and CPU time.

An Empirical Study of Global Malware Encounters

Ghita Mezzour, Kathleen M. Carley, and L. Richard Carley

The authors empirically test alternative hypotheses about factors behind international variation in the number of trojans, worm, and virus encounters using Symantec Anti-Virus (AV) telemetry data collected from more than 10 million Symantec global customer computers. They used regression analysis to test for the effect of computing and monetary resources, web browsing behavior, computer piracy, cyber security expertise, and international relations on international variation in malware encounters and found that trojans, worms, and viruses are most prevalent in Sub-Saharan African and Asian countries. The main factor that explains the high malware exposure of these countries is widespread computer piracy, especially when combined with poverty.

An Integrated Computer-Aided Cognitive Task Analysis Method for Tracing Cyber-Attack Analysis Processes

Chen Zhong, John Yen, Peng Liu, Rob Erbacher, Renee Etoty, and Christopher Garneau

Cyber-attack analysts are required to process large amounts of network data and to reason under uncertainty to detect cyber-attacks. Capturing and studying the fine-grained analysts' cognitive processes helps researchers gain deep understanding of how they conduct analytical reasoning and elicit their procedure knowledge and experience to further improve their performance. To conduct cognitive task analysis studies in cyber-attack analysis, the authors proposed an integrated computer-aided data collection method for cognitive task analysis (CTA) with three building elements: a trace representation of the fine-grained cyber-attack analysis process, a computer tool supporting process tracing, and a laboratory experiment for collecting traces of analysts' cognitive processes in conducting a cyber-attack analysis task.

All Signals Go: Investigating How Individual Differences Affect Performance on a Medical Diagnosis Task Designed to Parallel a Signals Intelligence Analyst Task

Allaire K. Welk and Christopher B. Mayhorn

Signals intelligence analysts perform complex decision-making tasks that involve gathering, sorting, and analyzing information. This study aimed to evaluate how individual differences influence performance in an Internet search-based medical diagnosis task designed to simulate a signals analyst task. Individual differences included working memory capacity and previous experience with elements of the task, prior experience using the Internet, and prior experience conducting Internet searches. Results indicated that working memory significantly predicted performance on this medical diagnosis task and other factors were not significant predictors of performance. These results provide additional evidence that working memory capacity greatly influences performance on cognitively complex decision-making tasks, whereas experience with elements of the task may not. These findings suggest that working memory capacity should be considered when screening individuals for signals intelligence analyst positions.

Detecting Abnormal User Behavior Through Pattern-Mining Input Device Analytics

Ignacio X. Domínguez, Alok Goel, David L. Roberts, and Robert St. Amant

This paper presents a method for detecting patterns in the usage of a computer mouse that can give insights into user's cognitive processes. The authors conducted a study using a computer version of the Memory game (also known as the Concentration game) that allowed some participants to reveal the content of the tiles, expecting their low-level mouse interaction patterns to deviate from those of normal players with no access to this information. They then trained models to detect these differences using task-independent input device features. The models detected cheating with 98.73% accuracy for players who cheated or did not cheat consistently for entire rounds of the game, and with 89.18% accuracy for cases in which players enabled and then disabled cheating within rounds.

Understanding Sanction under Variable Observability in a Secure, Collaborative Environment

Hongying Du, Bennett Narron, Nirav Ajmeri, Emily Berglund, Jon Doyle, and Munindar P. Singh

Many aspects of norm-governance remain poorly understood, inhibiting adoption in real-life collaborative systems. This work focuses on the combined effects of sanction and the observability of the sanctioner in a secure, collaborative environment using a simulation consisting of agents maintaining "compliance" to enforced security norms while remaining "motivated" as researchers. They tested whether delayed observability of the environment would lead to greater motivation of agents to complete research tasks than immediate observability, and if sanctioning a group for a violation would lead to greater compliance to security norms than sanctioning an individual. They found that only the latter hypothesis is supported.

Measuring the Security Impacts of Password Policies Using Cognitive Behavioral Agent-Based Modeling

Vijay Kothari, Jim Blythe, Sean W. Smith, and Ross Koppel

Agent-based modeling can serve as a valuable asset to security personnel who wish to better understand the security landscape within their organization, especially as it relates to user behavior and circumvention. The authors argue in favor of cognitive behavioral agent-based modeling for usable security, report on their work on developing an agent-based model for a password management scenario, perform a sensitivity analysis, which provides them with valuable insights into improving security, and provides directions for future work.

TUTORIALS

Tutorial 1:

Social Network Analysis for Science of Security

Kathleen Carley, Carnegie Mellon University

The tutorial provided a brief introduction to the area of network science, covering analytics and visualization. Dr. Carley described the core ideas, most common metrics, critical theories, and an overview of key tools. She drew illustrative examples from three security-related issues: insider threat analysis, resilient organizational designs, and global cyber-security attacks.

Tutorial 2:

Understanding and Accounting for Human Behavior

Sean Smith, Dartmouth College and Jim Blythe, University of Southern California

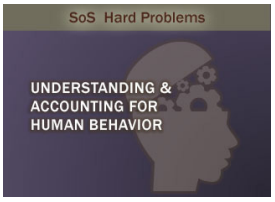

Since computers are machines, it's tempting to think of computer security as purely a technical problem. However, computing systems are created, used, and maintained by humans and exist to serve the goals of human and institutional stakeholders. Consequently, effectively addressing the security problem requires understanding this human dimension. The presenters discussed this challenge and the principal research approaches to it.

Tutorial 3:

Policy-Governed Secure Collaboration

Munindar Singh, North Carolina State University

The envisioned Science of Security can be understood as a systemic body of knowledge

with theoretical and empirical underpinnings that inform the engineering of secure information systems. The presentation addressed the underpinnings pertaining to the hard problem of secure collaboration, approaching cybersecurity from a sociotechnical perspective, and understanding systems through the interplay of human behavior with technical architecture on the one hand and social architecture on the other. The presentation emphasized the social architecture and modeled it in terms of a formalization based on organizations and normative relationships. Dr. Singh described how norms provide a basis for specifying security requirements at a high level, a basis for accountability, and a semantic basis for trust. He concluded the presentation by providing some directions and challenges for future research, including formalization and empirical study.

Tutorial 4:

Security-Metrics-Driven Evaluation, Design, Development and Deployment
William Sanders, University of Illinois at Urbana-Champaign

Making sound security decisions when designing, operating, and maintaining a complex system is a challenging task. Analysts need to be able to understand and predict how different factors affect overall system security. During system design, security analysts want to compare the security of multiple proposed system architectures. After a system is deployed, analysts want to determine where security enhancement should be focused by examining how the system is most likely to be successfully penetrated. Additionally, when several security enhancement options are being considered, analysts would like to evaluate the relative merit of each. In each of these scenarios, quantitative security metrics should provide insight on system security and aid security decisions. The tutorial provided a survey of existing quantitative security evaluation techniques and described new work being done at the University of Illinois at Urbana-Champaign in this field.

Tutorial 5:

Resilient Architectures
Ravishankar Iyer, University of Illinois at Urbana-Champaign

Resilience brings together experts in security, fault tolerance, human factors, and high integrity computing for the design and validation of systems that are expected to continue to deliver critical services in the event of attacks and failures. The tutorial highlighted issues and challenges in designing systems that are resilient to both malicious attacks and accidental failures, provided both cyber and cyber-physical examples, and concluded by addressing the challenges and opportunities from both a theoretical and practical perspective.

SCIENCE OF SECURITY VIRTUAL ORGANIZATION

The Science of Security Virtual Organization (**SoS-VO**) promotes collaborative research with a web-based clearinghouse for sharing research, publications, events, funding opportunities, collaboration, and news about cybersecurity science. Researchers, students, educators, and industry partners can obtain the latest information about the growing field of Science of Security and participate in a range of collaborative activities, as well as learn about recent, ongoing, and planned research. Work from the four Lablets and the SURE consortium are posted regularly, as are upcoming events, news items, and relevant publications. There are also active forums, discussion groups, and blogs. Over the past year, membership in the SoS-VO has grown from approximately 150 members to over 900.

For more information scan this link or visit
http://sos-vo.org

SCIENCE OF SECURITY NEWSLETTERS

The **Science of Security Newsletter** showcases research programs of interest to the Science of Security (SoS) Community. All of the newsletter's content is based on open sources and in many cases links to the original work or to the web page for a specific program. The newsletters published have included research papers, news articles of interest, SoS lablet activities, upcoming events, conference reports, and other items of interest to the SoS Community. There have been 14 newsletters published containing almost 6,500 unique articles and nearly 400 bibliographies with an average of 15-20 papers each. Based on newsletter articles, we have identified almost 700 keywords and phrases associated with the hard problems and related topics under multiple headings. Over 87 related conferences, both inside and outside the US, were curated in the newsletters.

Word cloud of key terms found in titles of "Publications of Interest" topics.

Source: Permission granted by WordItOut at http://worditout.com for use of the generated word cloud.

OUTREACH ACTIVITIES

The Science of Security initiative continues to grow the SoS community, as part of our mission statement. SoS continues to gain partnerships and momentum in research. A large part of the SoS program team engage with the academic community researchers involved in the 4 SoS lablets, 26 sub lablets and 53 collaborating institutions. With over 100+ institutions engaged in the Science of Security efforts, including several industry partners. Partnerships have formed with Army Research Office (ARO), Army Research Lab (ARL), U.S. Naval Academy, Air Force Research Lab (AFRL), increasing the need for foundational approaches to secure and steering the way forward for an increased community awareness. However, the Science of Security initiative has grown much farther in collaboration, reaching many scientific researchers internationally, with the Science of Security Virtual Organization (SoS-VO). Over 2,000 published researchers have been cited on the SoS-VO. Science of Security lablet quarterly meetings, the annual HotSoS presentations, and Annual Best Scientific Cybersecurity Paper Competition is also captured on the SoS-VO. Several research projects were featured in the 2015 cybersecurity edition of the Next Wave magazine. In FY16, Science of Security will continue to expand the outreach and partnerships with academia and industry partners.

produced by Cyber Pack Ventures, Inc.

SCIENCE OF SECURITY